

# **MySQL 3.23, 4.0, 4.1 Reference Manual**

---

# MySQL 3.23, 4.0, 4.1 Reference Manual

## Abstract

This is the MySQL Reference Manual. It documents MySQL 3.23 through MySQL 4.1.22.

Document generated on: 2006-10-01 (revision: 3517)

Copyright 1997-2006 MySQL AB

This documentation is NOT distributed under a GPL license. Use of this documentation is subject to the following terms: You may create a printed copy of this documentation solely for your own personal use. Conversion to other formats is allowed as long as the actual content is not altered or edited in any way. You shall not publish or distribute this documentation in any form or on any media, except if you distribute the documentation in a manner similar to how MySQL disseminates it (that is, electronically for download on a website with the software) or on a CD-ROM or similar medium, provided however that the documentation is disseminated together with the software on the same medium. Any other use, such as any dissemination of printed copies or use of this documentation, in whole or in part, in another publication, requires the prior written consent from an authorized representative of MySQL AB. MySQL AB reserves any and all rights to this documentation not expressly granted above.

Please email [<docs@mysql.com>](mailto:docs@mysql.com) for more information.

---

---

---

---

# Table of Contents

Preface .....	xxiv
1. General Information .....	1
1.1. About This Manual .....	1
1.2. Conventions Used in This Manual .....	2
1.3. Overview of MySQL AB .....	4
1.4. Overview of the MySQL Database Management System .....	5
1.4.1. History of MySQL .....	6
1.4.2. The Main Features of MySQL .....	6
1.4.3. MySQL Stability .....	9
1.4.4. How Large MySQL Tables Can Be .....	10
1.4.5. Year 2000 Compliance .....	11
1.5. Overview of the MaxDB Database Management System .....	12
1.5.1. What is MaxDB? .....	12
1.5.2. History of MaxDB .....	13
1.5.3. Features of MaxDB .....	13
1.5.4. Licensing and Support .....	14
1.5.5. Feature Differences Between MaxDB and MySQL .....	14
1.5.6. Interoperability Features Between MaxDB and MySQL .....	15
1.5.7. MaxDB-Related Links .....	15
1.6. MySQL Development Roadmap .....	15
1.6.1. MySQL 4.0 in a Nutshell .....	16
1.6.2. MySQL 4.1 in a Nutshell .....	18
1.6.3. What's New in MySQL 5.0 .....	19
1.7. MySQL Information Sources .....	20
1.7.1. MySQL Mailing Lists .....	20
1.7.2. MySQL Community Support at the MySQL Forums .....	23
1.7.3. MySQL Community Support on Internet Relay Chat (IRC) .....	23
1.8. How to Report Bugs or Problems .....	24
1.9. MySQL Standards Compliance .....	28
1.9.1. What Standards MySQL Follows .....	28
1.9.2. Selecting SQL Modes .....	28
1.9.3. Running MySQL in ANSI Mode .....	29
1.9.4. MySQL Extensions to Standard SQL .....	29
1.9.5. MySQL Differences from Standard SQL .....	32
1.9.6. How MySQL Deals with Constraints .....	38
2. Installing and Upgrading MySQL .....	41
2.1. General Installation Issues .....	41
2.1.1. Operating Systems Supported by MySQL .....	42
2.1.2. Choosing Which MySQL Distribution to Install .....	43
2.1.3. How to Get MySQL .....	54
2.1.4. Verifying Package Integrity Using MD5 Checksums or <a href="#">GnuPG</a> .....	54
2.1.5. Installation Layouts .....	57
2.2. Standard MySQL Installation Using a Binary Distribution .....	58
2.3. Installing MySQL on Windows .....	59
2.3.1. Choosing An Installation Package .....	60
2.3.2. Installing MySQL with the Automated Installer .....	60
2.3.3. Using the MySQL Installation Wizard .....	61
2.3.4. Using the Configuration Wizard .....	64
2.3.5. Installing MySQL from a Noinstall Zip Archive .....	68
2.3.6. Extracting the Install Archive .....	69
2.3.7. Creating an Option File .....	69
2.3.8. Selecting a MySQL Server type .....	70
2.3.9. Starting the Server for the First Time .....	71

---

2.3.10. Starting MySQL from the Windows Command Line .....	73
2.3.11. Starting MySQL as a Windows Service .....	73
2.3.12. Testing The MySQL Installation .....	76
2.3.13. Troubleshooting a MySQL Installation Under Windows .....	76
2.3.14. Upgrading MySQL on Windows .....	78
2.3.15. MySQL on Windows Compared to MySQL on Unix .....	79
2.4. Installing MySQL on Linux .....	81
2.5. Installing MySQL on Mac OS X .....	84
2.6. Installing MySQL on Solaris .....	86
2.7. Installing MySQL on NetWare .....	87
2.8. Installing MySQL on Other Unix-Like Systems .....	89
2.9. MySQL Installation Using a Source Distribution .....	92
2.9.1. Source Installation Overview .....	92
2.9.2. Typical <code>configure</code> Options .....	95
2.9.3. Installing from the Development Source Tree .....	98
2.9.4. Dealing with Problems Compiling MySQL .....	101
2.9.5. MIT-pthreads Notes .....	104
2.9.6. Installing MySQL from Source on Windows .....	105
2.9.7. Compiling MySQL Clients on Windows .....	109
2.10. Post-Installation Setup and Testing .....	109
2.10.1. Windows Post-Installation Procedures .....	110
2.10.2. Unix Post-Installation Procedures .....	111
2.10.3. Securing the Initial MySQL Accounts .....	121
2.11. Upgrading MySQL .....	124
2.11.1. Upgrading from MySQL 4.0 to 4.1 .....	125
2.11.2. Upgrading from MySQL 3.23 to 4.0 .....	132
2.11.3. Copying MySQL Databases to Another Machine .....	137
2.12. Downgrading MySQL .....	138
2.12.1. Downgrading to MySQL 4.0 .....	138
2.13. Operating System-Specific Notes .....	139
2.13.1. Linux Notes .....	139
2.13.2. Mac OS X Notes .....	146
2.13.3. Solaris Notes .....	147
2.13.4. BSD Notes .....	151
2.13.5. Other Unix Notes .....	154
2.13.6. OS/2 Notes .....	170
2.14. Perl Installation Notes .....	170
2.14.1. Installing Perl on Unix .....	171
2.14.2. Installing ActiveState Perl on Windows .....	172
2.14.3. Problems Using the Perl <code>DBI/DBD</code> Interface .....	172
3. Tutorial .....	175
3.1. Connecting to and Disconnecting from the Server .....	175
3.2. Entering Queries .....	176
3.3. Creating and Using a Database .....	179
3.3.1. Creating and Selecting a Database .....	180
3.3.2. Creating a Table .....	181
3.3.3. Loading Data into a Table .....	182
3.3.4. Retrieving Information from a Table .....	183
3.4. Getting Information About Databases and Tables .....	196
3.5. Using <code>mysql</code> in Batch Mode .....	196
3.6. Examples of Common Queries .....	198
3.6.1. The Maximum Value for a Column .....	198
3.6.2. The Row Holding the Maximum of a Certain Column .....	199
3.6.3. Maximum of Column per Group .....	199
3.6.4. The Rows Holding the Group-wise Maximum of a Certain Field .....	200
3.6.5. Using User-Defined Variables .....	201
3.6.6. Using Foreign Keys .....	201
3.6.7. Searching on Two Keys .....	202

---

---

3.6.8. Calculating Visits Per Day .....	203
3.6.9. Using <code>AUTO_INCREMENT</code> .....	203
3.7. Queries from the Twin Project .....	205
3.7.1. Find All Non-distributed Twins .....	205
3.7.2. Show a Table of Twin Pair Status .....	207
3.8. Using MySQL with Apache .....	208
4. Using MySQL Programs .....	209
4.1. Overview of MySQL Programs .....	209
4.2. Invoking MySQL Programs .....	210
4.3. Specifying Program Options .....	211
4.3.1. Using Options on the Command Line .....	212
4.3.2. Using Option Files .....	214
4.3.3. Using Environment Variables to Specify Options .....	218
4.3.4. Using Options to Set Program Variables .....	219
5. Database Administration .....	221
5.1. Overview of Server-Side Programs .....	221
5.2. <code>mysqld</code> — The MySQL Server .....	222
5.2.1. Command Options .....	222
5.2.2. System Variables .....	234
5.2.3. Using System Variables .....	257
5.2.4. Status Variables .....	266
5.2.5. SQL Modes .....	273
5.2.6. The Shutdown Process .....	276
5.2.7. Server-Side Help .....	277
5.3. The <code>mysqld-max</code> Extended MySQL Server .....	277
5.4. MySQL Server Startup Programs .....	280
5.4.1. <code>mysqld_safe</code> — MySQL Server Startup Script .....	280
5.4.2. <code>mysql.server</code> — MySQL Server Startup Script .....	284
5.4.3. <code>mysqld_multi</code> — Manage Multiple MySQL Servers .....	284
5.5. Installation-Related Programs .....	288
5.5.1. <code>mysql_fix_privilege_tables</code> — Upgrade MySQL System Tables .....	288
5.6. General Security Issues .....	289
5.6.1. General Security Guidelines .....	289
5.6.2. Making MySQL Secure Against Attackers .....	291
5.6.3. Security-Related <code>mysqld</code> Options .....	293
5.6.4. Security Issues with <code>LOAD DATA LOCAL</code> .....	294
5.6.5. How to Run MySQL as a Normal User .....	295
5.7. The MySQL Access Privilege System .....	296
5.7.1. What the Privilege System Does .....	296
5.7.2. How the Privilege System Works .....	296
5.7.3. Privileges Provided by MySQL .....	300
5.7.4. Connecting to the MySQL Server .....	303
5.7.5. Access Control, Stage 1: Connection Verification .....	305
5.7.6. Access Control, Stage 2: Request Verification .....	308
5.7.7. When Privilege Changes Take Effect .....	310
5.7.8. Causes of <code>Access denied</code> Errors .....	311
5.7.9. Password Hashing as of MySQL 4.1 .....	316
5.8. MySQL User Account Management .....	321
5.8.1. MySQL Usernames and Passwords .....	321
5.8.2. Adding New User Accounts to MySQL .....	322
5.8.3. Removing User Accounts from MySQL .....	325
5.8.4. Limiting Account Resources .....	325
5.8.5. Assigning Account Passwords .....	327
5.8.6. Keeping Your Password Secure .....	328
5.8.7. Using Secure Connections .....	329
5.9. Backup and Recovery .....	336
5.9.1. Database Backups .....	336
5.9.2. Example Backup and Recovery Strategy .....	338

---

---

5.9.3. Point-in-Time Recovery .....	341
5.9.4. Table Maintenance and Crash Recovery .....	343
5.10. MySQL Localization and International Usage .....	354
5.10.1. The Character Set Used for Data and Sorting .....	354
5.10.2. Setting the Error Message Language .....	356
5.10.3. Adding a New Character Set .....	356
5.10.4. The Character Definition Arrays .....	358
5.10.5. String Collating Support .....	358
5.10.6. Multi-Byte Character Support .....	358
5.10.7. Problems With Character Sets .....	359
5.10.8. MySQL Server Time Zone Support .....	359
5.11. MySQL Server Logs .....	361
5.11.1. The Error Log .....	361
5.11.2. The General Query Log .....	362
5.11.3. The Update Log .....	362
5.11.4. The Binary Log .....	363
5.11.5. The Slow Query Log .....	367
5.11.6. Server Log Maintenance .....	367
5.12. Running Multiple MySQL Servers on the Same Machine .....	368
5.12.1. Running Multiple Servers on Windows .....	370
5.12.2. Running Multiple Servers on Unix .....	373
5.12.3. Using Client Programs in a Multiple-Server Environment .....	374
5.13. The MySQL Query Cache .....	375
5.13.1. How the Query Cache Operates .....	376
5.13.2. Query Cache <code>SELECT</code> Options .....	377
5.13.3. Query Cache Configuration .....	377
5.13.4. Query Cache Status and Maintenance .....	379
6. Replication .....	381
6.1. Introduction to Replication .....	381
6.2. Replication Implementation Overview .....	381
6.3. Replication Implementation Details .....	382
6.3.1. Replication Master Thread States .....	383
6.3.2. Replication Slave I/O Thread States .....	383
6.3.3. Replication Slave SQL Thread States .....	385
6.3.4. Replication Relay and Status Files .....	385
6.4. How to Set Up Replication .....	387
6.5. Replication Compatibility Between MySQL Versions .....	391
6.6. Upgrading a Replication Setup .....	392
6.6.1. Upgrading Replication to 4.0 or 4.1 .....	392
6.7. Replication Features and Known Problems .....	392
6.8. Replication Startup Options .....	398
6.9. How Servers Evaluate Replication Rules .....	405
6.10. Replication FAQ .....	408
6.11. Troubleshooting Replication .....	413
6.12. How to Report Replication Bugs or Problems .....	414
7. Optimization .....	416
7.1. Optimization Overview .....	416
7.1.1. MySQL Design Limitations and Tradeoffs .....	416
7.1.2. Designing Applications for Portability .....	417
7.1.3. What We Have Used MySQL For .....	418
7.1.4. The MySQL Benchmark Suite .....	418
7.1.5. Using Your Own Benchmarks .....	419
7.2. Optimizing <code>SELECT</code> and Other Statements .....	420
7.2.1. Optimizing Queries with <code>EXPLAIN</code> .....	420
7.2.2. Estimating Query Performance .....	428
7.2.3. Speed of <code>SELECT</code> Queries .....	428
7.2.4. <code>WHERE</code> Clause Optimization .....	429
7.2.5. Range Optimization .....	431

---

---

7.2.6. IS NULL Optimization .....	434
7.2.7. DISTINCT Optimization .....	435
7.2.8. LEFT JOIN and RIGHT JOIN Optimization .....	435
7.2.9. ORDER BY Optimization .....	436
7.2.10. GROUP BY Optimization .....	439
7.2.11. LIMIT Optimization .....	440
7.2.12. How to Avoid Table Scans .....	440
7.2.13. Speed of INSERT Statements .....	441
7.2.14. Speed of UPDATE Statements .....	443
7.2.15. Speed of DELETE Statements .....	443
7.2.16. Other Optimization Tips .....	443
7.3. Locking Issues .....	446
7.3.1. Locking Methods .....	446
7.3.2. Table Locking Issues .....	448
7.3.3. Concurrent Inserts .....	450
7.4. Optimizing Database Structure .....	450
7.4.1. Design Choices .....	450
7.4.2. Make Your Data as Small as Possible .....	450
7.4.3. Column Indexes .....	451
7.4.4. Multiple-Column Indexes .....	452
7.4.5. How MySQL Uses Indexes .....	453
7.4.6. The MyISAM Key Cache .....	455
7.4.7. MyISAM Index Statistics Collection .....	460
7.4.8. How MySQL Opens and Closes Tables .....	462
7.4.9. Drawbacks to Creating Many Tables in the Same Database .....	463
7.5. Optimizing the MySQL Server .....	463
7.5.1. System Factors and Startup Parameter Tuning .....	463
7.5.2. Tuning Server Parameters .....	464
7.5.3. How Compiling and Linking Affects the Speed of MySQL .....	466
7.5.4. How MySQL Uses Memory .....	467
7.5.5. How MySQL Uses DNS .....	469
7.6. Disk Issues .....	469
7.6.1. Using Symbolic Links .....	470
8. Client and Utility Programs .....	474
8.1. Overview of Client and Utility Programs .....	474
8.2. myisam_ftdump — Display Full-Text Index information .....	476
8.3. myisamchk — MyISAM Table-Maintenance Utility .....	477
8.3.1. myisamchk General Options .....	478
8.3.2. myisamchk Check Options .....	480
8.3.3. myisamchk Repair Options .....	481
8.3.4. Other myisamchk Options .....	483
8.3.5. myisamchk Memory Usage .....	484
8.4. myisamlog — Display MyISAM Log File Contents .....	484
8.5. myisampack — Generate Compressed, Read-Only MyISAM Tables .....	486
8.6. mysql — The MySQL Command-Line Tool .....	492
8.6.1. mysql Options .....	492
8.6.2. mysql Commands .....	497
8.6.3. mysql Server-Side Help .....	500
8.6.4. Executing SQL Statements from a Text File .....	501
8.6.5. mysql Tips .....	502
8.7. mysql_explain_log — Use EXPLAIN on Statements in Query Log .....	503
8.8. mysqlaccess — Client for Checking Access Privileges .....	504
8.9. mysqladmin — Client for Administering a MySQL Server .....	506
8.10. mysqlbinlog — Utility for Processing Binary Log Files .....	511
8.11. mysqlcheck — A Table Maintenance and Repair Program .....	516
8.12. mysqldump — A Database Backup Program .....	519
8.13. mysqlhotcopy — A Database Backup Program .....	527
8.14. mysqlimport — A Data Import Program .....	530

---



---

8.15. <code>mysqlshow</code> — Display Database, Table, and Column Information .....	532
8.16. <code>mysql_zap</code> — Kill Processes That Match a Pattern .....	534
8.17. <code>pererror</code> — Explain Error Codes .....	535
8.18. <code>replace</code> — A String-Replacement Utility .....	536
9. Language Structure .....	538
9.1. Literal Values .....	538
9.1.1. Strings .....	538
9.1.2. Numbers .....	540
9.1.3. Hexadecimal Values .....	540
9.1.4. Boolean Values .....	541
9.1.5. <code>NULL</code> Values .....	541
9.2. Database, Table, Index, Column, and Alias Names .....	541
9.2.1. Identifier Qualifiers .....	543
9.2.2. Identifier Case Sensitivity .....	544
9.3. User-Defined Variables .....	545
9.4. Comment Syntax .....	546
9.5. Treatment of Reserved Words in MySQL .....	547
10. Character Set Support .....	551
10.1. Character Sets and Collations in General .....	551
10.2. Character Sets and Collations in MySQL .....	552
10.3. Specifying Character Sets and Collations .....	553
10.3.1. Server Character Set and Collation .....	553
10.3.2. Database Character Set and Collation .....	554
10.3.3. Table Character Set and Collation .....	555
10.3.4. Column Character Set and Collation .....	555
10.3.5. Character String Literal Character Set and Collation .....	556
10.3.6. National Character Set .....	557
10.3.7. Examples of Character Set and Collation Assignment .....	557
10.3.8. Compatibility with Other DBMSs .....	558
10.4. Connection Character Sets and Collations .....	558
10.5. Collation Issues .....	560
10.5.1. Using <code>COLLATE</code> in SQL Statements .....	560
10.5.2. <code>COLLATE</code> Clause Precedence .....	561
10.5.3. <code>BINARY</code> Operator .....	561
10.5.4. Some Special Cases Where the Collation Determination Is Tricky .....	562
10.5.5. Collations Must Be for the Right Character Set .....	563
10.5.6. An Example of the Effect of Collation .....	564
10.6. Operations Affected by Character Set Support .....	565
10.6.1. Result Strings .....	565
10.6.2. <code>CONVERT()</code> and <code>CAST()</code> .....	566
10.6.3. <code>SHOW</code> Statements and <code>INFORMATION_SCHEMA</code> .....	566
10.7. Unicode Support .....	568
10.8. UTF-8 for Metadata .....	568
10.9. Upgrading Character Sets from MySQL 4.0 .....	569
10.9.1. 4.0 Character Sets and Corresponding 4.1 Character Set/Collation Pairs ..	570
10.9.2. Converting 4.0 Character Columns to 4.1 Format .....	571
10.10. Character Sets and Collations That MySQL Supports .....	572
10.10.1. Unicode Character Sets .....	573
10.10.2. West European Character Sets .....	575
10.10.3. Central European Character Sets .....	577
10.10.4. South European and Middle East Character Sets .....	578
10.10.5. Baltic Character Sets .....	578
10.10.6. Cyrillic Character Sets .....	579
10.10.7. Asian Character Sets .....	579
11. Data Types .....	584
11.1. Data Type Overview .....	584
11.1.1. Overview of Numeric Types .....	584
11.1.2. Overview of Date and Time Types .....	587

---

---

11.1.3. Overview of String Types .....	589
11.1.4. Data Type Default Values .....	592
11.2. Numeric Types .....	592
11.3. Date and Time Types .....	595
11.3.1. The <code>DATETIME</code> , <code>DATE</code> , and <code>TIMESTAMP</code> Types .....	596
11.3.2. The <code>TIME</code> Type .....	603
11.3.3. The <code>YEAR</code> Type .....	604
11.3.4. Y2K Issues and Date Types .....	604
11.4. String Types .....	605
11.4.1. The <code>CHAR</code> and <code>VARCHAR</code> Types .....	605
11.4.2. The <code>BINARY</code> and <code>VARBINARY</code> Types .....	606
11.4.3. The <code>BLOB</code> and <code>TEXT</code> Types .....	607
11.4.4. The <code>ENUM</code> Type .....	609
11.4.5. The <code>SET</code> Type .....	611
11.5. Data Type Storage Requirements .....	613
11.6. Choosing the Right Type for a Column .....	615
11.7. Using Data Types from Other Database Engines .....	615
12. Functions and Operators .....	617
12.1. Operators .....	617
12.1.1. Operator Precedence .....	617
12.1.2. Type Conversion in Expression Evaluation .....	618
12.1.3. Comparison Functions and Operators .....	619
12.1.4. Logical Operators .....	623
12.2. Control Flow Functions .....	625
12.3. String Functions .....	627
12.3.1. String Comparison Functions .....	636
12.4. Numeric Functions .....	638
12.4.1. Arithmetic Operators .....	638
12.4.2. Mathematical Functions .....	640
12.5. Date and Time Functions .....	646
12.6. What Calendar Is Used By MySQL? .....	661
12.7. Full-Text Search Functions .....	662
12.7.1. Boolean Full-Text Searches .....	665
12.7.2. Full-Text Searches with Query Expansion .....	667
12.7.3. Full-Text Stopwords .....	668
12.7.4. Full-Text Restrictions .....	671
12.7.5. Fine-Tuning MySQL Full-Text Search .....	671
12.8. Cast Functions and Operators .....	673
12.9. Other Functions .....	676
12.9.1. Bit Functions .....	676
12.9.2. Encryption and Compression Functions .....	677
12.9.3. Information Functions .....	681
12.9.4. Miscellaneous Functions .....	687
12.10. Functions and Modifiers for Use with <code>GROUP BY</code> Clauses .....	690
12.10.1. <code>GROUP BY</code> (Aggregate) Functions .....	690
12.10.2. <code>GROUP BY</code> Modifiers .....	693
12.10.3. <code>GROUP BY</code> and <code>HAVING</code> with Hidden Fields .....	696
13. SQL Statement Syntax .....	698
13.1. Data Definition Statements .....	698
13.1.1. <code>ALTER DATABASE</code> Syntax .....	698
13.1.2. <code>ALTER TABLE</code> Syntax .....	698
13.1.3. <code>CREATE DATABASE</code> Syntax .....	704
13.1.4. <code>CREATE INDEX</code> Syntax .....	704
13.1.5. <code>CREATE TABLE</code> Syntax .....	706
13.1.6. <code>DROP DATABASE</code> Syntax .....	717
13.1.7. <code>DROP INDEX</code> Syntax .....	718
13.1.8. <code>DROP TABLE</code> Syntax .....	718
13.1.9. <code>RENAME TABLE</code> Syntax .....	718

---

---

13.2. Data Manipulation Statements .....	719
13.2.1. <code>DELETE</code> Syntax .....	719
13.2.2. <code>DO</code> Syntax .....	722
13.2.3. <code>HANDLER</code> Syntax .....	722
13.2.4. <code>INSERT</code> Syntax .....	724
13.2.5. <code>LOAD DATA INFILE</code> Syntax .....	730
13.2.6. <code>REPLACE</code> Syntax .....	738
13.2.7. <code>SELECT</code> Syntax .....	739
13.2.8. Subquery Syntax .....	749
13.2.9. <code>TRUNCATE</code> Syntax .....	759
13.2.10. <code>UPDATE</code> Syntax .....	760
13.3. MySQL Utility Statements .....	761
13.3.1. <code>DESCRIBE</code> Syntax .....	761
13.3.2. <code>HELP</code> Syntax .....	762
13.3.3. <code>USE</code> Syntax .....	764
13.4. MySQL Transactional and Locking Statements .....	765
13.4.1. <code>START TRANSACTION</code> , <code>COMMIT</code> , and <code>ROLLBACK</code> Syntax .....	765
13.4.2. Statements That Cannot Be Rolled Back .....	766
13.4.3. Statements That Cause an Implicit Commit .....	767
13.4.4. <code>SAVEPOINT</code> and <code>ROLLBACK TO SAVEPOINT</code> Syntax .....	767
13.4.5. <code>LOCK TABLES</code> and <code>UNLOCK TABLES</code> Syntax .....	767
13.4.6. <code>SET TRANSACTION</code> Syntax .....	770
13.5. Database Administration Statements .....	770
13.5.1. Account Management Statements .....	771
13.5.2. Table Maintenance Statements .....	779
13.5.3. <code>SET</code> Syntax .....	785
13.5.4. <code>SHOW</code> Syntax .....	790
13.5.5. Other Administrative Statements .....	809
13.6. Replication Statements .....	813
13.6.1. SQL Statements for Controlling Master Servers .....	813
13.6.2. SQL Statements for Controlling Slave Servers .....	816
13.7. SQL Syntax for Prepared Statements .....	825
14. Storage Engines and Table Types .....	828
14.1. The <code>MyISAM</code> Storage Engine .....	830
14.1.1. <code>MyISAM</code> Startup Options .....	832
14.1.2. Space Needed for Keys .....	833
14.1.3. <code>MyISAM</code> Table Storage Formats .....	834
14.1.4. <code>MyISAM</code> Table Problems .....	836
14.2. The <code>InnoDB</code> Storage Engine .....	838
14.2.1. <code>InnoDB</code> Overview .....	838
14.2.2. <code>InnoDB</code> Contact Information .....	838
14.2.3. <code>InnoDB</code> in MySQL 3.23 .....	838
14.2.4. <code>InnoDB</code> Configuration .....	839
14.2.5. <code>InnoDB</code> Startup Options and System Variables .....	845
14.2.6. Creating the <code>InnoDB</code> Tablespace .....	851
14.2.7. Creating and Using <code>InnoDB</code> Tables .....	852
14.2.8. Adding and Removing <code>InnoDB</code> Data and Log Files .....	860
14.2.9. Backing Up and Recovering an <code>InnoDB</code> Database .....	861
14.2.10. Moving an <code>InnoDB</code> Database to Another Machine .....	864
14.2.11. <code>InnoDB</code> Transaction Model and Locking .....	864
14.2.12. <code>InnoDB</code> Performance Tuning Tips .....	874
14.2.13. Implementation of Multi-Versioning .....	880
14.2.14. <code>InnoDB</code> Table and Index Structures .....	880
14.2.15. <code>InnoDB</code> File Space Management and Disk I/O .....	882
14.2.16. <code>InnoDB</code> Error Handling .....	884
14.2.17. Restrictions on <code>InnoDB</code> Tables .....	890
14.2.18. <code>InnoDB</code> Troubleshooting .....	893
14.3. The <code>MERGE</code> Storage Engine .....	894

---

---

14.3.1. <a href="#">MERGE</a> Table Problems .....	896
14.4. The <a href="#">MEMORY (HEAP)</a> Storage Engine .....	897
14.5. The <a href="#">BDB (BerkeleyDB)</a> Storage Engine .....	899
14.5.1. Operating Systems Supported by <a href="#">BDB</a> .....	900
14.5.2. Installing <a href="#">BDB</a> .....	901
14.5.3. <a href="#">BDB</a> Startup Options .....	901
14.5.4. Characteristics of <a href="#">BDB</a> Tables .....	902
14.5.5. Restrictions on <a href="#">BDB</a> Tables .....	904
14.5.6. Errors That May Occur When Using <a href="#">BDB</a> Tables .....	904
14.6. The <a href="#">EXAMPLE</a> Storage Engine .....	905
14.7. The <a href="#">ARCHIVE</a> Storage Engine .....	905
14.8. The <a href="#">CSV</a> Storage Engine .....	906
14.9. The <a href="#">BLACKHOLE</a> Storage Engine .....	907
14.10. The <a href="#">ISAM</a> Storage Engine .....	908
15. MySQL Cluster .....	910
15.1. MySQL Cluster Overview .....	911
15.2. Basic MySQL Cluster Concepts .....	912
15.2.1. MySQL Cluster Nodes, Node Groups, Replicas, and Partitions .....	913
15.3. Simple Multi-Computer How-To .....	915
15.3.1. Hardware, Software, and Networking .....	917
15.3.2. Multi-Computer Installation .....	918
15.3.3. Multi-Computer Configuration .....	920
15.3.4. Initial Startup .....	921
15.3.5. Loading Sample Data and Performing Queries .....	922
15.3.6. Safe Shutdown and Restart .....	926
15.4. MySQL Cluster Configuration .....	926
15.4.1. Building MySQL Cluster from Source Code .....	927
15.4.2. Installing the Software .....	927
15.4.3. Quick Test Setup of MySQL Cluster .....	927
15.4.4. Configuration File .....	929
15.4.5. Overview of Cluster Configuration Parameters .....	958
15.4.6. Configuring Parameters for Local Checkpoints .....	968
15.5. Upgrading and Downgrading MySQL Cluster .....	970
15.5.1. Performing a Rolling Restart of the Cluster .....	970
15.5.2. Cluster Upgrade and Downgrade Compatibility .....	971
15.6. Process Management in MySQL Cluster .....	973
15.6.1. MySQL Server Process Usage for MySQL Cluster .....	974
15.6.2. <a href="#">ndbd</a> , the Storage Engine Node Process .....	975
15.6.3. <a href="#">ndb_mgmd</a> , the Management Server Process .....	976
15.6.4. <a href="#">ndb_mgm</a> , the Management Client Process .....	977
15.6.5. Command Options for MySQL Cluster Processes .....	977
15.7. Management of MySQL Cluster .....	980
15.7.1. MySQL Cluster Startup Phases .....	980
15.7.2. Commands in the Management Client .....	983
15.7.3. Event Reports Generated in MySQL Cluster .....	984
15.7.4. Single-User Mode .....	991
15.8. On-line Backup of MySQL Cluster .....	991
15.8.1. Cluster Backup Concepts .....	992
15.8.2. Using The Management Client to Create a Backup .....	992
15.8.3. How to Restore a Cluster Backup .....	993
15.8.4. Configuration for Cluster Backup .....	995
15.8.5. Backup Troubleshooting .....	996
15.9. Cluster Utility Programs .....	996
15.9.1. <a href="#">ndb_delete_all</a> .....	998
15.9.2. <a href="#">ndb_desc</a> .....	998
15.9.3. <a href="#">ndb_drop_index</a> .....	999
15.9.4. <a href="#">ndb_drop_table</a> .....	1000
15.9.5. <a href="#">ndb_error_reporter</a> .....	1000

---

---

15.9.6. <code>ndb_print_backup_file</code> .....	1001
15.9.7. <code>ndb_print_schema_file</code> .....	1001
15.9.8. <code>ndb_print_sys_file</code> .....	1001
15.9.9. <code>ndb_select_all</code> .....	1002
15.9.10. <code>ndb_select_count</code> .....	1004
15.9.11. <code>ndb_show_tables</code> .....	1004
15.9.12. <code>ndb_size.pl</code> .....	1005
15.9.13. <code>ndb_waiter</code> .....	1007
15.10. Using High-Speed Interconnects with MySQL Cluster .....	1009
15.10.1. Configuring MySQL Cluster to use SCI Sockets .....	1009
15.10.2. Understanding the Impact of Cluster Interconnects .....	1013
15.11. Known Limitations of MySQL Cluster .....	1014
15.12. MySQL Cluster FAQ .....	1020
15.13. MySQL Cluster Glossary .....	1027
16. Spatial Extensions .....	1032
16.1. Introduction to MySQL Spatial Support .....	1032
16.2. The OpenGIS Geometry Model .....	1033
16.2.1. The Geometry Class Hierarchy .....	1033
16.2.2. Class <code>Geometry</code> .....	1034
16.2.3. Class <code>Point</code> .....	1035
16.2.4. Class <code>Curve</code> .....	1036
16.2.5. Class <code>LineString</code> .....	1036
16.2.6. Class <code>Surface</code> .....	1036
16.2.7. Class <code>Polygon</code> .....	1037
16.2.8. Class <code>GeometryCollection</code> .....	1037
16.2.9. Class <code>MultiPoint</code> .....	1037
16.2.10. Class <code>MultiCurve</code> .....	1038
16.2.11. Class <code>MultiLineString</code> .....	1038
16.2.12. Class <code>MultiSurface</code> .....	1038
16.2.13. Class <code>MultiPolygon</code> .....	1039
16.3. Supported Spatial Data Formats .....	1039
16.3.1. Well-Known Text (WKT) Format .....	1040
16.3.2. Well-Known Binary (WKB) Format .....	1040
16.4. Creating a Spatially Enabled MySQL Database .....	1041
16.4.1. MySQL Spatial Data Types .....	1041
16.4.2. Creating Spatial Values .....	1042
16.4.3. Creating Spatial Columns .....	1044
16.4.4. Populating Spatial Columns .....	1045
16.4.5. Fetching Spatial Data .....	1046
16.5. Analyzing Spatial Information .....	1046
16.5.1. Geometry Format Conversion Functions .....	1047
16.5.2. <code>Geometry</code> Functions .....	1047
16.5.3. Functions That Create New Geometries from Existing Ones .....	1053
16.5.4. Functions for Testing Spatial Relations Between Geometric Objects .....	1054
16.5.5. Relations on Geometry Minimal Bounding Rectangles (MBRs) .....	1054
16.5.6. Functions That Test Spatial Relationships Between Geometries .....	1055
16.6. Optimizing Spatial Analysis .....	1056
16.6.1. Creating Spatial Indexes .....	1057
16.6.2. Using a Spatial Index .....	1058
16.7. MySQL Conformance and Compatibility .....	1059
17. APIs and Libraries .....	1060
17.1. <code>libmysqld</code> , the Embedded MySQL Server Library .....	1060
17.1.1. Overview of the Embedded MySQL Server Library .....	1060
17.1.2. Compiling Programs with <code>libmysqld</code> .....	1060
17.1.3. Restrictions When Using the Embedded MySQL Server .....	1061
17.1.4. Options with the Embedded Server .....	1061
17.1.5. Embedded Server Examples .....	1062
17.1.6. Licensing the Embedded Server .....	1065

---

---

17.2. MySQL C API .....	1065
17.2.1. C API Data types .....	1065
17.2.2. C API Function Overview .....	1069
17.2.3. C API Function Descriptions .....	1073
17.2.4. C API Prepared Statements .....	1120
17.2.5. C API Prepared Statement Data types .....	1120
17.2.6. C API Prepared Statement Function Overview .....	1123
17.2.7. C API Prepared Statement Function Descriptions .....	1126
17.2.8. C API Prepared statement problems .....	1148
17.2.9. C API Handling of Multiple Statement Execution .....	1149
17.2.10. C API Handling of Date and Time Values .....	1149
17.2.11. C API Threaded Function Descriptions .....	1150
17.2.12. C API Embedded Server Function Descriptions .....	1151
17.2.13. Common Questions and Problems When Using the C API .....	1153
17.2.14. Building Client Programs .....	1155
17.2.15. How to Make a Threaded Client .....	1155
17.3. MySQL PHP API .....	1157
17.3.1. Common Problems with MySQL and PHP .....	1157
17.3.2. Enabling Both <code>mysql</code> and <code>mysql_i</code> in PHP .....	1158
17.4. MySQL Perl API .....	1158
17.5. MySQL C++ API .....	1159
17.5.1. Borland C++ .....	1159
17.6. MySQL Python API .....	1159
17.7. MySQL Tcl API .....	1159
17.8. MySQL Eiffel Wrapper .....	1159
17.9. MySQL Program Development Utilities .....	1159
17.9.1. <code>mysql2mysql</code> — Convert mSQL Programs for Use with MySQL .....	1160
17.9.2. <code>mysql_config</code> — Get Compile Options for Compiling Clients .....	1160
18. Connectors .....	1162
18.1. MySQL Connector/ODBC .....	1162
18.1.1. Introduction to MyODBC .....	1163
18.1.2. How to Install MyODBC .....	1166
18.1.3. MyODBC Configuration .....	1186
18.1.4. MyODBC Examples .....	1202
18.1.5. MyODBC Reference .....	1218
18.1.6. MyODBC Notes and Tips .....	1224
18.1.7. MyODBC Support .....	1234
18.2. Connector/NET .....	1235
18.2.1. Connector/NET Versions .....	1236
18.2.2. How to install Connector/NET .....	1236
18.2.3. Connector/NET Examples .....	1243
18.2.4. Connector/NET Reference .....	1292
18.2.5. Connector/NET Notes and Tips .....	1406
18.2.6. Connector/NET Support .....	1423
18.3. MySQL Connector/J .....	1424
18.3.1. Connector/J Versions .....	1425
18.3.2. Installing Connector/J .....	1426
18.3.3. Connector/J Examples .....	1430
18.3.4. Connector/J (JDBC) Reference .....	1430
18.3.5. Connector/J Notes and Tips .....	1448
18.3.6. Connector/J Support .....	1465
18.4. MySQL Connector/MXJ .....	1467
18.4.1. Introduction to Connector/MXJ .....	1467
18.4.2. Installing Connector/MXJ .....	1469
18.4.3. Connector/MXJ Configuration .....	1473
18.4.4. Connector/MXJ Reference .....	1476
18.4.5. Connector/MXJ Notes and Tips .....	1477
18.4.6. Connector/MXJ Support .....	1481

---

---

18.5. Connector/PHP .....	1482
19. Extending MySQL .....	1483
19.1. MySQL Internals .....	1483
19.1.1. MySQL Threads .....	1483
19.1.2. MySQL Test Suite .....	1483
19.2. Adding New Functions to MySQL .....	1484
19.2.1. Features of the User-Defined Function Interface .....	1485
19.2.2. CREATE FUNCTION Syntax .....	1485
19.2.3. DROP FUNCTION Syntax .....	1486
19.2.4. Adding a New User-Defined Function .....	1486
19.2.5. Adding a New Native Function .....	1495
19.3. Adding New Procedures to MySQL .....	1496
19.3.1. Procedure Analyse .....	1496
19.3.2. Writing a Procedure .....	1496
A. Problems and Common Errors .....	1498
A.1. How to Determine What Is Causing a Problem .....	1498
A.2. Common Errors When Using MySQL Programs .....	1499
A.2.1. Access denied .....	1499
A.2.2. Can't connect to [local] MySQL server .....	1499
A.2.3. Client does not support authentication protocol ...	1502
A.2.4. Password Fails When Entered Interactively .....	1503
A.2.5. Host 'host_name' is blocked .....	1504
A.2.6. Too many connections .....	1504
A.2.7. Out of memory .....	1504
A.2.8. MySQL server has gone away .....	1505
A.2.9. Packet too large .....	1507
A.2.10. Communication Errors and Aborted Connections .....	1508
A.2.11. The table is full .....	1509
A.2.12. Can't create/write to file .....	1509
A.2.13. Commands out of sync .....	1510
A.2.14. Ignoring user .....	1510
A.2.15. Table 'tbl_name' doesn't exist .....	1510
A.2.16. Can't initialize character set .....	1511
A.2.17. File Not Found .....	1511
A.3. Installation-Related Issues .....	1512
A.3.1. Problems Linking to the MySQL Client Library .....	1512
A.3.2. Problems with File Permissions .....	1513
A.4. Administration-Related Issues .....	1514
A.4.1. How to Reset the Root Password .....	1514
A.4.2. What to Do If MySQL Keeps Crashing .....	1516
A.4.3. How MySQL Handles a Full Disk .....	1518
A.4.4. Where MySQL Stores Temporary Files .....	1519
A.4.5. How to Protect or Change the MySQL Unix Socket File .....	1519
A.4.6. Time Zone Problems .....	1520
A.5. Query-Related Issues .....	1520
A.5.1. Case Sensitivity in Searches .....	1520
A.5.2. Problems Using DATE Columns .....	1521
A.5.3. Problems with NULL Values .....	1522
A.5.4. Problems with Column Aliases .....	1523
A.5.5. Rollback Failure for Non-Transactional Tables .....	1523
A.5.6. Deleting Rows from Related Tables .....	1524
A.5.7. Solving Problems with No Matching Rows .....	1524
A.5.8. Problems with Floating-Point Comparisons .....	1525
A.6. Optimizer-Related Issues .....	1527
A.7. Table Definition-Related Issues .....	1527
A.7.1. Problems with ALTER TABLE .....	1528
A.7.2. How to Change the Order of Columns in a Table .....	1528
A.7.3. TEMPORARY TABLE Problems .....	1529

---

---

A.8. Known Issues in MySQL .....	1529
A.8.1. Issues in MySQL 3.23 Fixed in a Later MySQL Version .....	1529
A.8.2. Issues in MySQL 4.0 Fixed in a Later Version .....	1530
A.8.3. Issues in MySQL 4.1 Fixed in a Later Version .....	1530
A.8.4. Open Issues in MySQL .....	1530
B. Error Codes and Messages .....	1535
B.1. Server Error Codes and Messages .....	1535
B.2. Client Error Codes and Messages .....	1556
C. Credits .....	1561
C.1. Developers at MySQL AB .....	1561
C.2. Contributors to MySQL .....	1566
C.3. Documenters and translators .....	1570
C.4. Libraries used by and included with MySQL .....	1572
C.5. Packages that support MySQL .....	1573
C.6. Tools that were used to create MySQL .....	1573
C.7. Supporters of MySQL .....	1574
D. MySQL Change History .....	1575
D.1. Changes in release 4.1.x (Production) .....	1575
D.1.1. Changes in release 4.1.22 (Not yet released) .....	1576
D.1.2. Changes in release 4.1.21 (19 July 2006) .....	1580
D.1.3. Changes in release 4.1.20 (24 May 2006) .....	1585
D.1.4. Changes in release 4.1.19 (29 April 2006) .....	1586
D.1.5. Changes in release 4.1.18 (27 January 2006) .....	1591
D.1.6. Changes in release 4.1.17 (Not released) .....	1592
D.1.7. Changes in release 4.1.16 (29 November 2005) .....	1594
D.1.8. Changes in release 4.1.15 (13 October 2005) .....	1598
D.1.9. Changes in release 4.1.14 (17 August 2005) .....	1603
D.1.10. Changes in release 4.1.13 (15 July 2005) .....	1606
D.1.11. Changes in release 4.1.12 (13 May 2005) .....	1613
D.1.12. Changes in release 4.1.11 (01 April 2005) .....	1618
D.1.13. Changes in release 4.1.10 (12 February 2005) .....	1624
D.1.14. Changes in release 4.1.9 (11 January 2005) .....	1629
D.1.15. Changes in release 4.1.8 (14 December 2004) .....	1631
D.1.16. Changes in release 4.1.7 (23 October 2004: Production) .....	1636
D.1.17. Changes in release 4.1.6 (10 October 2004) .....	1637
D.1.18. Changes in release 4.1.5 (16 September 2004) .....	1639
D.1.19. Changes in release 4.1.4 (26 August 2004: Gamma) .....	1640
D.1.20. Changes in release 4.1.3 (28 June 2004: Beta) .....	1643
D.1.21. Changes in release 4.1.2 (28 May 2004) .....	1646
D.1.22. Changes in release 4.1.1 (01 December 2003) .....	1656
D.1.23. Changes in release 4.1.0 (03 April 2003: Alpha) .....	1661
D.2. Changes in release 4.0.x (Recent; still supported) .....	1664
D.2.1. Changes in release 4.0.28 (Not yet released) .....	1665
D.2.2. Changes in release 4.0.27 (06 May 2006) .....	1665
D.2.3. Changes in release 4.0.26 (08 September 2005) .....	1667
D.2.4. Changes in release 4.0.25 (05 July 2005) .....	1668
D.2.5. Changes in release 4.0.24 (04 March 2005) .....	1669
D.2.6. Changes in release 4.0.23 (18 December 2004) .....	1672
D.2.7. Changes in release 4.0.22 (27 October 2004) .....	1673
D.2.8. Changes in release 4.0.21 (06 September 2004) .....	1675
D.2.9. Changes in release 4.0.20 (17 May 2004) .....	1678
D.2.10. Changes in release 4.0.19 (04 May 2004) .....	1678
D.2.11. Changes in release 4.0.18 (12 February 2004) .....	1682
D.2.12. Changes in release 4.0.17 (14 December 2003) .....	1685
D.2.13. Changes in release 4.0.16 (17 October 2003) .....	1688
D.2.14. Changes in release 4.0.15 (03 September 2003) .....	1690
D.2.15. Changes in release 4.0.14 (18 July 2003) .....	1695
D.2.16. Changes in release 4.0.13 (16 May 2003) .....	1698

---



---

D.2.17. Changes in release 4.0.12 (15 March 2003: Production) .....	1702
D.2.18. Changes in release 4.0.11 (20 February 2003) .....	1704
D.2.19. Changes in release 4.0.10 (29 January 2003) .....	1705
D.2.20. Changes in release 4.0.9 (09 January 2003) .....	1707
D.2.21. Changes in release 4.0.8 (07 January 2003) .....	1707
D.2.22. Changes in release 4.0.7 (20 December 2002) .....	1708
D.2.23. Changes in release 4.0.6 (14 December 2002: Gamma) .....	1708
D.2.24. Changes in release 4.0.5 (13 November 2002) .....	1710
D.2.25. Changes in release 4.0.4 (29 September 2002) .....	1712
D.2.26. Changes in release 4.0.3 (26 August 2002: Beta) .....	1714
D.2.27. Changes in release 4.0.2 (01 July 2002) .....	1716
D.2.28. Changes in release 4.0.1 (23 December 2001) .....	1720
D.2.29. Changes in release 4.0.0 (October 2001: Alpha) .....	1721
D.3. Changes in release 3.23.x (Recent; still supported) .....	1723
D.3.1. Changes in release 3.23.59 (Not yet released) .....	1723
D.3.2. Changes in release 3.23.58 (11 September 2003) .....	1724
D.3.3. Changes in release 3.23.57 (06 June 2003) .....	1725
D.3.4. Changes in release 3.23.56 (13 March 2003) .....	1726
D.3.5. Changes in release 3.23.55 (23 January 2003) .....	1727
D.3.6. Changes in release 3.23.54 (05 December 2002) .....	1728
D.3.7. Changes in release 3.23.53 (09 October 2002) .....	1729
D.3.8. Changes in release 3.23.52 (14 August 2002) .....	1730
D.3.9. Changes in release 3.23.51 (31 May 2002) .....	1731
D.3.10. Changes in release 3.23.50 (21 April 2002) .....	1731
D.3.11. Changes in release 3.23.49 (14 February 2002) .....	1732
D.3.12. Changes in release 3.23.48 (07 February 2002) .....	1733
D.3.13. Changes in release 3.23.47 (27 December 2001) .....	1734
D.3.14. Changes in release 3.23.46 (29 November 2001) .....	1734
D.3.15. Changes in release 3.23.45 (22 November 2001) .....	1735
D.3.16. Changes in release 3.23.44 (31 October 2001) .....	1735
D.3.17. Changes in release 3.23.43 (04 October 2001) .....	1736
D.3.18. Changes in release 3.23.42 (08 September 2001) .....	1737
D.3.19. Changes in release 3.23.41 (11 August 2001) .....	1738
D.3.20. Changes in release 3.23.40 (18 July 2001) .....	1738
D.3.21. Changes in release 3.23.39 (12 June 2001) .....	1739
D.3.22. Changes in release 3.23.38 (09 May 2001) .....	1740
D.3.23. Changes in release 3.23.37 (17 April 2001) .....	1740
D.3.24. Changes in release 3.23.36 (27 March 2001) .....	1741
D.3.25. Changes in release 3.23.35 (15 March 2001) .....	1742
D.3.26. Changes in release 3.23.34a (11 March 2001) .....	1742
D.3.27. Changes in release 3.23.34 (10 March 2001) .....	1742
D.3.28. Changes in release 3.23.33 (09 February 2001) .....	1743
D.3.29. Changes in release 3.23.32 (22 January 2001) .....	1745
D.3.30. Changes in release 3.23.31 (17 January 2001: Production) .....	1745
D.3.31. Changes in release 3.23.30 (04 January 2001) .....	1746
D.3.32. Changes in release 3.23.29 (16 December 2000) .....	1747
D.3.33. Changes in release 3.23.28 (22 November 2000: Gamma) .....	1749
D.3.34. Changes in release 3.23.27 (24 October 2000) .....	1751
D.3.35. Changes in release 3.23.26 (18 October 2000) .....	1751
D.3.36. Changes in release 3.23.25 (29 September 2000) .....	1752
D.3.37. Changes in release 3.23.24 (08 September 2000) .....	1753
D.3.38. Changes in release 3.23.23 (01 September 2000) .....	1754
D.3.39. Changes in release 3.23.22 (31 July 2000) .....	1755
D.3.40. Changes in release 3.23.21 (04 July 2000) .....	1756
D.3.41. Changes in release 3.23.20 (28 June 2000: Beta) .....	1757
D.3.42. Changes in release 3.23.19 .....	1757
D.3.43. Changes in release 3.23.18 (11 June 2000) .....	1757
D.3.44. Changes in release 3.23.17 (07 June 2000) .....	1758

---

---

D.3.45. Changes in release 3.23.16 (16 May 2000)	1758
D.3.46. Changes in release 3.23.15 (08 May 2000)	1759
D.3.47. Changes in release 3.23.14 (09 April 2000)	1760
D.3.48. Changes in release 3.23.13 (14 March 2000)	1761
D.3.49. Changes in release 3.23.12 (07 March 2000)	1761
D.3.50. Changes in release 3.23.11 (16 February 2000)	1762
D.3.51. Changes in release 3.23.10 (30 January 2000)	1763
D.3.52. Changes in release 3.23.9 (29 January 2000)	1763
D.3.53. Changes in release 3.23.8 (02 January 2000)	1764
D.3.54. Changes in release 3.23.7 (10 December 1999)	1764
D.3.55. Changes in release 3.23.6 (15 December 1999)	1765
D.3.56. Changes in release 3.23.5 (20 October 1999)	1766
D.3.57. Changes in release 3.23.4 (28 September 1999)	1767
D.3.58. Changes in release 3.23.3 (13 September 1999)	1767
D.3.59. Changes in release 3.23.2 (09 August 1999)	1768
D.3.60. Changes in release 3.23.1 (08 July 1999)	1769
D.3.61. Changes in release 3.23.0 (05 July 1999: Alpha)	1769
D.4. Changes in <a href="#">InnoDB</a>	1771
D.4.1. Changes in MySQL/InnoDB-4.0.21, September 10, 2004	1771
D.4.2. Changes in MySQL/InnoDB-4.1.4, August 31, 2004	1773
D.4.3. Changes in MySQL/InnoDB-4.1.3, June 28, 2004	1774
D.4.4. Changes in MySQL/InnoDB-4.1.2, May 30, 2004	1775
D.4.5. Changes in MySQL/InnoDB-4.0.20, May 18, 2004	1776
D.4.6. Changes in MySQL/InnoDB-4.0.19, May 4, 2004	1776
D.4.7. Changes in MySQL/InnoDB-4.0.18, February 13, 2004	1777
D.4.8. Changes in MySQL/InnoDB-5.0.0, December 24, 2003	1778
D.4.9. Changes in MySQL/InnoDB-4.0.17, December 17, 2003	1778
D.4.10. Changes in MySQL/InnoDB-4.1.1, December 4, 2003	1778
D.4.11. Changes in MySQL/InnoDB-4.0.16, October 22, 2003	1779
D.4.12. Changes in MySQL/InnoDB-3.23.58, September 15, 2003	1779
D.4.13. Changes in MySQL/InnoDB-4.0.15, September 10, 2003	1779
D.4.14. Changes in MySQL/InnoDB-4.0.14, July 22, 2003	1780
D.4.15. Changes in MySQL/InnoDB-3.23.57, June 20, 2003	1781
D.4.16. Changes in MySQL/InnoDB-4.0.13, May 20, 2003	1782
D.4.17. Changes in MySQL/InnoDB-4.1.0, April 3, 2003	1783
D.4.18. Changes in MySQL/InnoDB-3.23.56, March 17, 2003	1783
D.4.19. Changes in MySQL/InnoDB-4.0.12, March 18, 2003	1783
D.4.20. Changes in MySQL/InnoDB-4.0.11, February 25, 2003	1783
D.4.21. Changes in MySQL/InnoDB-4.0.10, February 4, 2003	1783
D.4.22. Changes in MySQL/InnoDB-3.23.55, January 24, 2003	1784
D.4.23. Changes in MySQL/InnoDB-4.0.9, January 14, 2003	1785
D.4.24. Changes in MySQL/InnoDB-4.0.8, January 7, 2003	1785
D.4.25. Changes in MySQL/InnoDB-4.0.7, December 26, 2002	1785
D.4.26. Changes in MySQL/InnoDB-4.0.6, December 19, 2002	1785
D.4.27. Changes in MySQL/InnoDB-3.23.54, December 12, 2002	1786
D.4.28. Changes in MySQL/InnoDB-4.0.5, November 18, 2002	1786
D.4.29. Changes in MySQL/InnoDB-3.23.53, October 9, 2002	1787
D.4.30. Changes in MySQL/InnoDB-4.0.4, October 2, 2002	1788
D.4.31. Changes in MySQL/InnoDB-4.0.3, August 28, 2002	1789
D.4.32. Changes in MySQL/InnoDB-3.23.52, August 16, 2002	1789
D.4.33. Changes in MySQL/InnoDB-4.0.2, July 10, 2002	1790
D.4.34. Changes in MySQL/InnoDB-3.23.51, June 12, 2002	1791
D.4.35. Changes in MySQL/InnoDB-3.23.50, April 23, 2002	1791
D.4.36. Changes in MySQL/InnoDB-3.23.49, February 17, 2002	1792
D.4.37. Changes in MySQL/InnoDB-3.23.48, February 9, 2002	1792
D.4.38. Changes in MySQL/InnoDB-3.23.47, December 28, 2001	1793
D.4.39. Changes in MySQL/InnoDB-4.0.1, December 23, 2001	1793
D.4.40. Changes in MySQL/InnoDB-3.23.46, November 30, 2001	1793

---

---

D.4.41. Changes in MySQL/InnoDB-3.23.45, November 23, 2001 .....	1794
D.4.42. Changes in MySQL/InnoDB-3.23.44, November 2, 2001 .....	1794
D.4.43. Changes in MySQL/InnoDB-3.23.43, October 4, 2001 .....	1795
D.4.44. Changes in MySQL/InnoDB-3.23.42, September 9, 2001 .....	1795
D.4.45. Changes in MySQL/InnoDB-3.23.41, August 13, 2001 .....	1795
D.4.46. Changes in MySQL/InnoDB-3.23.40, July 16, 2001 .....	1796
D.4.47. Changes in MySQL/InnoDB-3.23.39, June 13, 2001 .....	1796
D.4.48. Changes in MySQL/InnoDB-3.23.38, May 12, 2001 .....	1796
D.5. Changes in MySQL Cluster .....	1796
D.5.1. Changes in MySQL Cluster-5.0.7 (10 June 2005) .....	1796
D.5.2. Changes in MySQL Cluster-5.0.6 (26 May 2005) .....	1797
D.5.3. Changes in MySQL Cluster-5.0.5 (Not released) .....	1797
D.5.4. Changes in MySQL Cluster-5.0.4 (16 April 2005) .....	1798
D.5.5. Changes in MySQL Cluster-5.0.3 (23 March 2005: Beta) .....	1798
D.5.6. Changes in MySQL Cluster-5.0.1 (27 July 2004) .....	1798
D.5.7. Changes in MySQL Cluster-4.1.13 (15 July 2005) .....	1799
D.5.8. Changes in MySQL Cluster-4.1.12 (13 May 2005) .....	1799
D.5.9. Changes in MySQL Cluster-4.1.11 (01 April 2005) .....	1800
D.5.10. Changes in MySQL Cluster-4.1.10 (12 February 2005) .....	1801
D.5.11. Changes in MySQL Cluster-4.1.9 (13 January 2005) .....	1801
D.5.12. Changes in MySQL Cluster-4.1.8 (14 December 2004) .....	1802
D.5.13. Changes in MySQL Cluster-4.1.7 (23 October 2004) .....	1805
D.5.14. Changes in MySQL Cluster-4.1.6 (10 October 2004) .....	1805
D.5.15. Changes in MySQL Cluster-4.1.5 (16 September 2004) .....	1806
D.5.16. Changes in MySQL Cluster-4.1.4 (31 August 2004) .....	1808
D.5.17. Changes in MySQL Cluster-4.1.3 (28 June 2004) .....	1808
D.6. MySQL Connector/ODBC (MyODBC) Change History .....	1808
D.6.1. Changes in MyODBC 3.51.13 .....	1808
D.6.2. Changes in MyODBC 3.51.12 .....	1809
D.6.3. Changes in MyODBC 3.51.11 .....	1809
D.7. MySQL Connector/NET Change History .....	1809
D.7.1. Version 5.0.1 (not yet released) .....	1809
D.7.2. Version 5.0.0 .....	1810
D.7.3. Version 1.0.8 (not yet released) .....	1811
D.7.4. Version 1.0.7 .....	1812
D.7.5. Version 1.0.6 .....	1812
D.7.6. Version 1.0.5 .....	1812
D.7.7. Version 1.0.4 1-20-05 .....	1813
D.7.8. Version 1.0.3-gamma 12-10-04 .....	1814
D.7.9. Version 1.0.2-gamma 04-11-15 .....	1815
D.7.10. Version 1.0.1-beta2 04-10-27 .....	1815
D.7.11. Version 1.0.0 04-09-01 .....	1817
D.7.12. Version 0.9.0 04-08-30 .....	1817
D.7.13. Version 0.76 .....	1821
D.7.14. Version 0.75 .....	1822
D.7.15. Version 0.74 .....	1822
D.7.16. Version 0.71 .....	1824
D.7.17. Version 0.70 .....	1825
D.7.18. Version 0.68 .....	1827
D.7.19. Version 0.65 .....	1827
D.7.20. Version 0.60 .....	1828
D.7.21. Version 0.50 .....	1828
D.8. MySQL Connector/J Change History .....	1828
D.8.1. Changes in MySQL Connector/J 5.0.3 (26 July 2006) .....	1828
D.8.2. Changes in MySQL Connector/J 5.0.2-beta (11 July 2006) .....	1828
D.8.3. Changes in MySQL Connector/J 5.0.1-beta (Not Released) .....	1829
D.8.4. Changes in MySQL Connector/J 5.0.0-beta (22 December 2005) .....	1829
D.8.5. Changes in MySQL Connector/J 3.1.14 (not yet released) .....	1830

---

---

D.8.6. Changes in MySQL Connector/J 3.1.13 (26 May 2006) .....	1831
D.8.7. Changes in MySQL Connector/J 3.1.12 (30 November 2005) .....	1833
D.8.8. Changes in MySQL Connector/J 3.1.11-stable (07 October 2005) .....	1834
D.8.9. Changes in MySQL Connector/J 3.1.10-stable (23 June 2005) .....	1837
D.8.10. Changes in MySQL Connector/J 3.1.9-stable (22 June 2005) .....	1838
D.8.11. Changes in MySQL Connector/J 3.1.8-stable (14 April 2005) .....	1840
D.8.12. Changes in MySQL Connector/J 3.1.7-stable (18 February 2005) .....	1842
D.8.13. Changes in MySQL Connector/J 3.1.6-stable (23 December 2004) .....	1843
D.8.14. Changes in MySQL Connector/J 3.1.5-gamma (02 December 2004) .....	1844
D.8.15. Changes in MySQL Connector/J 3.1.4-beta (04 September 2004) .....	1845
D.8.16. Changes in MySQL Connector/J 3.1.3-beta (07 July 2004) .....	1846
D.8.17. Changes in MySQL Connector/J 3.1.2-alpha (09 June 2004) .....	1847
D.8.18. Changes in MySQL Connector/J 3.1.1-alpha (14 February 2004) .....	1848
D.8.19. Changes in MySQL Connector/J 3.1.0-alpha (18 February 2003) .....	1850
D.8.20. Changes in MySQL Connector/J 3.0.17-ga (23 June 2005) .....	1850
D.8.21. Changes in MySQL Connector/J 3.0.16-ga (15 November 2004) .....	1852
D.8.22. Changes in MySQL Connector/J 3.0.15-production (04 September 2004) .....	1852
D.8.23. Changes in MySQL Connector/J 3.0.14-production (28 May 2004) .....	1853
D.8.24. Changes in MySQL Connector/J 3.0.13-production (27 May 2004) .....	1853
D.8.25. Changes in MySQL Connector/J 3.0.12-production (18 May 2004) .....	1853
D.8.26. Changes in MySQL Connector/J 3.0.11-stable (19 February 2004) .....	1855
D.8.27. Changes in MySQL Connector/J 3.0.10-stable (13 January 2004) .....	1855
D.8.28. Changes in MySQL Connector/J 3.0.9-stable (07 October 2003) .....	1857
D.8.29. Changes in MySQL Connector/J 3.0.8-stable (23 May 2003) .....	1858
D.8.30. Changes in MySQL Connector/J 3.0.7-stable (08 April 2003) .....	1859
D.8.31. Changes in MySQL Connector/J 3.0.6-stable (18 February 2003) .....	1860
D.8.32. Changes in MySQL Connector/J 3.0.5-gamma (22 January 2003) .....	1860
D.8.33. Changes in MySQL Connector/J 3.0.4-gamma (06 January 2003) .....	1861
D.8.34. Changes in MySQL Connector/J 3.0.3-dev (17 December 2002) .....	1861
D.8.35. Changes in MySQL Connector/J 3.0.2-dev (08 November 2002) .....	1862
D.8.36. Changes in MySQL Connector/J 3.0.1-dev (21 September 2002) .....	1864
D.8.37. Changes in MySQL Connector/J 3.0.0-dev (31 July 2002) .....	1864
D.8.38. Changes in MySQL Connector/J 2.0.14 (16 May 2002) .....	1865
D.8.39. Changes in MySQL Connector/J 2.0.13 (24 April 2002) .....	1865
D.8.40. Changes in MySQL Connector/J 2.0.12 (07 April 2002) .....	1866
D.8.41. Changes in MySQL Connector/J 2.0.11 (27 January 2002) .....	1866
D.8.42. Changes in MySQL Connector/J 2.0.10 (24 January 2002) .....	1867
D.8.43. Changes in MySQL Connector/J 2.0.9 (13 January 2002) .....	1867
D.8.44. Changes in MySQL Connector/J 2.0.8 (25 November 2001) .....	1867
D.8.45. Changes in MySQL Connector/J 2.0.7 (24 October 2001) .....	1868
D.8.46. Changes in MySQL Connector/J 2.0.6 (16 June 2001) .....	1868
D.8.47. Changes in MySQL Connector/J 2.0.5 (13 June 2001) .....	1868
D.8.48. Changes in MySQL Connector/J 2.0.3 (03 December 2000) .....	1869
D.8.49. Changes in MySQL Connector/J 2.0.1 (06 April 2000) .....	1869
D.8.50. Changes in MySQL Connector/J 2.0.0pre5 (21 February 2000) .....	1870
D.8.51. Changes in MySQL Connector/J 2.0.0pre4 (10 January 2000) .....	1870
D.8.52. Changes in MySQL Connector/J 2.0.0pre (17 August 1999) .....	1870
D.8.53. Changes in MySQL Connector/J 1.2b (04 July 1999) .....	1871
D.8.54. Changes in MySQL Connector/J 1.2a (14 April 1999) .....	1871
D.8.55. Changes in MySQL Connector/J 1.1i (24 March 1999) .....	1871
D.8.56. Changes in MySQL Connector/J 1.1h (08 March 1999) .....	1872
D.8.57. Changes in MySQL Connector/J 1.1g (19 February 1999) .....	1872
D.8.58. Changes in MySQL Connector/J 1.1f (31 December 1998) .....	1872
D.8.59. Changes in MySQL Connector/J 1.1b (03 November 1998) .....	1873
D.8.60. Changes in MySQL Connector/J 1.1 (02 September 1998) .....	1873
D.8.61. Changes in MySQL Connector/J 1.0 (24 August 1998) .....	1874
D.8.62. Changes in MySQL Connector/J 0.9d (04 August 1998) .....	1874
D.8.63. Changes in MySQL Connector/J 0.9 (28 July 1998) .....	1874

D.8.64. Changes in MySQL Connector/J 0.8 (06 July 1998) .....	1875
D.8.65. Changes in MySQL Connector/J 0.7 (01 July 1998) .....	1875
D.8.66. Changes in MySQL Connector/J 0.6 (21 May 1998) .....	1875
E. Porting to Other Systems .....	1876
E.1. Debugging a MySQL Server .....	1877
E.1.1. Compiling MySQL for Debugging .....	1877
E.1.2. Creating Trace Files .....	1878
E.1.3. Debugging <code>mysqld</code> under <code>gdb</code> .....	1878
E.1.4. Using a Stack Trace .....	1880
E.1.5. Using Server Logs to Find Causes of Errors in <code>mysqld</code> .....	1881
E.1.6. Making a Test Case If You Experience Table Corruption .....	1881
E.2. Debugging a MySQL Client .....	1882
E.3. The DBUG Package .....	1882
E.4. Comments about RTS Threads .....	1884
E.5. Differences Between Thread Packages .....	1885
F. Environment Variables .....	1886
G. Regular Expressions .....	1887
H. Limits in MySQL .....	1891
H.1. Limits of Joins .....	1891
I. Feature Restrictions .....	1892
I.1. Restrictions on Subqueries .....	1892
J. GNU General Public License .....	1895
K. MySQL FLOSS License Exception .....	1900
Index .....	1903

---

## List of Tables

15.1. Steps for Cluster rolling restarts — by type .....

---

## List of Examples

18.1. Obtaining a connection from the <code>DriverManager</code> .....	1448
18.2. Using <code>java.sql.Statement</code> to execute a <code>SELECT</code> query .....	1449
18.3. Stored Procedures .....	1450
18.4. Using <code>Connection.prepareStatement()</code> .....	1450
18.5. Registering output parameters .....	1451
18.6. Setting <code>CallableStatement</code> input parameters .....	1451
18.7. Retrieving results and output parameter values .....	1452
18.8. Retrieving <code>AUTO_INCREMENT</code> column values using <code>Statement.getGeneratedKeys()</code> .....	1453
18.9. Retrieving <code>AUTO_INCREMENT</code> column values using <code>SELECT LAST_INSERT_ID()</code> .....	1454
18.10. Retrieving <code>AUTO_INCREMENT</code> column values in <code>Updatable ResultSets</code> .....	1455
18.11. Using a connection pool with a J2EE application server .....	1457
18.12. Example of transaction with retry logic .....	1463

---

# Preface

This is the Reference Manual for all releases of the MySQL Database System through version 4.1.22. It is applicable for older versions of the MySQL software (such as 3.23 or 4.0-production) because functional changes are indicated with reference to a version number. For later MySQL releases, see the appropriately-numbered edition of this manual.



---

# Chapter 1. General Information

The MySQL® software delivers a very fast, multi-threaded, multi-user, and robust SQL (Structured Query Language) database server. MySQL Server is intended for mission-critical, heavy-load production systems as well as for embedding into mass-deployed software. MySQL is a registered trademark of MySQL AB.

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

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

- For a discussion about the capabilities of the MySQL Database Server, see [Section 1.4.2, “The Main Features of MySQL”](#).
- For installation instructions, see [Chapter 2, \*Installing and Upgrading MySQL\*](#). For information about upgrading MySQL, see [Section 2.11, “Upgrading MySQL”](#).
- For information about configuring and administering MySQL Server, see [Chapter 5, \*Database Administration\*](#).
- For information about setting up replication servers, see [Chapter 6, \*Replication\*](#).
- For tips on porting the MySQL Database Software to new architectures or operating systems, see [Appendix E, \*Porting to Other Systems\*](#).
- For a tutorial introduction to the MySQL Database Server, see [Chapter 3, \*Tutorial\*](#).
- For benchmarking information, see the `sql-bench` benchmarking directory in your MySQL distribution.
- For a history of new features and bugfixes, see [Appendix D, \*MySQL Change History\*](#).
- For a list of currently known bugs and misfeatures, see [Section A.8, “Known Issues in MySQL”](#).
- For a list of all the contributors to this project, see [Appendix C, \*Credits\*](#).

## **Important:**

To report errors (often called “bugs”), please use the instructions at [Section 1.8, “How to Report Bugs or Problems”](#).

If you have found a sensitive security bug in MySQL Server, please let us know immediately by sending an email message to [<security@mysql.com>](mailto:security@mysql.com).

## 1.1. About This Manual

This is the Reference Manual for all releases of the MySQL Database System from version 3.23 through release 4.1.22. It is also applicable for versions of the MySQL software previous to 4.1 (such as 3.23 or 4.0) because functional changes are indicated with reference to version numbers. For later MySQL releases, see the appropriately-numbered edition of this manual.

Because this manual serves as a reference, it does not provide general instruction on SQL or relational

database concepts. It also does not teach you how to use your operating system or command-line interpreter.

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

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

The DocBook XML sources of this manual are available from <http://dev.mysql.com/tech-resources/sources.html>. You can check out a copy of the documentation repository with this command:

```
svn checkout http://svn.mysql.com/svnpublic/mysql/doc/
```

If you have any suggestions concerning additions or corrections to this manual, please send them to the documentation team at [<docs@mysql.com>](mailto:docs@mysql.com).

This manual was originally written by David Axmark and Michael “Monty” Widenius. It is maintained by the MySQL Documentation Team, consisting of Paul DuBois, Stefan Hinz, Mike Hillyer, and Jon Stephens. For the many other contributors, see [Appendix C, Credits](#).

The copyright to this manual is owned by the Swedish company MySQL AB. MySQL® and the MySQL logo are registered trademarks of MySQL AB. Other trademarks and registered trademarks referred to in this manual are the property of their respective owners, and are used for identification purposes only.

## 1.2. Conventions Used in This Manual

This manual uses certain typographical conventions:

- *Text in this style* is used for SQL statements; database, table, and column names; program listings and source code; and environment variables. Example: “To reload the grant tables, use the `FLUSH PRIVILEGES` statement.”
- **Text in this style** indicates input that you type in examples.
- *Text in this style* indicates the names of executable programs and scripts, examples being `mysql` (the MySQL command line client program) and `mysqld` (the MySQL server executable).
- *Text in this style* is used for variable input for which you should substitute a value of your own choosing.
- Filenames and directory names are written like this: “The global `my.cnf` file is located in the `/etc` directory.”
- Character sequences are written like this: “To specify a wildcard, use the ‘`%`’ character.”
- *Text in this style* is used for emphasis.
- **Text in this style** is used in table headings and to convey especially strong emphasis.

When commands are shown that are meant to be executed from within a particular program, the prompt

shown preceding the command indicates which command to use. For example, `shell>` indicates a command that you execute from your login shell, and `mysql>` indicates a statement that you execute from the `mysql` client program:

```
shell> type a shell command here
mysql> type a mysql statement here
```

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

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

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

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

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

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

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

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

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

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

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

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

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

An ellipsis (‘...’) indicates the omission of a section of a statement, typically to provide a shorter version of more complex syntax. For example, `INSERT ... SELECT` is shorthand for the form of `INSERT` statement that is followed by a `SELECT` statement.

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

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

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

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

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

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

## 1.3. Overview of MySQL AB

MySQL AB is the company of the MySQL founders and main developers. MySQL AB was originally established in Sweden by David Axmark, Allan Larsson, and Michael “Monty” Widenius.

We are dedicated to developing the MySQL database software and promoting it to new users. MySQL AB owns the copyright to the MySQL source code, the MySQL logo and (registered) trademark, and this manual. See [Section 1.4, “Overview of the MySQL Database Management System”](#).

The MySQL core values show our dedication to MySQL and Open Source.

These core values direct how MySQL AB works with the MySQL server software:

- To be the best and the most widely used database in the world
- To be available and affordable by all
- To be easy to use
- To be continuously improved while remaining fast and safe
- To be fun to use and improve
- To be free from bugs

These are the core values of the company MySQL AB and its employees:

- We subscribe to the Open Source philosophy and support the Open Source community
- We aim to be good citizens
- We prefer partners that share our values and mindset
- We answer email and provide support
- We are a virtual company, networking with others
- We work against software patents

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

By the way, the “AB” part of the company name is the acronym for the Swedish “aktiebolag,” or “stock company.” It translates to “MySQL, Inc.” In fact, MySQL, Inc. and MySQL GmbH are examples of MySQL AB subsidiaries. They are located in the United States and Germany, respectively.

## 1.4. Overview of the MySQL Database Management

## System

MySQL, the most popular Open Source SQL database management system, is developed, distributed, and supported by MySQL AB. MySQL AB is a commercial company, founded by the MySQL developers. It is a second generation Open Source company that unites Open Source values and methodology with a successful business model.

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

- MySQL is a database management system.

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

- MySQL is a relational database management system.

A relational database stores data in separate tables rather than putting all the data in one big store-room. This adds speed and flexibility. The SQL part of “MySQL” stands for “Structured Query Language.” SQL is the most common standardized language used to access databases and is defined by the ANSI/ISO SQL Standard. The SQL standard has been evolving since 1986 and several versions exist. In this manual, “SQL-92” refers to the standard released in 1992, “SQL:1999” refers to the standard released in 1999, and “SQL:2003” refers to the current version of the standard. We use the phrase “the SQL standard” to mean the current version of the SQL Standard at any time.

- MySQL software is Open Source.

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

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

If that is what you are looking for, you should give it a try. MySQL Server also has a practical set of features developed in close cooperation with our users. You can find a performance comparison of MySQL Server with other database managers on our benchmark page. See [Section 7.1.4, “The MySQL Benchmark Suite”](#).

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

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

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

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

- A large amount of contributed MySQL software is available.

It is very likely that your favorite application or language supports the MySQL Database Server.

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

## 1.4.1. History of MySQL

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

The derivation of the name MySQL is not clear. Our base directory and a large number of our libraries and tools have had the prefix “my” for well over 10 years. However, co-founder Monty Widenius's daughter is also named My. Which of the two gave its name to MySQL is still a mystery, even for us.

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

## 1.4.2. The Main Features of MySQL

The following list describes some of the important characteristics of the MySQL Database Software. See also [Section 1.6, “MySQL Development Roadmap”](#), for more information about current and upcoming features.

Internals and Portability:

- Written in C and C++.
- Tested with a broad range of different compilers.
- Works on many different platforms. See [Section 2.1.1, “Operating Systems Supported by MySQL”](#).
- Uses GNU Automake, Autoconf, and Libtool for portability.
- APIs for C, C++, Eiffel, Java, Perl, PHP, Python, Ruby, and Tcl are available. See [Chapter 17, APIs and Libraries](#).
- Fully multi-threaded using kernel threads. It can easily use multiple CPUs if they are available.
- Provides transactional and non-transactional storage engines.
- Uses very fast B-tree disk tables ([MyISAM](#)) with index compression.
- Relatively easy to add other storage engines. This is useful if you want to add an SQL interface to an in-house database.

- A very fast thread-based memory allocation system.
- Very fast joins using an optimized one-sweep multi-join.
- In-memory hash tables, which are used as temporary tables.
- SQL functions are implemented using a highly optimized class library and should be as fast as possible. Usually there is no memory allocation at all after query initialization.
- The MySQL code is tested with Purify (a commercial memory leakage detector) as well as with Valgrind, a GPL tool (<http://developer.kde.org/~sewardj/>).
- The server is available as a separate program for use in a client/server networked environment. It is also available as a library that can be embedded (linked) into standalone applications. Such applications can be used in isolation or in environments where no network is available.

Data Types:

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

Statements and Functions:

- Full operator and function support in the [SELECT](#) and [WHERE](#) clauses of queries. For example:

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

- Full support for SQL [GROUP BY](#) and [ORDER BY](#) clauses. Support for group functions ([COUNT\(\)](#), [COUNT\(DISTINCT ...\)](#), [AVG\(\)](#), [STD\(\)](#), [SUM\(\)](#), [MAX\(\)](#), [MIN\(\)](#), and [GROUP\\_CONCAT\(\)](#)).
- Support for [LEFT OUTER JOIN](#) and [RIGHT OUTER JOIN](#) with both standard SQL and ODBC syntax.
- Support for aliases on tables and columns as required by standard SQL.
- [DELETE](#), [INSERT](#), [REPLACE](#), and [UPDATE](#) return the number of rows that were changed (affected). It is possible to return the number of rows matched instead by setting a flag when connecting to the server.
- The MySQL-specific [SHOW](#) statement can be used to retrieve information about databases, storage engines, tables, and indexes.

The [EXPLAIN](#) statement can be used to determine how the optimizer resolves a query.

- Function names do not clash with table or column names. For example, [ABS](#) is a valid column name. The only restriction is that for a function call, no spaces are allowed between the function name and the '(' that follows it. See [Section 9.5, "Treatment of Reserved Words in MySQL"](#).
- You can mix tables from different databases in the same query (as of MySQL 3.22).

Security:

- A privilege and password system that is very flexible and secure, and that allows host-based verification. Passwords are secure because all password traffic is encrypted when you connect to a server.

Scalability and Limits:

- Handles large databases. We use MySQL Server with databases that contain 50 million records. We also know of users who use MySQL Server with 60,000 tables and about 5,000,000,000 rows.
- Up to 64 indexes per table are allowed (32 before MySQL 4.1.2). Each index may consist of 1 to 16 columns or parts of columns. The maximum index width is 1000 bytes (767 for [InnoDB](#)); before MySQL 4.1.2, the limit is 500 bytes. An index may use a prefix of a column for [CHAR](#), [VARCHAR](#), [BLOB](#), or [TEXT](#) column types.

Connectivity:

- Clients can connect to the MySQL server using TCP/IP sockets on any platform. On Windows systems in the NT family (NT, 2000, XP, 2003, or Vista), clients can connect using named pipes. On Unix systems, clients can connect using Unix domain socket files.
- In MySQL 4.1 and higher, Windows servers also support shared-memory connections if started with the `--shared-memory` option. Clients can connect through shared memory by using the `--protocol=memory` option.
- The Connector/ODBC (MyODBC) interface provides MySQL support for client programs that use ODBC (Open Database Connectivity) connections. For example, you can use MS Access to connect to your MySQL server. Clients can be run on Windows or Unix. MyODBC source is available. All ODBC 2.5 functions are supported, as are many others. See [Chapter 18, Connectors](#).
- The Connector/J interface provides MySQL support for Java client programs that use JDBC connections. Clients can be run on Windows or Unix. Connector/J source is available. See [Chapter 18, Connectors](#).
- MySQL Connector/NET enables developers to easily create .NET applications that require secure, high-performance data connectivity with MySQL. It implements the required ADO.NET interfaces and integrates into ADO.NET aware tools. Developers can build applications using their choice of .NET languages. MySQL Connector/NET is a fully managed ADO.NET driver written in 100% pure C#. See [Chapter 18, Connectors](#).

Localization:

- The server can provide error messages to clients in many languages. See [Section 5.10.2, “Setting the Error Message Language”](#).
- Full support for several different character sets, including [latin1](#) (cp1252), [german](#), [big5](#), [ujis](#), and more. For example, the Scandinavian characters ‘å’, ‘ä’ and ‘ö’ are allowed in table and column names. Unicode support is available as of MySQL 4.1.
- All data is saved in the chosen character set. All comparisons for normal string columns are case-insensitive.



- Sorting is done according to the chosen character set (using Swedish collation by default). It is possible to change this when the MySQL server is started. To see an example of very advanced sorting, look at the Czech sorting code. MySQL Server supports many different character sets that can be specified at compile time and runtime.

Clients and Tools:

- MySQL Server has built-in support for SQL statements to check, optimize, and repair tables. These statements are available from the command line through the `mysqlcheck` client. MySQL also includes `myisamchk`, a very fast command-line utility for performing these operations on `MyISAM` tables. See [Chapter 5, Database Administration](#).
- All MySQL programs can be invoked with the `--help` or `-?` options to obtain online assistance.

### 1.4.3. MySQL Stability

This section addresses the questions, “*How stable is MySQL Server?*” and, “*Can I depend on MySQL Server in this project?*” We will try to clarify these issues and answer some important questions that concern many potential users. The information in this section is based on data gathered from the mailing lists, which are very active in identifying problems as well as reporting types of use.

The original code stems back to the early 1980s. It provides a stable code base, and the `ISAM` table format used by the original storage engine remains backward-compatible. At `TcX`, the predecessor of MySQL AB, MySQL code has worked in projects since mid-1996, without any problems. When the MySQL Database Software initially was released to a wider public, our new users quickly found some pieces of untested code. Each new release since then has had fewer portability problems, even though each new release has also had many new features.

Each release of the MySQL Server has been usable. Problems have occurred only when users try code from the “gray zones.” Naturally, new users don’t know what the gray zones are; this section therefore attempts to document those areas that are currently known. The descriptions mostly deal with Versions 3.23 and later of MySQL Server. All known and reported bugs are fixed in the latest version, with the exception of those listed in the bugs section, which are design-related. See [Section A.8, “Known Issues in MySQL”](#).

The MySQL Server design is multi-layered with independent modules. Some of the newer modules are listed here with an indication of how well-tested each of them is:

- Replication (Stable)

Large groups of servers using replication are in production use, with good results. Work on enhanced replication features is continuing.

- `InnoDB` tables (Stable)

The `InnoDB` transactional storage engine has been stable since version 3.23.49. `InnoDB` is being used in large, heavy-load production systems.

- Full-text searches (Stable)

Full-text searching is widely used. Important feature enhancements were added in MySQL 4.0 and 4.1.

- `MyODBC` 3.51 (Stable)

[MyODBC 3.51](#) uses ODBC SDK 3.51 and is in wide production use. Some issues brought up appear to be application-related and independent of the ODBC driver or underlying database server.

## 1.4.4. How Large MySQL Tables Can Be

MySQL 3.22 had a 4GB (4 gigabyte) limit on table size. With the [MyISAM](#) storage engine in MySQL 3.23, the maximum table size was increased to 65536 terabytes ( $256^7 - 1$  bytes). With this larger allowed table size, the maximum effective table size for MySQL databases is usually determined by operating system constraints on file sizes, not by MySQL internal limits.

The [InnoDB](#) storage engine maintains [InnoDB](#) tables within a tablespace that can be created from several files. This allows a table to exceed the maximum individual file size. The tablespace can include raw disk partitions, which allows extremely large tables. The maximum tablespace size is 64TB.

The following table lists some examples of operating system file-size limits. This is only a rough guide and is not intended to be definitive. For the most up-to-date information, be sure to check the documentation specific to your operating system.

Operating System	File-size Limit
Linux 2.2-Intel 32-bit	2GB (LFS: 4GB)
Linux 2.4+	(using ext3 filesystem) 4TB
Solaris 9/10	16TB
NetWare w/NSS filesystem	8TB
Win32 w/ FAT/FAT32	2GB/4GB
Win32 w/ NTFS	2TB (possibly larger)
MacOS X w/ HFS+	2TB

On Linux 2.2, you can get [MyISAM](#) tables larger than 2GB in size by using the Large File Support (LFS) patch for the ext2 filesystem. On Linux 2.4, patches also exist for ReiserFS to get support for big files (up to 2TB). Most current Linux distributions are based on kernel 2.4 or higher and include all the required LFS patches. With JFS and XFS, petabyte and larger files are possible on Linux. However, the maximum available file size still depends on several factors, one of them being the filesystem used to store MySQL tables.

For a detailed overview about LFS in Linux, have a look at Andreas Jaeger's *Large File Support in Linux* page at [http://www.suse.de/~aj/linux\\_lfs.html](http://www.suse.de/~aj/linux_lfs.html).

Windows users please note: FAT and VFAT (FAT32) are *not* considered suitable for production use with MySQL. Use NTFS instead.

By default, MySQL creates [MyISAM](#) tables with an internal structure that allows a maximum size of about 4GB. You can check the maximum table size for a [MyISAM](#) table with the `SHOW TABLE STATUS` statement or with `myisamchk -dv tbl_name`. See [Section 13.5.4, "SHOW Syntax"](#).

If you need a [MyISAM](#) table that is larger than 4GB and your operating system supports large files, the `CREATE TABLE` statement supports `AVG_ROW_LENGTH` and `MAX_ROWS` options. See [Section 13.1.5, "CREATE TABLE Syntax"](#). You can also change these options with `ALTER TABLE` to increase a table's maximum allowable size after the table has been created. See [Section 13.1.2, "ALTER TABLE Syntax"](#).

Other ways to work around file-size limits for [MyISAM](#) tables are as follows:

- If your large table is read-only, you can use `myisampack` to compress it. `myisampack` usually compresses a table by at least 50%, so you can have, in effect, much bigger tables. `myisampack` also can merge multiple tables into a single table. See [Section 8.5, “myisampack — Generate Compressed, Read-Only MyISAM Tables”](#).
- MySQL includes a `MERGE` library that allows you to handle a collection of `MyISAM` tables that have identical structure as a single `MERGE` table. See [Section 14.3, “The MERGE Storage Engine”](#).

## 1.4.5. Year 2000 Compliance

The MySQL Server itself has no problems with Year 2000 (Y2K) compliance:

- MySQL Server uses Unix time functions that handle dates into the year 2037 for `TIMESTAMP` values. For `DATE` and `DATETIME` values, dates through the year 9999 are accepted.
- All MySQL date functions are implemented in one source file, `sql/time.cc`, and are coded very carefully to be year 2000-safe.
- In MySQL, the `YEAR` data type can store the years 0 and 1901 to 2155 in one byte and display them using two or four digits. All two-digit years are considered to be in the range 1970 to 2069, which means that if you store 01 in a `YEAR` column, MySQL Server treats it as 2001.

The following simple demonstration illustrates that MySQL Server has no problems with `DATE` or `DATETIME` values through the year 9999, and no problems with `TIMESTAMP` values until after the year 2030:

```
mysql> DROP TABLE IF EXISTS y2k;
Query OK, 0 rows affected (0.00 sec)

mysql> CREATE TABLE y2k (date DATE,
->                        date_time DATETIME,
->                        time_stamp TIMESTAMP);
Query OK, 0 rows affected (0.01 sec)

mysql> INSERT INTO y2k VALUES
-> ('1998-12-31','1998-12-31 23:59:59','1998-12-31 23:59:59'),
-> ('1999-01-01','1999-01-01 00:00:00','1999-01-01 00:00:00'),
-> ('1999-09-09','1999-09-09 23:59:59','1999-09-09 23:59:59'),
-> ('2000-01-01','2000-01-01 00:00:00','2000-01-01 00:00:00'),
-> ('2000-02-28','2000-02-28 00:00:00','2000-02-28 00:00:00'),
-> ('2000-02-29','2000-02-29 00:00:00','2000-02-29 00:00:00'),
-> ('2000-03-01','2000-03-01 00:00:00','2000-03-01 00:00:00'),
-> ('2000-12-31','2000-12-31 23:59:59','2000-12-31 23:59:59'),
-> ('2001-01-01','2001-01-01 00:00:00','2001-01-01 00:00:00'),
-> ('2004-12-31','2004-12-31 23:59:59','2004-12-31 23:59:59'),
-> ('2005-01-01','2005-01-01 00:00:00','2005-01-01 00:00:00'),
-> ('2030-01-01','2030-01-01 00:00:00','2030-01-01 00:00:00'),
-> ('2040-01-01','2040-01-01 00:00:00','2040-01-01 00:00:00'),
-> ('9999-12-31','9999-12-31 23:59:59','9999-12-31 23:59:59');
Query OK, 14 rows affected, 2 warnings (0.00 sec)
Records: 14 Duplicates: 0 Warnings: 2

mysql> SELECT * FROM y2k;
```

date	date_time	time_stamp
1998-12-31	1998-12-31 23:59:59	1998-12-31 23:59:59
1999-01-01	1999-01-01 00:00:00	1999-01-01 00:00:00
1999-09-09	1999-09-09 23:59:59	1999-09-09 23:59:59
2000-01-01	2000-01-01 00:00:00	2000-01-01 00:00:00
2000-02-28	2000-02-28 00:00:00	2000-02-28 00:00:00
2000-02-29	2000-02-29 00:00:00	2000-02-29 00:00:00
2000-03-01	2000-03-01 00:00:00	2000-03-01 00:00:00
2000-12-31	2000-12-31 23:59:59	2000-12-31 23:59:59
2001-01-01	2001-01-01 00:00:00	2001-01-01 00:00:00
2004-12-31	2004-12-31 23:59:59	2004-12-31 23:59:59

```

+-----+-----+-----+
| 2005-01-01 | 2005-01-01 00:00:00 | 2005-01-01 00:00:00 |
| 2030-01-01 | 2030-01-01 00:00:00 | 2030-01-01 00:00:00 |
| 2040-01-01 | 2040-01-01 00:00:00 | 0000-00-00 00:00:00 |
| 9999-12-31 | 9999-12-31 23:59:59 | 0000-00-00 00:00:00 |
+-----+-----+-----+
14 rows in set (0.00 sec)

```

The final two `TIMESTAMP` column values are zero because the year values (2040, 9999) exceed the `TIMESTAMP` maximum. The `TIMESTAMP` data type, which is used to store the current time, supports values that range from '1970-01-01 00:00:00' to '2030-01-01 00:00:00' on 32-bit machines (signed value). On 64-bit machines, `TIMESTAMP` handles values up to 2106 (unsigned value).

Although MySQL Server itself is Y2K-safe, you may run into problems if you use it with applications that are not Y2K-safe. For example, many old applications store or manipulate years using two-digit values (which are ambiguous) rather than four-digit values. This problem may be compounded by applications that use values such as 00 or 99 as “missing” value indicators. Unfortunately, these problems may be difficult to fix because different applications may be written by different programmers, each of whom may use a different set of conventions and date-handling functions.

Thus, even though MySQL Server has no Y2K problems, *it is the application's responsibility to provide unambiguous input*. See [Section 11.3.4, “Y2K Issues and Date Types”](#), for MySQL Server's rules for dealing with ambiguous date input data that contains two-digit year values.

## 1.5. Overview of the MaxDB Database Management System

MaxDB is a heavy-duty enterprise database. The database management system is SAP-certified.

MaxDB is the new name of a database management system formerly called SAP DB. In 2003 SAP AG and MySQL AB joined a partnership and re-branded the database system to MaxDB. The development of MaxDB has continued since then as it was done before—through the SAP developer team.

MySQL AB cooperates closely with the MaxDB team at SAP around delivering improvements to the MaxDB product. Joint efforts include development of new native drivers to enable more efficient usage of MaxDB in the Open Source community, and improvement of documentation to expand the MaxDB user base. Interoperability features between MySQL and MaxDB database also are seen as important. For example, the new MaxDB Synchronization Manager supports data synchronization from MaxDB to MySQL.

The MaxDB database management system does not share a common code-base with the MySQL database management system. The MaxDB and MySQL database management systems are independent products provided by MySQL AB.

MySQL AB offers a complete portfolio of Professional Services for MaxDB.

### 1.5.1. What is MaxDB?

MaxDB is an ANSI SQL-92 (entry level) compliant relational database management system (RDBMS) from SAP AG, that is delivered by MySQL AB as well. MaxDB fulfills the needs for enterprise usage: safety, scalability, high concurrency, and performance. It runs on all major operating systems. Over the years it has proven able to run SAP R/3 and terabytes of data in 24×7 operation.

The database development started in 1977 as a research project at the Technical University of Berlin. In the early 1980s it became a database product that subsequently was owned by Nixdorf, Siemens Nixdorf, Software AG, and today by SAP AG. Along the way, it has been named VDN, Reflex, Supra 2, DDB/4, Entire SQL-DB-Server, and ADABAS D. In 1997, SAP took over the software from Software AG and renamed it to SAP DB. Since October 2000, SAP DB sources additionally were released as Open Source under the GNU General Public License (see [Appendix J, GNU General Public License](#)).

In 2003, SAP AG and MySQL AB formed a partnership and re-branded the database system to MaxDB.

## 1.5.2. History of MaxDB

The history of MaxDB goes back to SAP DB, SAP AG's DBMS. That is, MaxDB is a re-branded and enhanced version of SAP DB. For many years, MaxDB has been used for small, medium, and large installations of the mySAP Business Suite and other demanding SQL applications requiring an enterprise-class DBMS with regard to the number of users, the transactional workload, and the size of the database.

SAP DB was meant to provide an alternative to third-party database systems such as Oracle, Microsoft SQL Server, and DB2 by IBM. In October 2000, SAP AG released SAP DB under the GNU GPL license (see [Appendix J, GNU General Public License](#)), thus making it Open Source software.

Today, MaxDB is used in about 3,500 SAP customer installations worldwide. Moreover, the majority of all DBMS installations on Unix and Linux within SAP's IT department rely on MaxDB. MaxDB is tuned toward heavy-duty online transaction processing (OLTP) with several thousand users and database sizes ranging from several hundred GB to multiple TB.

In 2003, SAP and MySQL concluded a partnership and development cooperation agreement. As a result, SAP's database system SAP DB has been delivered under the name of MaxDB by MySQL since the release of version 7.5 (November 2003).

Version 7.5 of MaxDB is a direct advancement of the SAP DB 7.4 code base. Therefore, the MaxDB software version 7.5 can be used as a direct upgrade of previous SAP DB versions starting 7.2.04 and higher.

The former SAP DB development team at SAP AG is responsible, now as before, for developing and supporting MaxDB. MySQL AB cooperates closely with the MaxDB team at SAP around delivering improvements to the MaxDB product, see [Section 1.5, "Overview of the MaxDB Database Management System"](#). Both SAP AG and MySQL AB handle the sale and distribution of MaxDB. The advancement of MaxDB and the MySQL Server leverages synergies that benefit both product lines.

MaxDB is subjected to SAP AG's complete quality assurance process before it is shipped with SAP solutions or provided as a download from the MySQL site.

## 1.5.3. Features of MaxDB

MaxDB is a heavy-duty, SAP-certified Open Source database for OLTP and OLAP usage which offers high reliability, availability, scalability, and a very comprehensive feature set. It is targeted for large mySAP Business Suite environments and other applications that require maximum enterprise-level database functionality and complements the MySQL database server.

MaxDB operates as a client/server product. It was developed to meet the needs of installations in OLTP and Data Warehouse/OLAP/Decision Support scenarios and offers these benefits:

- **Easy configuration and administration:** GUI-based Installation Manager and Database Manager as single administration tools for DBMS operations
- **Around-the-clock operation, no planned downtimes, no permanent attendance required:** Automatic space management, no need for reorganizations
- **Sophisticated backup and restore capabilities:** Online and incremental backups, recovery wizard to guide you through the recovery scenario
- **Supports large number of users, database sizes in the terabytes, and demanding workloads:** Proven reliability, performance, and scalability

- **High availability:** Cluster support, standby configuration, hot standby configuration

## 1.5.4. Licensing and Support

MaxDB can be used under the same licenses available for the other products distributed by MySQL AB. Thus, MaxDB is available under the GNU General Public License, and a commercial license. For more information on licensing, see <http://www.mysql.com/company/legal/licensing/>.

MySQL AB offers MaxDB technical support to non-SAP customers. MaxDB support is available on various levels (Basic, Silver, and Gold), which expand from unlimited email/web-support to 24x7 phone support for business critical systems.

MySQL AB also offers Licenses and Support for MaxDB when used with SAP Applications, like SAP NetWeaver and mySAP Business Suite. For more information on licenses and support for your needs, please contact MySQL AB. (See <http://www.mysql.com/company/contact/>.)

Consulting and training services are available. MySQL gives classes on MaxDB at regular intervals. See <http://www.mysql.com/training/> for a list of classes.

## 1.5.5. Feature Differences Between MaxDB and MySQL

MaxDB is MySQL AB's SAP-certified database. The MaxDB database server complements the MySQL AB product portfolio. Some MaxDB features are not available on the MySQL database management server and vice versa.

The following list summarizes the main differences between MaxDB and MySQL; it is not complete.

- MaxDB runs as a client/server system. MySQL can run as a client/server system or as an embedded system.
- MaxDB might not run on all platforms supported by MySQL.
- MaxDB uses a proprietary network protocol for client/server communication. MySQL uses either TCP/IP (with or without SSL encryption), sockets (under Unix-like systems), or named pipes or shared memory (under Windows NT-family systems).
- MaxDB supports stored procedures and functions. MySQL 5.0 and up also supports stored procedures and functions. MaxDB supports programming of triggers through an SQL extension. MySQL 5.0 supports triggers. MaxDB contains a debugger for stored procedure languages, can cascade nested triggers, and supports multiple triggers per action and row.
- MaxDB is distributed with user interfaces that are text-based, graphical, or Web-based. MySQL is distributed with text-based user interfaces only; graphical user interfaces such as MySQL Query Browser or MySQL Administrator are shipped separately from the main distributions. Web-based user interfaces for MySQL are offered by third parties.
- MaxDB supports a number of programming interfaces that also are supported by MySQL. For developing with MaxDB, the MaxDB ODBC Driver, SQL Database Connectivity (SQLDBC), JDBC Driver, Perl and Python modules and a MaxDB PHP extension, which provides access to MySQL MaxDB databases using PHP, are available. Third Party Programming Interfaces: Support for OLE DB, ADO, DAO, RDO and .NET through ODBC. MaxDB supports embedded SQL with C/C++.
- MaxDB includes administrative features that MySQL does not have: job scheduling by time (included in MySQL as of 5.1), event, and alert, and sending messages to a database administrator on alert thresholds. (MySQL has scheduling support starting with version 5.1.6.)

## 1.5.6. Interoperability Features Between MaxDB and MySQL

MaxDB and MySQL are independent database management servers. The interoperation of the systems is possible in a way that the systems can exchange their data. To exchange data between MaxDB and MySQL, you can use the import and export tools of the systems or the MaxDB Synchronization Manager. The import and export tools can be used to transfer data in an infrequent, manual fashion. The MaxDB Synchronization Manager offers faster, automatic data transfer capabilities.

The MaxDB Loader can be used to export data and object definitions. The Loader can export data using MaxDB internal, binary formats and text formats (CSV). Data exported from MaxDB in text formats can be imported into MySQL using the `mysqlimport` client program. To export MySQL data, you can use either `mysqldump` to create `INSERT` statements or `SELECT ... INTO OUTFILE` to create a text file (CSV). Use the MaxDB Loader to import the data files generated by MySQL.

Object definitions can be exchanged between the systems using MaxDB Loader and the MySQL tool `mysqldump`. As the SQL dialects of both systems differ slightly and MaxDB has features currently not supported by MySQL like SQL constraints, we recommend to hand-tune the definition files. The `mysqldump` tool offers an option `--compatible=maxdb` to produce output that is compatible with MaxDB to make porting easier.

The MaxDB Synchronization Manager is available as part of MaxDB 7.6. The Synchronization Manager supports creation of asynchronous replication scenarios between several MaxDB instances. However, interoperability features also are planned, so that the Synchronization Manager supports replication to and from a MySQL server.

## 1.5.7. MaxDB-Related Links

The main page for MaxDB information is <http://www.mysql.com/products/maxdb>, which provides details about the features of the MaxDB database management systems and has pointers to available documentation.

The MySQL Reference Manual does not contain any MaxDB documentation other than the introduction given in this section. MaxDB has its own documentation, which is called the MaxDB library and is available at <http://dev.mysql.com/doc/maxdb/index.html>.

MySQL AB runs a community mailing list on MaxDB; see <http://lists.mysql.com/maxdb>. The list shows a vivid community discussion. Many of the core developers contribute to it. Product announcements are sent to the list.

A Web forum on MaxDB is available at <http://forums.mysql.com/>. The forum focuses on MaxDB questions not related to SAP applications.

## 1.6. MySQL Development Roadmap

This section provides a snapshot of the MySQL development roadmap, including major features implemented in or planned for various MySQL releases. The following sections provide information for each release series.

The current production release series is MySQL 5.0, which was declared stable for production use as of MySQL 5.0.15, released in October 2005. The previous production release series was MySQL 4.1, which was declared stable for production use as of MySQL 4.1.7, released in October 2004. “Production status” means that future 5.0 and 4.1 development is limited only to bugfixes. For the older MySQL 4.0 and 3.23 series, only critical bugfixes are made.

Active MySQL development currently is taking place in the MySQL 5.0 and 5.1 release series; and new features are being added only to the latter.

Before upgrading from one release series to the next, please see the notes in [Section 2.11, “Upgrading MySQL”](#).

The most requested features and the versions in which they were implemented or are scheduled for implementation are summarized in the following table:

Feature	MySQL Series
Foreign keys	3.23 (for the <a href="#">InnoDB</a> storage engine)
Unions	4.0
Subqueries	4.1
R-trees	4.1 (for the <a href="#">MyISAM</a> storage engine)
Stored procedures	5.0
Views	5.0
Cursors	5.0
XA transactions	5.0
Foreign keys	5.2 (implemented in 3.23 for <a href="#">InnoDB</a> )
Triggers	5.0 and 5.1
Partitioning	5.1
Row-Based Replication	5.1

## 1.6.1. MySQL 4.0 in a Nutshell

MySQL 4.0 is available for download at <http://dev.mysql.com/> and from our mirrors. MySQL 4.0 has been tested by a large number of users and is in production use at many large sites.

### 1.6.1.1. Features Available in MySQL 4.0

- Speed enhancements
  - MySQL 4.0 implemented a query cache that can give a major speed boost to applications with repetitive queries. See [Section 5.13, “The MySQL Query Cache”](#).
  - MySQL 4.0 further increased the speed of MySQL Server in a number of areas, such as bulk [INSERT](#) statements, searching on packed indexes, full-text searching (using [FULLTEXT](#) indexes), and [COUNT\(DISTINCT\)](#).
- Introduction of Embedded MySQL Server
  - The Embedded Server library added in this release can easily be used to create standalone and embedded applications. The embedded server provides an alternative to using MySQL in a client/server environment. See [Section 1.6.1.2, “The Embedded MySQL Server”](#).
- [InnoDB](#) storage engine as standard
  - The [InnoDB](#) storage engine began to be offered as a standard feature of the MySQL server. This provided full support for ACID transactions, foreign keys with cascading [UPDATE](#) and [DELETE](#), and row-level locking as standard features. See [Section 14.2, “The InnoDB Storage Engine”](#).
- New functionality
  - The enhanced [FULLTEXT](#) search capabilities of MySQL Server 4.0 enabled [FULLTEXT](#) index-



ing of large text masses with both binary and natural-language searching logic. It became possible to customize minimal word length and define your own stop word lists in most human languages, enabling a broader class of applications to be built with MySQL Server. See [Section 12.7, “Full-Text Search Functions”](#).

- Standards compliance, portability, and migration
  - MySQL Server added support for the `UNION` statement, a standard SQL feature.
  - Starting with version 4.0, MySQL runs natively on Novell NetWare 6.0 and higher. See [Section 2.7, “Installing MySQL on NetWare”](#).
  - Features to simplify migration from other database systems to MySQL Server include `TRUNCATE TABLE` (as in Oracle) .
- Internationalization
  - German-speaking users should note that MySQL 4.0 added support for a new character set, `latin1_de`, which ensures that words with umlauts are sorted in the same order as in German telephone books.
- **Usability enhancements**
  - As of version 4.0, most `mysqld` parameters (startup options) can be set without taking down the server. This is a convenient feature for database administrators. See [Section 13.5.3, “SET Syntax”](#).
  - Multiple-table `DELETE` and `UPDATE` statements were added.
  - On Windows, symbolic link handling at the database level was enabled by default. On Unix, the `MyISAM` storage engine added support for symbolic linking at the table level (and not just the database level as before).
  - The addition of the `SQL_CALC_FOUND_ROWS` and `FOUND_ROWS()` functions made it possible to find out the number of rows a `SELECT` query that includes a `LIMIT` clause would have returned without that clause.

The news section of this manual includes a more in-depth list of MySQL 4.0 features. See [Section D.2, “Changes in release 4.0.x \(Recent; still supported\)”](#).

### 1.6.1.2. The Embedded MySQL Server

The `libmysqld` embedded server library made MySQL Server suitable for a wider range of applications. Using this library, developers can embed MySQL Server into various applications and electronics devices, where the end user has no knowledge of there actually being an underlying database. Embedded MySQL Server is ideal for use in Internet appliances, public kiosks, turnkey hardware/software combination units, high performance Internet servers, self-contained databases distributed on CD-ROM, and so on.

The embedded MySQL library uses the same interface as the normal client library. See [Section 17.1, “libmysqld, the Embedded MySQL Server Library”](#). Embedded MySQL is available under the same dual-licensing model as the MySQL Server; see <http://www.mysql.com/company/legal/licensing/> for more information.

On Windows there are two different libraries:

<code>libmysqld.lib</code>	Dynamic library for threaded applications.
----------------------------	--

<code>mysqldemb.lib</code>	Static library for not threaded applications.
----------------------------	---

## 1.6.2. MySQL 4.1 in a Nutshell

MySQL Server 4.0 laid the foundation for new features implemented in MySQL 4.1, such as subqueries and Unicode support, which were desired by many of our customers.

MySQL Server 4.1 is currently in production status, and binaries are available for download at <http://dev.mysql.com/downloads/mysql/4.1.html>. All binary releases pass our extensive test suite without any errors on the platforms on which we test. See [Section D.1, “Changes in release 4.1.x \(Production\)”](#).

For those wishing to use the most recent development source for MySQL 4.1, we also make our BitKeeper repositories publicly available. See [Section 2.9.3, “Installing from the Development Source Tree”](#).

### 1.6.2.1. Features Available in MySQL 4.1

This section lists features implemented in MySQL 4.1. Features that are available in MySQL 5.0 are described in [Section 1.6.3, “What’s New in MySQL 5.0”](#).

- **Support for subqueries and derived tables:**
  - A “subquery” is a `SELECT` statement nested within another statement. A “derived table” (an unnamed view) is a subquery in the `FROM` clause of another statement. See [Section 13.2.8, “Subquery Syntax”](#).
- **Speed enhancements:**
  - Faster binary client/server protocol with support for prepared statements and parameter binding. See [Section 17.2.4, “C API Prepared Statements”](#).
  - `BTREE` indexing is supported for `HEAP` tables, significantly improving response time for non-exact searches.
- **Added functionality:**
  - `CREATE TABLE tbl_name2 LIKE tbl_name1` allows you to create, with a single statement, a new table with a structure exactly like that of an existing table.
  - The `MyISAM` storage engine added support for OpenGIS spatial types for storing geographical data. See [Chapter 16, \*Spatial Extensions\*](#).
  - Support was added for replication over SSL connections.
  - Support for a number of additional storage engines was implemented in the MySQL 4.1 release series:
    - The `EXAMPLE` storage engine is a “stub” engine that serves as an example in the MySQL source code for writing new storage engines, and is primarily of interest to developers. See [Section 14.6, “The `EXAMPLE` Storage Engine”](#).
    - `NDB Cluster` is the storage engine used by MySQL Cluster to implement tables that are partitioned over many computers. See [Chapter 15, \*MySQL Cluster\*](#).
    - The `ARCHIVE` storage engine is used for storing large amounts of data without indexes in a very small footprint. See [Section 14.7, “The `ARCHIVE` Storage Engine”](#).

- The `CSV` storage engine stores data in text files using comma-separated values format. See [Section 14.8, “The `CSV` Storage Engine”](#).
- The `BLACKHOLE` storage engine accepts but does not store data, and always returns an empty result set. It is for use primarily in replication. See [Section 14.9, “The `BLACKHOLE` Storage Engine”](#).

**Note:** These engine were implemented at different points in the development of MySQL 4.1. Please see the indicated sections for particulars in each case.

- **Standards compliance, portability, and migration:**
  - The enhanced client/server protocol available beginning with MySQL 4.1.1 provides the ability to pass multiple warnings to the client, rather than only a single result, making it much easier to track problems that occur in operations such as bulk data loading.
  - `SHOW WARNINGS` shows warnings for the last command. See [Section 13.5.4.21, “`SHOW WARNINGS` Syntax”](#).
- **Internationalization and Localization:**
  - To support applications that require the use of local languages, the MySQL software added extensive Unicode support through the `utf8` and `ucs2` character sets.
  - Definition of character sets by column, table, and database. This allows for a high degree of flexibility in application design, particularly for multi-language Web sites. See [Chapter 10, \*Character Set Support\*](#).
  - Per-connection time zones support, allowing individual clients to select their own time zones when necessary.
- **Usability enhancements:**
  - The addition of a server-based `HELP` command that can be used to get help information for SQL statements. This information is always applicable to the particular server version being used. Because this information is available by issuing an SQL statement, any client can access it. For example, the `help` command of the `mysql` command-line client has been modified to have this capability.
  - The improved client/server protocol allows multiple statements to be issued with a single call, and for returning multiple result sets. See [Section 17.2.9, “C API Handling of Multiple Statement Execution”](#).
  - The syntax `INSERT ... ON DUPLICATE KEY UPDATE ...` was implemented. This allows you to update an existing row if the insert would have caused a duplicate value for a primary or unique index. See [Section 13.2.4, “`INSERT` Syntax”](#).
  - The aggregate function `GROUP_CONCAT( )`, added the capability to concatenate column values from grouped rows into a single result string. See [Section 12.10, “Functions and Modifiers for Use with `GROUP BY` Clauses”](#).

The News section of this manual includes a more in-depth list of MySQL 4.1 features. See [Section D.1, “Changes in release 4.1.x \(Production\)”](#).

### 1.6.3. What's New in MySQL 5.0

The following features are implemented in MySQL 5.0.

- **BIT Data Type:** Can be used to store numbers in binary notation.
- **Cursors:** Elementary support for server-side cursors.
- **Data Dictionary (Information Schema):** The introduction of the `INFORMATION_SCHEMA` database in MySQL 5.0 provided a standards-compliant means for accessing the MySQL Server's metadata; that is, data about the databases (schemas) on the server and the objects which they contain.
- **Instance Manager:** Can be used to start and stop the MySQL Server, even from a remote host.
- **Precision Math:** MySQL 5.0 introduced stricter criteria for acceptance or rejection of data, and implemented a new library for fixed-point arithmetic. These contributed to a much higher degree of accuracy for mathematical operations and greater control over invalid values.
- **Storage Engines:** Storage engines added in MySQL 5.0 include `ARCHIVE` and `FEDERATED`.
- **Stored Routines:** Support for named stored procedures and stored functions was implemented in MySQL 5.0.
- **Strict Mode and Standard Error Handling:** MySQL 5.0 added a strict mode where by it follows standard SQL in a number of ways in which it did not previously. Support for standard SQLSTATE error messages was also implemented.
- **Triggers:** MySQL 5.0 added limited support for triggers.
- **VARCHAR Data Type:** The maximum effective length of a `VARCHAR` column was increased to 65,532 bytes, and stripping of trailing whitespace was eliminated.
- **Views:** MySQL 5.0 added support for named, updatable views.

For those wishing to take a look at the bleeding edge of MySQL development, we make our BitKeeper repository for MySQL publicly available. See [Section 2.9.3, "Installing from the Development Source Tree"](#).

## 1.7. MySQL Information Sources

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

### 1.7.1. MySQL Mailing Lists

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

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

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

Your local site may have many subscribers to a MySQL mailing list. If so, the site may have a local

mailing list, so that messages sent from [lists.mysql.com](mailto:lists.mysql.com) to your site are propagated to the local list. In such cases, please contact your system administrator to be added to or dropped from the local MySQL list.

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

The MySQL mailing lists are as follows:

- [announce](#)

This list is for announcements of new versions of MySQL and related programs. This is a low-volume list to which all MySQL users should subscribe.
- [mysql](#)

This is the main list for general MySQL discussion. Please note that some topics are better discussed on the more-specialized lists. If you post to the wrong list, you may not get an answer.
- [bugs](#)

This list is for people who want to stay informed about issues reported since the last release of MySQL or who want to be actively involved in the process of bug hunting and fixing. See [Section 1.8, “How to Report Bugs or Problems”](#).
- [internals](#)

This list is for people who work on the MySQL code. This is also the forum for discussions on MySQL development and for posting patches.
- [mysqldoc](#)

This list is for people who work on the MySQL documentation: people from MySQL AB, translators, and other community members.
- [benchmarks](#)

This list is for anyone interested in performance issues. Discussions concentrate on database performance (not limited to MySQL), but also include broader categories such as performance of the kernel, filesystem, disk system, and so on.
- [packagers](#)

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

This list is for discussions about the MySQL server and Java. It is mostly used to discuss JDBC drivers such as MySQL Connector/J.
- [win32](#)

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

- [myodbc](#)

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

- [gui-tools](#)

This list is for all topics concerning MySQL graphical user interface tools such as [MySQL Administrator](#) and [MySQL Query Browser](#).

- [cluster](#)

This list is for discussion of MySQL Cluster.

- [dotnet](#)

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

- [plusplus](#)

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

- [perl](#)

This list is for all topics concerning Perl support for MySQL with [DBD: :mysql](#).

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

The following table shows some MySQL mailing lists in languages other than English. These lists are not operated by MySQL AB.

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

A French mailing list.

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

A Korean mailing list. To subscribe, email [subscribe mysql your@email.address](#) to this list.

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

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

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

A Portuguese mailing list. To subscribe, email [subscribe mysql-br your@email.address](#) to this list.

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

A Spanish mailing list. To subscribe, email [subscribe mysql your@email.address](#) to this list.

### 1.7.1.1. Guidelines for Using the Mailing Lists

Please don't post mail messages from your browser with HTML mode turned on. Many users don't read mail with a browser.

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

Try to summarize the essential part of the question in your reply. Don't feel obliged to quote the entire original message.

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

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

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

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

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

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

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

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

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

## 1.8. How to Report Bugs or Problems

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

- Start by searching the MySQL online manual at <http://dev.mysql.com/doc/>. We try to keep the manual up to date by updating it frequently with solutions to newly found problems. The change history (<http://dev.mysql.com/doc/mysql/en/news.html>) can be particularly useful since it is quite possible that a newer version contains a solution to your problem.
- If you get a parse error for a SQL statement, please check your syntax closely. If you can't find something wrong with it, it's extremely likely that your current version of MySQL Server doesn't support the syntax you are using. If you are using the current version and the manual doesn't cover the syntax that you are using, MySQL Server doesn't support your statement. In this case, your options are to implement the syntax yourself or email [<licensing@mysql.com>](mailto:licensing@mysql.com) and ask for an offer to implement it.

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

- For solutions to some common problems, see [Appendix A, Problems and Common Errors](#).
- Search the bugs database at <http://bugs.mysql.com/> to see whether the bug has been reported and fixed.
- Search the MySQL mailing list archives at <http://lists.mysql.com/>. See [Section 1.7.1, “MySQL Mailing Lists”](#).
- You can also use <http://www.mysql.com/search/> to search all the Web pages (including the manual) that are located at the MySQL AB Web site.

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

The normal way to report bugs is to visit <http://bugs.mysql.com/>, which is the address for our bugs database. This database is public and can be browsed and searched by anyone. If you log in to the system, you can enter new reports. If you have no Web access, you can generate a bug report by using the `mysqlbug` script described at the end of this section.

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

If you have found a sensitive security bug in MySQL, you can send email to [<security@mysql.com>](mailto:security@mysql.com).

To discuss problems with other users, you can use one of the MySQL mailing lists. [Section 1.7.1, “MySQL Mailing Lists”](#).

Writing a good bug report takes patience, but doing it right the first time saves time both for us and for yourself. A good bug report, containing a full test case for the bug, makes it very likely that we will fix



the bug in the next release. This section helps you write your report correctly so that you don't waste your time doing things that may not help us much or at all. Please read this section carefully and make sure that all the information described here is included in your report.

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

It is most helpful when a good description of the problem is included in the bug report. That is, give a good example of everything you did that led to the problem and describe, in exact detail, the problem itself. The best reports are those that include a full example showing how to reproduce the bug or problem. See [Section E.1.6, “Making a Test Case If You Experience Table Corruption”](#).

Remember that it is possible for us to respond to a report containing too much information, but not to one containing too little. People often omit facts because they think they know the cause of a problem and assume that some details don't matter. A good principle to follow is that if you are in doubt about stating something, state it. It is faster and less troublesome to write a couple more lines in your report than to wait longer for the answer if we must ask you to provide information that was missing from the initial report.

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

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

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

If you have a problem with Connector/ODBC (MyODBC), please try to generate a trace file and send it with your report. See the MyODBC section of [Chapter 18, \*Connectors\*](#).

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

Please include the following information in your report:

- The version number of the MySQL distribution you are using (for example, MySQL 5.0.19). You can find out which version you are running by executing `mysqladmin version`. The `mysql-admin` program can be found in the `bin` directory under your MySQL installation directory.
- The manufacturer and model of the machine on which you experience the problem.
- The operating system name and version. If you work with Windows, you can usually get the name and version number by double-clicking your My Computer icon and pulling down the “Help/About

Windows” menu. For most Unix-like operating systems, you can get this information by executing the command `uname -a`.

- Sometimes the amount of memory (real and virtual) is relevant. If in doubt, include these values.
- If you are using a source distribution of the MySQL software, include the name and version number of the compiler that you used. If you have a binary distribution, include the distribution name.
- If the problem occurs during compilation, include the exact error messages and also a few lines of context around the offending code in the file where the error occurs.
- If `mysqld` died, you should also report the statement that crashed `mysqld`. You can usually get this information by running `mysqld` with query logging enabled, and then looking in the log after `mysqld` crashes. See [Section E.1.5, “Using Server Logs to Find Causes of Errors in `mysqld`”](#).
- If a database table is related to the problem, include the output from the `SHOW CREATE TABLE db_name.tbl_name` statement in the bug report. This is a very easy way to get the definition of any table in a database. The information helps us create a situation matching the one that you have experienced.
- For performance-related bugs or problems with `SELECT` statements, you should always include the output of `EXPLAIN SELECT ...`, and at least the number of rows that the `SELECT` statement produces. You should also include the output from `SHOW CREATE TABLE tbl_name` for each table that is involved. The more information you provide about your situation, the more likely it is that someone can help you.

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

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

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

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

- If you can't produce a test case with only a few rows, or if the test table is too big to be included in the bug report (more than 10 rows), you should dump your tables using `mysqldump` and create a `README` file that describes your problem. Create a compressed archive of your files using `tar` and `gzip` or `zip`, and use FTP to transfer the archive to <ftp://ftp.mysql.com/pub/mysql/upload/>. Then enter the problem into our bugs database at <http://bugs.mysql.com/>.
- If you believe that the MySQL server produces a strange result from a statement, include not only the result, but also your opinion of what the result should be, and an explanation describing the basis for your opinion.

- When you provide an example of the problem, it's better to use the table names, variable names, and so forth that exist in your actual situation than to come up with new names. The problem could be related to the name of a table or variable. These cases are rare, perhaps, but it is better to be safe than sorry. After all, it should be easier for you to provide an example that uses your actual situation, and it is by all means better for us. If you have data that you don't want to be visible to others in the bug report, you can use FTP to transfer it to <ftp://ftp.mysql.com/pub/mysql/upload/>. If the information is really top secret and you don't want to show it even to us, go ahead and provide an example using other names, but please regard this as the last choice.
- Include all the options given to the relevant programs, if possible. For example, indicate the options that you use when you start the `mysqld` server, as well as the options that you use to run any MySQL client programs. The options to programs such as `mysqld` and `mysql`, and to the `configure` script, are often key to resolving problems and are very relevant. It is never a bad idea to include them. If your problem involves a program written in a language such as Perl or PHP, please include the language processor's version number, as well as the version for any modules that the program uses. For example, if you have a Perl script that uses the `DBI` and `DBD: :mysql` modules, include the version numbers for Perl, `DBI`, and `DBD: :mysql`.
- If your question is related to the privilege system, please include the output of `mysqlaccess`, the output of `mysqladmin reload`, and all the error messages you get when trying to connect. When you test your privileges, you should first run `mysqlaccess`. After this, execute `mysqladmin reload version` and try to connect with the program that gives you trouble. `mysqlaccess` can be found in the `bin` directory under your MySQL installation directory.
- If you have a patch for a bug, do include it. But don't assume that the patch is all we need, or that we can use it, if you don't provide some necessary information such as test cases showing the bug that your patch fixes. We might find problems with your patch or we might not understand it at all. If so, we can't use it.

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

- Guesses about what the bug is, why it occurs, or what it depends on are usually wrong. Even the MySQL team can't guess such things without first using a debugger to determine the real cause of a bug.
- Indicate in your bug report that you have checked the reference manual and mail archive so that others know you have tried to solve the problem yourself.
- If the problem is that your data appears corrupt or you get errors when you access a particular table, you should first check your tables and then try to repair them with `CHECK TABLE` and `REPAIR TABLE` or with `myisamchk`. See [Chapter 5, Database Administration](#).

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

- If you often get corrupted tables, you should try to find out when and why this happens. In this case, the error log in the MySQL data directory may contain some information about what happened. (This is the file with the `.err` suffix in the name.) See [Section 5.11.1, "The Error Log"](#). Please include any relevant information from this file in your bug report. Normally `mysqld` should *never* crash a table if nothing killed it in the middle of an update. If you can find the cause of `mysqld` dying, it's much easier for us to provide you with a fix for the problem. See [Section A.1, "How to Determine What Is Causing a Problem"](#).
- If possible, download and install the most recent version of MySQL Server and check whether it

solves your problem. All versions of the MySQL software are thoroughly tested and should work without problems. We believe in making everything as backward-compatible as possible, and you should be able to switch MySQL versions without difficulty. See [Section 2.1.2, “Choosing Which MySQL Distribution to Install”](#).

If you have no Web access and cannot report a bug by visiting <http://bugs.mysql.com/>, you can use the `mysqlbug` script to generate a bug report (or a report about any problem). `mysqlbug` helps you generate a report by determining much of the following information automatically, but if something important is missing, please include it with your message. `mysqlbug` can be found in the `scripts` directory (source distribution) and in the `bin` directory under your MySQL installation directory (binary distribution).

## 1.9. MySQL Standards Compliance

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

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

One of our main goals with the product is to continue to work toward compliance with the SQL standard, but without sacrificing speed or reliability. We are not afraid to add extensions to SQL or support for non-SQL features if this greatly increases the usability of MySQL Server for a large segment of our user base. The `HANDLER` interface is an example of this strategy. See [Section 13.2.3, “HANDLER Syntax”](#).

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

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

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

In MySQL 4.1.2 in later, high-availability database clustering is supported by the `NDBCluster` storage engine. See [Chapter 15, \*MySQL Cluster\*](#).

XML support is to be implemented in a future version of the database server.

### 1.9.1. What Standards MySQL Follows

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

ODBC levels 0-3.51.

### 1.9.2. Selecting SQL Modes

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

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

You can set the default SQL mode by starting `mysqld` with the `--sql-mode="mode_value"` option. Beginning with MySQL 4.1, you can also change the mode at runtime by setting the `sql_mode` system variable with a `SET [SESSION|GLOBAL] sql_mode='mode_value'` statement.

For more information on setting the SQL mode, see [Section 5.2.5, “SQL Modes”](#).

### 1.9.3. Running MySQL in ANSI Mode

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

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

As of MySQL 4.1.1, you can achieve the same effect at runtime by executing these two statements:

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

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

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

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

See [Section 5.2.1, “Command Options”](#), and [Section 1.9.2, “Selecting SQL Modes”](#).

### 1.9.4. MySQL Extensions to Standard SQL

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

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

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

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

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

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

The following descriptions list MySQL extensions, organized by category.

- Organization of data on disk

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

- Database and table names are case sensitive in MySQL Server on operating systems that have case-sensitive filenames (such as most Unix systems). See [Section 9.2.2, “Identifier Case Sensitivity”](#).
- You can use standard system commands to back up, rename, move, delete, and copy tables that are managed by the [MyISAM](#) or [ISAM](#) storage engines. For example, it is possible to rename a [MyISAM](#) table by renaming the `.MYD`, `.MYI`, and `.frm` files to which the table corresponds. (Nevertheless, it is preferable to use `RENAME TABLE` or `ALTER TABLE ... RENAME` and let the server rename the files.)

Database and table names cannot contain pathname separator characters (`/`, `\`).

- General language syntax

- By default, strings can be enclosed by either `"` or `'`, not just by `'`. (If the [ANSI\\_QUOTES](#) SQL mode is enabled, strings can be enclosed only by `'` and the server interprets strings enclosed by `"` as identifiers.)
- `\` is the escape character in strings.
- In SQL statements, you can access tables from different databases with the `db_name.tbl_name` syntax. Some SQL servers provide the same functionality but call this `User space`. MySQL Server doesn't support tablespaces such as used in statements like this: `CREATE TABLE ralph.my_table ... IN my_tablespace`.

- SQL statement syntax

- The `ANALYZE TABLE`, `CHECK TABLE`, `OPTIMIZE TABLE`, and `REPAIR TABLE` statements.
- The `CREATE DATABASE`, `DROP DATABASE`, and `ALTER DATABASE` statements. See [Section 13.1.3, “CREATE DATABASE Syntax”](#), [Section 13.1.6, “DROP DATABASE Syntax”](#), and [Section 13.1.1, “ALTER DATABASE Syntax”](#).
- The `DO` statement.
- `EXPLAIN SELECT` to obtain a description of how tables are processed by the query optimizer.
- The `FLUSH` and `RESET` statements.
- The `SET` statement. See [Section 13.5.3, “SET Syntax”](#).
- The `SHOW` statement. See [Section 13.5.4, “SHOW Syntax”](#).
- Use of `LOAD DATA INFILE`. In many cases, this syntax is compatible with Oracle's `LOAD DATA INFILE`. See [Section 13.2.5, “LOAD DATA INFILE Syntax”](#).
- Use of `RENAME TABLE`. See [Section 13.1.9, “RENAME TABLE Syntax”](#).

- Use of `REPLACE` instead of `DELETE` plus `INSERT`. See [Section 13.2.6, “REPLACE Syntax”](#).
- Use of `CHANGE col_name`, `DROP col_name`, or `DROP INDEX`, `IGNORE` or `RENAME` in `ALTER TABLE` statements. Use of multiple `ADD`, `ALTER`, `DROP`, or `CHANGE` clauses in an `ALTER TABLE` statement. See [Section 13.1.2, “ALTER TABLE Syntax”](#).
- Use of index names, indexes on a prefix of a column, and use of `INDEX` or `KEY` in `CREATE TABLE` statements. See [Section 13.1.5, “CREATE TABLE Syntax”](#).
- Use of `TEMPORARY` or `IF NOT EXISTS` with `CREATE TABLE`.
- Use of `IF EXISTS` with `DROP TABLE` and `DROP DATABASE`.
- The capability of dropping multiple tables with a single `DROP TABLE` statement.
- The `ORDER BY` and `LIMIT` clauses of the `UPDATE` and `DELETE` statements.
- `INSERT INTO tbl_name SET col_name = ...` syntax.
- The `DELAYED` clause of the `INSERT` and `REPLACE` statements.
- The `LOW_PRIORITY` clause of the `INSERT`, `REPLACE`, `DELETE`, and `UPDATE` statements.
- Use of `INTO OUTFILE` or `INTO DUMPFILE` in `SELECT` statements. See [Section 13.2.7, “SELECT Syntax”](#).
- Options such as `STRAIGHT_JOIN` or `SQL_SMALL_RESULT` in `SELECT` statements.
- You don't need to name all selected columns in the `GROUP BY` clause. This gives better performance for some very specific, but quite normal queries. See [Section 12.10, “Functions and Modifiers for Use with GROUP BY Clauses”](#).
- You can specify `ASC` and `DESC` with `GROUP BY`, not just with `ORDER BY`.
- The ability to set variables in a statement with the `:=` assignment operator:

```
mysql> SELECT @a:=SUM(total),@b=COUNT(*),@a/@b AS avg
-> FROM test_table;
mysql> SELECT @t1:=(@t2:=1)+@t3:=4,@t1,@t2,@t3;
```

- Data types
  - The `MEDIUMINT`, `SET`, and `ENUM` data types, and the various `BLOB` and `TEXT` data types.
  - The `AUTO_INCREMENT`, `BINARY`, `NULL`, `UNSIGNED`, and `ZEROFILL` data type attributes.
- Functions and operators
  - To make it easier for users who migrate from other SQL environments, MySQL Server supports aliases for many functions. For example, all string functions support both standard SQL syntax and ODBC syntax.
  - MySQL Server understands the `||` and `&&` operators to mean logical OR and AND, as in the C programming language. In MySQL Server, `||` and `OR` are synonyms, as are `&&` and `AND`. Because of this nice syntax, MySQL Server doesn't support the standard SQL `||` operator for string concatenation; use `CONCAT()` instead. Because `CONCAT()` takes any number of arguments, it's easy to convert use of the `||` operator to MySQL Server.
  - Use of `COUNT(DISTINCT value_list)` where `value_list` has more than one element.

- String comparisons are case-insensitive by default, with sort ordering determined by collation of the current character set, which is `latin1` (cp1252 West European) by default. If you don't like this, you should declare your columns with the `BINARY` attribute or use the `BINARY` cast, which causes comparisons to be done using the underlying character code values rather than a lexical ordering.
- The `%` operator is a synonym for `MOD()`. That is, `N % M` is equivalent to `MOD(N,M)`. `%` is supported for C programmers and for compatibility with PostgreSQL.
- The `=`, `<>`, `<=<`, `>=>`, `<<`, `>>`, `<=>`, `AND`, `OR`, or `LIKE` operators may be used in expressions in the output column list (to the left of the `FROM`) in `SELECT` statements. For example:

```
mysql> SELECT col1=1 AND col2=2 FROM my_table;
```

- The `LAST_INSERT_ID()` function returns the most recent `AUTO_INCREMENT` value. See [Section 12.9.3, “Information Functions”](#).
- `LIKE` is allowed on numeric values.
- The `REGEXP` and `NOT REGEXP` extended regular expression operators.
- `CONCAT()` or `CHAR()` with one argument or more than two arguments. (In MySQL Server, these functions can take a variable number of arguments.)
- The `BIT_COUNT()`, `CASE`, `ELT()`, `FROM_DAYS()`, `FORMAT()`, `IF()`, `PASSWORD()`, `ENCRYPT()`, `MD5()`, `ENCODE()`, `DECODE()`, `PERIOD_ADD()`, `PERIOD_DIFF()`, `TO_DAYS()`, and `WEEKDAY()` functions.
- Use of `TRIM()` to trim substrings. Standard SQL supports removal of single characters only.
- The `GROUP BY` functions `STD()`, `BIT_OR()`, `BIT_AND()`, `BIT_XOR()`, and `GROUP_CONCAT()`. See [Section 12.10, “Functions and Modifiers for Use with GROUP BY Clauses”](#).

For a prioritized list indicating when new extensions are added to MySQL Server, you should consult the online MySQL development roadmap at <http://dev.mysql.com/doc/mysql/en/roadmap.html>.

## 1.9.5. MySQL Differences from Standard SQL

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

- For `VARCHAR` columns, trailing spaces are removed when the value is stored. See [Section A.8, “Known Issues in MySQL”](#).
- In some cases, `CHAR` columns are silently converted to `VARCHAR` columns when you define a table or alter its structure. See [Section 13.1.5.1, “Silent Column Specification Changes”](#).
- There are several differences between the MySQL and standard SQL privilege systems. For example, in MySQL, privileges for a table are not automatically revoked when you delete a table. You must explicitly issue a `REVOKE` statement to revoke privileges for a table. For more information, see [Section 13.5.1.3, “REVOKE Syntax”](#).
- The `CAST()` function does not support cast to `REAL` or `BIGINT`. See [Section 12.8, “Cast Functions and Operators”](#).



### 1.9.5.1. Subquery Support

MySQL 4.1 and up supports subqueries and derived tables. A “subquery” is a `SELECT` statement nested within another statement. A “derived table” (an unnamed view) is a subquery in the `FROM` clause of another statement. See [Section 13.2.8, “Subquery Syntax”](#).

For MySQL versions older than 4.1, most subqueries can be rewritten using joins or other methods. See [Section 13.2.8.11, “Rewriting Subqueries as Joins for Earlier MySQL Versions”](#), for examples that show how to do this.

### 1.9.5.2. `SELECT INTO TABLE`

MySQL Server doesn't support the `SELECT ... INTO TABLE` Sybase SQL extension. Instead, MySQL Server supports the `INSERT INTO ... SELECT` standard SQL syntax, which is basically the same thing. See [Section 13.2.4.1, “INSERT ... SELECT Syntax”](#). For example:

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

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

### 1.9.5.3. Transactions and Atomic Operations

MySQL Server (version 3.23-max and all versions 4.0 and above) supports transactions with the `InnoDB` and `BDB` transactional storage engines. `InnoDB` provides *full ACID* compliance. MySQL Cluster is also a transaction-safe storage engine. See [Chapter 14, \*Storage Engines and Table Types\*](#). For information about `InnoDB` differences from standard SQL with regard to treatment of transaction errors, see [Section 14.2.16, “InnoDB Error Handling”](#).

The other non-transactional storage engines in MySQL Server (such as `MyISAM`) follow a different paradigm for data integrity called “atomic operations.” In transactional terms, `MyISAM` tables effectively always operate in `AUTOCOMMIT=1` mode. Atomic operations often offer comparable integrity with higher performance.

Because MySQL Server supports both paradigms, you can decide whether your applications are best served by the speed of atomic operations or the use of transactional features. This choice can be made on a per-table basis.

As noted, the trade-off for transactional versus non-transactional storage engines lies mostly in performance. Transactional tables have significantly higher memory and disk space requirements, and more CPU overhead. On the other hand, transactional storage engines such as `InnoDB` also offer many significant features. MySQL Server's modular design allows the concurrent use of different storage engines to suit different requirements and deliver optimum performance in all situations.

But how do you use the features of MySQL Server to maintain rigorous integrity even with the non-transactional `MyISAM` tables, and how do these features compare with the transactional storage engines?

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

If you use non-transactional tables, MySQL Server in almost all cases allows you to resolve potential problems by including simple checks before updates and by running simple scripts that check the databases for inconsistencies and automatically repair or warn if such an inconsistency occurs. Note that just by using the MySQL log or even adding one extra log, you can normally fix tables perfectly

with no data integrity loss.

- More often than not, critical transactional updates can be rewritten to be atomic. Generally speaking, all integrity problems that transactions solve can be done with `LOCK TABLES` or atomic updates, ensuring that there are no automatic aborts from the server, which is a common problem with transactional database systems.
- To be safe with MySQL Server, regardless of whether you use transactional tables, you only need to have backups and have binary logging turned on. When that is true, you can recover from any situation that you could with any other transactional database system. It is always good to have backups, regardless of which database system you use.

The transactional paradigm has its benefits and its drawbacks. Many users and application developers depend on the ease with which they can code around problems where an abort appears to be necessary, or is necessary. However, even if you are new to the atomic operations paradigm, or more familiar with transactions, do consider the speed benefit that non-transactional tables can offer on the order of three to five times the speed of the fastest and most optimally tuned transactional tables.

In situations where integrity is of highest importance, MySQL Server offers transaction-level reliability and integrity even for non-transactional tables. If you lock tables with `LOCK TABLES`, all updates stall until integrity checks are made. If you obtain a `READ LOCAL` lock (as opposed to a write lock) for a table that allows concurrent inserts at the end of the table, reads are allowed, as are inserts by other clients. The newly inserted records are not be seen by the client that has the read lock until it releases the lock. With `INSERT DELAYED`, you can write inserts that go into a local queue until the locks are released, without having the client wait for the insert to complete. See [Section 7.3.3, “Concurrent Inserts”](#), and [Section 13.2.4.2, “INSERT DELAYED Syntax”](#).

“Atomic,” in the sense that we mean it, is nothing magical. It only means that you can be sure that while each specific update is running, no other user can interfere with it, and there can never be an automatic rollback (which can happen with transactional tables if you are not very careful). MySQL Server also guarantees that there are no dirty reads.

Following are some techniques for working with non-transactional tables:

- Loops that need transactions normally can be coded with the help of `LOCK TABLES`, and you don't need cursors to update records on the fly.
- To avoid using `ROLLBACK`, you can employ the following strategy:
  1. Use `LOCK TABLES` to lock all the tables you want to access.
  2. Test the conditions that must be true before performing the update.
  3. Update if the conditions are satisfied.
  4. Use `UNLOCK TABLES` to release your locks.

This is usually a much faster method than using transactions with possible rollbacks, although not always. The only situation this solution doesn't handle is when someone kills the threads in the middle of an update. In that case, all locks are released but some of the updates may not have been executed.

- You can also use functions to update records in a single operation. You can get a very efficient application by using the following techniques:
  - Modify columns relative to their current value.

- Update only those columns that actually have changed.

For example, when we are updating customer information, we update only the customer data that has changed and test only that none of the changed data, or data that depends on the changed data, has changed compared to the original row. The test for changed data is done with the `WHERE` clause in the `UPDATE` statement. If the record wasn't updated, we give the client a message: "Some of the data you have changed has been changed by another user." Then we show the old row versus the new row in a window so that the user can decide which version of the customer record to use.

This gives us something that is similar to column locking but is actually even better because we only update some of the columns, using values that are relative to their current values. This means that typical `UPDATE` statements look something like these:

```
UPDATE tablename SET pay_back=pay_back+125;

UPDATE customer
  SET
    customer_date='current_date',
    address='new address',
    phone='new phone',
    money_owed_to_us=money_owed_to_us-125
  WHERE
    customer_id=id AND address='old address' AND phone='old phone';
```

This is very efficient and works even if another client has changed the values in the `pay_back` or `money_owed_to_us` columns.

- In many cases, users have wanted `LOCK TABLES` or `ROLLBACK` for the purpose of managing unique identifiers. This can be handled much more efficiently without locking or rolling back by using an `AUTO_INCREMENT` column and either the `LAST_INSERT_ID()` SQL function or the `mysql_insert_id()` C API function. See [Section 12.9.3, "Information Functions"](#), and [Section 17.2.3.35, "mysql\\_insert\\_id\(\)"](#).

You can generally code around the need for row-level locking. Some situations really do need it, and `InnoDB` tables support row-level locking. Otherwise, with `MyISAM` tables, you can use a flag column in the table and do something like the following:

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

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

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

### 1.9.5.4. Stored Routines and Triggers

Stored procedures and functions are implemented beginning with MySQL 5.0.

Basic trigger functionality is implemented beginning with MySQL 5.0.2, with further development planned for MySQL 5.1.

### 1.9.5.5. Foreign Keys

In MySQL Server 3.23.44 and up, the `InnoDB` storage engine supports checking of foreign key constraints, including `CASCADE`, `ON DELETE`, and `ON UPDATE`. See [Section 14.2.7.4, "FOREIGN KEY Constraints"](#).

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

Foreign key enforcement offers several benefits to database developers:

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

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

MySQL gives database developers the choice of which approach to use. If you don't need foreign keys and want to avoid the overhead associated with enforcing referential integrity, you can choose another storage engine instead, such as [MyISAM](#). (For example, the [MyISAM](#) storage engine offers very fast performance for applications that perform only [INSERT](#) and [SELECT](#) operations. In this case, the table has no holes in the middle and the inserts can be performed concurrently with retrievals. See [Section 7.3.3](#), “[Concurrent Inserts](#)”.)

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

- In the absence of server-side foreign key relationship checking, the application itself must handle relationship issues. For example, it must take care to insert rows into tables in the proper order, and to avoid creating orphaned child records. It must also be able to recover from errors that occur in the middle of multiple-record insert operations.
- If [ON DELETE](#) is the only referential integrity capability an application needs, you can achieve a similar effect as of MySQL Server 4.0 by using multiple-table [DELETE](#) statements to delete rows from many tables with a single statement. See [Section 13.2.1](#), “[DELETE Syntax](#)”.
- A workaround for the lack of [ON DELETE](#) is to add the appropriate [DELETE](#) statements to your application when you delete records from a table that has a foreign key. In practice, this is often as quick as using foreign keys and is more portable.

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

- Foreign key support addresses many referential integrity issues, but it is still necessary to design key relationships carefully to avoid circular rules or incorrect combinations of cascading deletes.
- It is not uncommon for a DBA to create a topology of relationships that makes it difficult to restore individual tables from a backup. (MySQL alleviates this difficulty by allowing you to temporarily disable foreign key checks when reloading a table that depends on other tables. See [Section 14.2.7.4](#),

“[FOREIGN KEY Constraints](#)”. As of MySQL 4.1.1, `mysqldump` generates dump files that take advantage of this capability automatically when they are reloaded.)

Note that foreign keys in SQL are used to check and enforce referential integrity, not to join tables. If you want to get results from multiple tables from a `SELECT` statement, you do this by performing a join between them:

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

See [Section 13.2.7.1, “JOIN Syntax”](#), and [Section 3.6.6, “Using Foreign Keys”](#).

The `FOREIGN KEY` syntax without `ON DELETE . . .` is often used by ODBC applications to produce automatic `WHERE` clauses.

### 1.9.5.6. Views

Views (including updatable views) are implemented beginning with MySQL Server 5.0.1.

Views are useful for allowing users to access a set of relations (tables) as if it were a single table, and limiting their access to just that. Views can also be used to restrict access to rows (a subset of a particular table). For access control to columns, you can also use the sophisticated privilege system in MySQL Server. See [Section 5.7, “The MySQL Access Privilege System”](#).

In designing an implementation of views, our ambitious goal, as much as is possible within the confines of SQL, has been full compliance with “Codd’s Rule #6” for relational database systems: “All views that are theoretically updatable, should in practice also be updatable.”

### 1.9.5.7. ‘--’ as the Start of a Comment

Standard SQL uses the C syntax `/* this is a comment */` for comments, and MySQL Server supports this syntax as well. MySQL also support extensions to this syntax that allow MySQL-specific SQL to be embedded in the comment, as described in [Section 9.4, “Comment Syntax”](#).

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

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

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

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

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

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

The statement produces no change in value at all. This illustrates that allowing comments to start with ‘--’ can have serious consequences.

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

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

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

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

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

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

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

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

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

Change them back with this command:

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

See [Section 8.18](#), “`replace` — A String-Replacement Utility”.

## 1.9.6. How MySQL Deals with Constraints

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

The basic philosophy is that MySQL Server tries to produce an error for anything that it can detect while parsing a statement to be executed, and tries to recover from any errors that occur while executing the statement. We do this in most cases.

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

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

### 1.9.6.1. PRIMARY KEY and UNIQUE Index Constraints

Normally, an error occurs when you try to `INSERT` or `UPDATE` a row that causes a primary key, unique key, or foreign key violation. If you are using a transactional storage engine such as `InnoDB`, MySQL automatically rolls back the statement. If you are using a non-transactional storage engine, MySQL stops processing the statement at the row for which the error occurred and leaves any remaining rows unprocessed.

If you want to ignore such key violations, MySQL supports an `IGNORE` keyword for `INSERT` and `UPDATE`. In this case, MySQL ignores any key violations and continues processing with the next row. See [Section 13.2.4](#), “`INSERT Syntax`”, and [Section 13.2.10](#), “`UPDATE Syntax`”.

You can get information about the number of rows actually inserted or updated with the `mysql_info()` C API function. In MySQL 4.1 and up, you also can use the `SHOW WARNINGS` statement. See [Section 17.2.3.33](#), “`mysql_info()`”, and [Section 13.5.4.21](#), “`SHOW WARNINGS Syntax`”.

Currently, only `InnoDB` tables support foreign keys. See [Section 14.2.7.4, “FOREIGN KEY Constraints”](#). (Foreign key support in `MyISAM` tables is scheduled for implementation in MySQL 5.2. See [Section 1.6, “MySQL Development Roadmap”](#).)

### 1.9.6.2. Constraints on Invalid Data

Through version 4.1, MySQL is forgiving of illegal or improper data values and coerces them to legal values for data entry. When you insert an “incorrect” value into a column, such as a `NULL` into a `NOT NULL` column or a too-large numeric value into a numeric column, MySQL sets the column to the “best possible value” instead of producing an error. The following rules describe in more detail how this works:

- If you try to store an out of range value into a numeric column, MySQL Server instead stores zero, the smallest possible value, or the largest possible value, whichever is closest to the invalid value.
- For strings, MySQL stores either the empty string or as much of the string as can be stored in the column.
- If you try to store a string that doesn't start with a number into a numeric column, MySQL Server stores 0.
- Invalid values for `ENUM` and `SET` columns are handled as described in [Section 1.9.6.3, “ENUM and SET Constraints”](#).
- MySQL allows you to store certain incorrect date values into `DATE` and `DATETIME` columns (such as `'2000-02-31'` or `'2000-02-00'`). The idea is that it's not the job of the SQL server to validate dates. If MySQL can store a date value and retrieve exactly the same value, MySQL stores it as given. If the date is totally wrong (outside the server's ability to store it), the special “zero” date value `'0000-00-00'` is stored in the column instead.
- If you try to store `NULL` into a column that doesn't take `NULL` values, an error occurs for single-row `INSERT` statements. For multiple-row `INSERT` statements or for `INSERT INTO ... SELECT` statements, MySQL Server stores the implicit default value for the column data type. In general, this is 0 for numeric types, the empty string (`' '`) for string types, and the “zero” value for date and time types. Implicit default values are discussed in [Section 11.1.4, “Data Type Default Values”](#).
- If an `INSERT` statement specifies no value for a column, MySQL inserts its default value if the column definition includes an explicit `DEFAULT` clause. If the definition has no such `DEFAULT` clause, MySQL inserts the implicit default value for the column data type.

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

### 1.9.6.3. `ENUM` and `SET` Constraints

`ENUM` and `SET` columns provide an efficient way to define columns that can contain only a given set of values. See [Section 11.4.4, “The `ENUM` Type”](#), and [Section 11.4.5, “The `SET` Type”](#). However, in MySQL 4.1 and earlier, `ENUM` and `SET` columns do not provide true constraints on entry of invalid data:

- `ENUM` columns always have a default value. If you specify no default value, then it is `NULL` for columns that can have `NULL`, otherwise it is the first enumeration value in the column definition.

- If you insert an incorrect value into an `ENUM` column or if you force a value into an `ENUM` column with `IGNORE`, it is set to the reserved enumeration value of `0`, which is displayed as an empty string in string context.
- If you insert an incorrect value into a `SET` column, the incorrect value is ignored. For example, if the column can contain the values `'a'`, `'b'`, and `'c'`, an attempt to assign `'a,x,b,y'` results in a value of `'a,b'`.



---

# Chapter 2. Installing and Upgrading MySQL

This chapter describes how to obtain and install MySQL. A summary of the procedure follows and later sections provide the details. If you plan to upgrade an existing version of MySQL to a newer version rather than install MySQL for the first time, see [Section 2.11, “Upgrading MySQL”](#), for information about upgrade procedures and about issues that you should consider before upgrading.

1. **Determine whether your platform is supported.** Please note that not all supported systems are equally suitable for running MySQL. On some platforms it is much more robust and efficient than others. See [Section 2.1.1, “Operating Systems Supported by MySQL”](#), for details.
2. **Choose which distribution to install.** Several versions of MySQL are available, and most are available in several distribution formats. You can choose from pre-packaged distributions containing binary (precompiled) programs or source code. When in doubt, use a binary distribution. We also provide public access to our current source tree for those who want to see our most recent developments and help us test new code. To determine which version and type of distribution you should use, see [Section 2.1.2, “Choosing Which MySQL Distribution to Install”](#).
3. **Download the distribution that you want to install.** For instructions, see [Section 2.1.3, “How to Get MySQL”](#). To verify the integrity of the distribution, use the instructions in [Section 2.1.4, “Verifying Package Integrity Using MD5 Checksums or GnuPG”](#).
4. **Install the distribution.** To install MySQL from a binary distribution, use the instructions in [Section 2.2, “Standard MySQL Installation Using a Binary Distribution”](#). To install MySQL from a source distribution or from the current development source tree, use the instructions in [Section 2.9, “MySQL Installation Using a Source Distribution”](#).

If you encounter installation difficulties, see [Section 2.13, “Operating System-Specific Notes”](#), for information on solving problems for particular platforms.

5. **Perform any necessary post-installation setup.** After installing MySQL, read [Section 2.10, “Post-Installation Setup and Testing”](#). This section contains important information about making sure the MySQL server is working properly. It also describes how to secure the initial MySQL user accounts, *which have no passwords* until you assign passwords. The section applies whether you install MySQL using a binary or source distribution.
6. If you want to run the MySQL benchmark scripts, Perl support for MySQL must be available. See [Section 2.14, “Perl Installation Notes”](#).

## 2.1. General Installation Issues

Before installing MySQL, you should do the following:

1. Determine whether MySQL runs on your platform.
2. Choose a distribution to install.
3. Download the distribution and verify its integrity.

This section contains the information necessary to carry out these steps. After doing so, you can use the instructions in later sections of the chapter to install the distribution that you choose.

## 2.1.1. Operating Systems Supported by MySQL

This section lists the operating systems on which you can expect to be able to run MySQL.

We use GNU Autoconf, so it is possible to port MySQL to all modern systems that have a C++ compiler and a working implementation of POSIX threads. (Thread support is needed for the server. To compile only the client code, the only requirement is a C++ compiler.) We use and develop the software ourselves primarily on Linux (SuSE and Red Hat), FreeBSD, and Sun Solaris (versions 8 and 9).

MySQL has been reported to compile successfully on the following combinations of operating system and thread package. Note that for many operating systems, native thread support works only in the latest versions.

- AIX 4.x, 5.x with native threads. See [Section 2.13.5.3, “IBM-AIX notes”](#).
- Amiga.
- BSDI 2.x with the MIT-pthreads package. See [Section 2.13.4.4, “BSD/OS Version 2.x Notes”](#).
- BSDI 3.0, 3.1 and 4.x with native threads. See [Section 2.13.4.4, “BSD/OS Version 2.x Notes”](#).
- Digital Unix 4.x with native threads. See [Section 2.13.5.5, “Alpha-DEC-UNIX Notes \(Tru64\)”](#).
- FreeBSD 2.x with the MIT-pthreads package. See [Section 2.13.4.1, “FreeBSD Notes”](#).
- FreeBSD 3.x and 4.x with native threads. See [Section 2.13.4.1, “FreeBSD Notes”](#).
- FreeBSD 4.x with LinuxThreads. See [Section 2.13.4.1, “FreeBSD Notes”](#).
- HP-UX 10.20 with the DCE threads or the MIT-pthreads package. See [Section 2.13.5.1, “HP-UX Version 10.20 Notes”](#).
- HP-UX 11.x with the native threads. See [Section 2.13.5.2, “HP-UX Version 11.x Notes”](#).
- Linux 2.0+ with LinuxThreads 0.7.1+ or [glibc 2.0.7+](#) for various CPU architectures. See [Section 2.13.1, “Linux Notes”](#).
- Mac OS X. See [Section 2.13.2, “Mac OS X Notes”](#).
- NetBSD 1.3/1.4 Intel and NetBSD 1.3 Alpha (requires GNU make). See [Section 2.13.4.2, “NetBSD Notes”](#).
- Novell NetWare 6.0 and 6.5. See [Section 2.7, “Installing MySQL on NetWare”](#).
- OpenBSD 2.5 and with native threads. OpenBSD earlier than 2.5 with the MIT-pthreads package. See [Section 2.13.4.3, “OpenBSD 2.5 Notes”](#).
- OS/2 Warp 3, FixPack 29 and OS/2 Warp 4, FixPack 4. See [Section 2.13.6, “OS/2 Notes”](#).
- SCO OpenServer 5.0.X with a recent port of the FSU Pthreads package. See [Section 2.13.5.8, “SCO UNIX and OpenServer 5.0.x Notes”](#).
- SCO Openserver 6.0.x. See [Section 2.13.5.9, “SCO OpenServer 6.0.x Notes”](#).
- SCO UnixWare 7.1.x. See [Section 2.13.5.10, “SCO UnixWare 7.1.x and OpenUNIX 8.0.0 Notes”](#).
- SGI Irix 6.x with native threads. See [Section 2.13.5.7, “SGI Irix Notes”](#).
- Solaris 2.5 and above with native threads on SPARC and x86. See [Section 2.13.3, “Solaris Notes”](#).

- SunOS 4.x with the MIT-pthreads package. See [Section 2.13.3, “Solaris Notes”](#).
- Tru64 Unix. See [Section 2.13.5.5, “Alpha-DEC-UNIX Notes \(Tru64\)”](#).
- Windows 9x, Me, NT, 2000, XP, and Windows Server 2003. See [Section 2.3, “Installing MySQL on Windows”](#).

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

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

Based on the preceding criteria, the best platforms for running MySQL at this point are x86 with SuSE Linux using a 2.4 or 2.6 kernel, and ReiserFS (or any similar Linux distribution) and SPARC with Solaris (2.7-9). FreeBSD comes third, but we really hope it joins the top club once the thread library is improved. We also hope that at some point we are able to include into the top category all other platforms on which MySQL currently compiles and runs, but not quite with the same level of stability and performance. This requires some effort on our part in cooperation with the developers of the operating systems and library components that MySQL depends on. If you are interested in improving one of those components, are in a position to influence its development, and need more detailed instructions on what MySQL needs to run better, send an email message to the MySQL [internals](#) mailing list. See [Section 1.7.1, “MySQL Mailing Lists”](#).

Please note that the purpose of the preceding comparison is not to say that one operating system is better or worse than another in general. We are talking only about choosing an OS for the specific purpose of running MySQL. With this in mind, the result of this comparison might be different if other factors were considered. In some cases, the reason one OS is better for MySQL than another might simply be that we have been able to put more effort into testing and optimizing for a particular platform. We are just stating our observations to help you decide which platform to use for running MySQL.

## 2.1.2. Choosing Which MySQL Distribution to Install

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

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

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

- MySQL 5.1 is the current development release series.
- MySQL 5.0 is the current stable (production-quality) release series. New releases are issued for bugfixes only; no new features are being added that could effect stability.
- MySQL 4.1 is the previous stable (production-quality) release series. New releases are issued for critical bugfixes and security fixes. No significant new features are to be added to this series.
- MySQL 4.0 and 3.23 are the old stable (production-quality) release series. These versions are now retired, so new releases are issued only to fix extremely critical bugs (primarily security issues).

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

Normally, if you are beginning to use MySQL for the first time or trying to port it to some system for which there is no binary distribution, we recommend going with the production release series. Currently, this is MySQL 5.0. All MySQL releases, even those from development series, are checked with the MySQL benchmarks and an extensive test suite before being issued.

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

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

If you want to use the very latest sources containing all current patches and bugfixes, you can use one of our BitKeeper repositories. These are not “releases” as such, but are available as previews of the code on which future releases are to be based.

The MySQL naming scheme uses release names that consist of three numbers and a suffix; for example, `mysql-4.1.2-alpha`. The numbers within the release name are interpreted like this:

- The first number (4) is the major version and also describes the file format. All version 4 releases have the same file format.
- The second number (1) is the release level. Taken together, the major version and release level constitute the release series number.
- The third number (2) is the version number within the release series. This is incremented for each new release. Usually you want the latest version for the series you have chosen.

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

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

- [alpha](#) indicates that the release contains some large section of new code that hasn't been 100% tested. Known bugs should be documented in the News section. See [Appendix D, \*MySQL Change History\*](#). There are also new commands and extensions in most alpha releases. Active development that may involve major code changes can occur in an alpha release, but everything is tested before issuing a release.
- [beta](#) means that we are feature complete and that all new code has been tested. No major new features that could cause corruption in old code are added. There should be no known critical bugs. A version changes from alpha to beta when there have not been any reported fatal bugs within an alpha version for at least a month and we have no plans to add any features that could make any old command unreliable.

All APIs, externally visible structures, and columns for SQL statements will not change during future beta, release candidate, or production releases.

- [rc](#) is a release candidate; that is, a beta that has been around a while and seems to work fine. Only minor fixes are added. (A release candidate is what formerly was known as a gamma release.)
- If there is no suffix, it means that the version has been run for a while at many different sites with no reports of critical repeatable bugs other than platform-specific bugs. Only critical bugfixes are applied to the release. This is what we call a production (stable) or “General Availability” (GA) release.

MySQL uses a naming scheme that is slightly different from most other products. In general, it is usually safe to use any version that has been out for a couple of weeks without being replaced by a new version within the same release series.

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

All releases have been tested at least with these tools:

- An internal test suite

The `mysql-test` directory contains an extensive set of test cases. We run these tests for virtually every server binary. See [Section 19.1.2, “MySQL Test Suite”](#), for more information about this test suite.

- The MySQL benchmark suite

This suite runs a range of common queries. It is also a test to determine whether the latest batch of optimizations actually made the code faster. See [Section 7.1.4, “The MySQL Benchmark Suite”](#).

- The `crash-me` test

This test tries to determine what features the database supports and what its capabilities and limitations are. See [Section 7.1.4, “The MySQL Benchmark Suite”](#).

We also test the newest MySQL version in our internal production environment, on at least one machine. We have more than 100GB of data to work with.

### 2.1.2.2. Choosing a Distribution Format

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

Reasons to choose a binary distribution include the following:

- Binary distributions generally are easier to install than source distributions.
- To satisfy different user requirements, we provide two different binary versions. One is compiled with the core feature set. The other (MySQL-Max) is compiled with an extended feature set. Both versions are compiled from the same source distribution. All native MySQL clients can connect to servers from either MySQL version.

The extended MySQL binary distribution is identified by the `-max` suffix and is configured with the same options as `mysqld-max`. See [Section 5.3, “The `mysqld-max` Extended MySQL Server”](#).

For RPM distributions, if you want to use the `MySQL-Max` RPM, you must first install the standard `MySQL-server` RPM.

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

- You want to install MySQL at some explicit location. The standard binary distributions are ready to run at any installation location, but you might require even more flexibility to place MySQL components where you want.
- You want to configure `mysqld` to ensure that features are available that might not be included in the standard binary distributions. Here is a list of the most common extra options that you may want to use to ensure feature availability:
  - `--with-berkeley-db` (not available on all platforms)
  - `--with-raid`
  - `--with-libwrap`
  - `--with-named-z-libs` (this is done for some of the binaries)
  - `--with-debug[=full]`
- You want to configure `mysqld` without some features that are included in the standard binary distributions. For example, distributions normally are compiled with support for all character sets. If you want a smaller MySQL server, you can recompile it with support for only the character sets you need.
- You have a special compiler (such as `pgcc`) or want to use compiler options that are better optimized for your processor. Binary distributions are compiled with options that should work on a variety of processors from the same processor family.
- You want to use the latest sources from one of the BitKeeper repositories to have access to all cur-

rent bugfixes. For example, if you have found a bug and reported it to the MySQL development team, the bugfix is committed to the source repository and you can access it there. The bugfix does not appear in a release until a release actually is issued.

- You want to read (or modify) the C and C++ code that makes up MySQL. For this purpose, you should get a source distribution, because the source code is always the ultimate manual.
- Source distributions contain more tests and examples than binary distributions.

### 2.1.2.3. How and When Updates Are Released

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

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

No one is *required* to download a new release. The News section helps you determine whether the new release has something you really want. See [Appendix D, MySQL Change History](#).

We use the following policy when updating MySQL:

- Releases are issued within each series. For each release, the last number in the version is one more than the previous release within the same series.
- Production (stable) releases are meant to appear about 1-2 times a year. However, if small bugs are found, a release with only bugfixes is issued.
- Working releases/bugfixes to old releases are meant to appear about every 4-8 weeks.
- Binary distributions for some platforms are made by us for major releases. Other people may make binary distributions for other systems, but probably less frequently.
- We make fixes available as soon as we have identified and corrected small or non-critical but annoying bugs. The fixes are available immediately from our public BitKeeper repositories, and are included in the next release.
- If by any chance a fatal bug is found in a release, our policy is to fix it in a new release as soon as possible. (We would like other companies to do this, too!)

### 2.1.2.4. Release Philosophy—No Known Bugs in Releases

We put considerable time and effort into making our releases bug-free. Our policy is never to release a version of MySQL intended for production use that has any known fatal, repeatable bugs.

We have documented all open problems, bugs, and issues that are dependent on design decisions. See [Section A.8, “Known Issues in MySQL”](#).

Our aim is to fix everything that is fixable without making a stable MySQL version less stable. In certain cases, this means we can fix an issue in the development versions, but not in the stable (production) version. Naturally, we document such issues so that users are aware of them.

Here is a description of our build process:

- We monitor bugs from our customer support list, the bugs database at <http://bugs.mysql.com/>, and the MySQL external mailing lists.
- All reported bugs for live versions are entered into the bugs database.
- When we fix a bug, we always try to make a test case for it and include it into our test system to ensure that the bug can never recur without being detected. (About 90% of all fixed bugs have test cases.)
- We create test cases for each new feature that we add to MySQL.
- Before we start to build a new MySQL release, we ensure that all reported repeatable bugs for that MySQL version (3.23.x, 4.0.x, 4.1.x, 5.0.x, 5.1.x, and so on) are fixed. If something is impossible to fix due to some internal design decision in MySQL, we document this in the manual. See [Section A.8, “Known Issues in MySQL”](#).
- We do a build on all platforms for which we support binaries and run our test suite and benchmark suite on all of them.
- We do not publish a binary for a platform for which the test or benchmark suite fails. If the problem is due to a general error in the source, we fix it and do the build plus tests on all systems again from scratch.
- The build and test process takes a week. If we receive a report regarding a fatal bug during this process (for example, one that causes a core dump), we fix the problem and restart the build process.
- After publishing the binaries on <http://dev.mysql.com/>, we send out an announcement message to the [mysql](#) and [announce](#) mailing lists. See [Section 1.7.1, “MySQL Mailing Lists”](#). The announcement message contains a list of all changes to the release and any known problems with the release. The **Known Problems** section in the release notes has been needed for only a handful of releases.
- To quickly give our users access to the latest MySQL features, we try to produce a new MySQL release every 4-8 weeks. Source code snapshots are built daily and are available at <http://downloads.mysql.com/snapshots.php>.
- If, despite our best efforts, we receive any bug reports after a release is issued that a critical problem exists for the build on a specific platform, we fix it at once and build a new 'a' release for that platform. Thanks to our large user base, problems are found and resolved very quickly.
- Our track record for making stable releases is quite good. In the last 150 releases, we had to do a new build for fewer than 10 of them. In three of these cases, the bug was a faulty [glibc](#) library on one of our build machines that took us a long time to track down.

### 2.1.2.5. MySQL Binaries Compiled by MySQL AB

As a service of MySQL AB, we provide a set of binary distributions of MySQL that are compiled on systems at our site or on systems where supporters of MySQL kindly have given us access to their machines.

In addition to the binaries provided in platform-specific package formats, we offer binary distributions for a number of platforms in the form of compressed `tar` files (`.tar.gz` files). See [Section 2.2, “Standard MySQL Installation Using a Binary Distribution”](#).

RPM distributions prior to MySQL 3.22 are user-contributed. Beginning with MySQL 3.22, RPM distributions that we make available through our Web site are generated by MySQL AB.

For Windows distributions, see [Section 2.3, “Installing MySQL on Windows”](#).



These distributions are generated using the script `Build-tools/Do-compile`, which compiles the source code and creates the binary `tar.gz` archive using `scripts/make_binary_distribution`.

These binaries are configured and built with the following compilers and options. This information can also be obtained by looking at the variables `COMP_ENV_INFO` and `CONFIGURE_LINE` inside the script `bin/mysqlbug` of every binary `tar` file distribution.

Anyone who has more optimal options for any of the following `configure` commands can mail them to the MySQL `internals` mailing list. See [Section 1.7.1, “MySQL Mailing Lists”](#).

If you want to compile a debug version of MySQL, you should add `--with-debug` or `--with-debug=full` to the following `configure` commands and remove any `fomit-frame-pointer` options.

The following binaries are built on MySQL AB development systems:

- Linux 2.4.xx x86 with `gcc` 2.95.3:

```
CFLAGS="-O2 -mcpu=pentiumpro" CXX=gcc CXXFLAGS="-O2 -mcpu=pentiumpro
-felide-constructors" ./configure --prefix=/usr/local/mysql
--with-extra-charsets=complex --enable-thread-safe-client
--enable-local-infile --enable-asm-asm --disable-shared
--with-client-ldflags=-all-static --with-mysqld-ldflags=-all-static
```

- Linux 2.4.x x86 with `icc` (Intel C++ Compiler 8.1 or later releases):

```
CC=icc CXX=icpc CFLAGS="-O3 -unroll2 -ip -mp -no-gcc -restrict"
CXXFLAGS="-O3 -unroll2 -ip -mp -no-gcc -restrict" ./configure
--prefix=/usr/local/mysql --localstatedir=/usr/local/mysql/data
--libexecdir=/usr/local/mysql/bin --with-extra-charsets=complex
--enable-thread-safe-client --enable-local-infile --enable-asm-asm
--disable-shared --with-client-ldflags=-all-static
--with-mysqld-ldflags=-all-static --with-embedded-server --with-innodb
```

Note that versions 8.1 and newer of the Intel compiler have separate drivers for 'pure' C (`icc`) and C++ (`icpc`); if you use `icc` version 8.0 or older for building MySQL, you will need to set `CXX=icc`.

- Linux 2.4.xx Intel Itanium 2 with `ecc` (Intel C++ Itanium Compiler 7.0):

```
CC=ecc CFLAGS="-O2 -tpp2 -ip -nolib_inline" CXX=ecc CXXFLAGS="-O2
-tpp2 -ip -nolib_inline" ./configure --prefix=/usr/local/mysql
--with-extra-charsets=complex --enable-thread-safe-client
--enable-local-infile
```

- Linux 2.4.xx Intel Itanium with `ecc` (Intel C++ Itanium Compiler 7.0):

```
CC=ecc CFLAGS=-tpp1 CXX=ecc CXXFLAGS=-tpp1 ./configure
--prefix=/usr/local/mysql --with-extra-charsets=complex
--enable-thread-safe-client --enable-local-infile
```

- Linux 2.4.xx alpha with `ccc` (Compaq C V6.2-505 / Compaq C++ V6.3-006):

```
CC=ccc CFLAGS="-fast -arch generic" CXX=cxx CXXFLAGS="-fast -arch
generic -noexceptions -nortti" ./configure --prefix=/usr/local/mysql
--with-extra-charsets=complex --enable-thread-safe-client
--enable-local-infile --with-mysqld-ldflags=-non_shared
--with-client-ldflags=-non_shared --disable-shared
```

- Linux 2.x.xx ppc with `gcc` 2.95.4:

```
CC=gcc CFLAGS="-O3 -fno-omit-frame-pointer" CXX=gcc CXXFLAGS="-O3
-fno-omit-frame-pointer -felide-constructors -fno-exceptions
-fno-rtti" ./configure --prefix=/usr/local/mysql
--localstatedir=/usr/local/mysql/data --libexecdir=/usr/local/mysql/bin
--with-extra-charsets=complex --enable-thread-safe-client
--enable-local-infile --disable-shared --with-embedded-server
--with-innodb
```

- Linux 2.4.xx s390 with [gcc 2.95.3](#):

```
CFLAGS="-O2" CXX=gcc CXXFLAGS="-O2 -felide-constructors" ./configure
--prefix=/usr/local/mysql --with-extra-charsets=complex
--enable-thread-safe-client --enable-local-infile --disable-shared
--with-client-ldflags=-all-static --with-mysqld-ldflags=-all-static
```

- Linux 2.4.xx x86\_64 (AMD64) with [gcc 3.2.1](#):

```
CXX=gcc ./configure --prefix=/usr/local/mysql
--with-extra-charsets=complex --enable-thread-safe-client
--enable-local-infile --disable-shared
```

- Sun Solaris 8 x86 with [gcc 3.2.3](#):

```
CC=gcc CFLAGS="-O3 -fno-omit-frame-pointer" CXX=gcc CXXFLAGS="-O3
-fno-omit-frame-pointer -felide-constructors -fno-exceptions
-fno-rtti" ./configure --prefix=/usr/local/mysql
--localstatedir=/usr/local/mysql/data --libexecdir=/usr/local/mysql/bin
--with-extra-charsets=complex --enable-thread-safe-client
--enable-local-infile --disable-shared --with-innodb
```

- Sun Solaris 8 SPARC with [gcc 3.2](#):

```
CC=gcc CFLAGS="-O3 -fno-omit-frame-pointer" CXX=gcc CXXFLAGS="-O3
-fno-omit-frame-pointer -felide-constructors -fno-exceptions
-fno-rtti" ./configure --prefix=/usr/local/mysql
--with-extra-charsets=complex --enable-thread-safe-client
--enable-local-infile --enable-assembler --with-named-z-libs=no
--with-named-curses-libs=-lcurses --disable-shared
```

- Sun Solaris 8 SPARC 64-bit with [gcc 3.2](#):

```
CC=gcc CFLAGS="-O3 -m64 -fno-omit-frame-pointer" CXX=gcc CXXFLAGS="-O3
-m64 -fno-omit-frame-pointer -felide-constructors -fno-exceptions
-fno-rtti" ./configure --prefix=/usr/local/mysql
--with-extra-charsets=complex --enable-thread-safe-client
--enable-local-infile --with-named-z-libs=no
--with-named-curses-libs=-lcurses --disable-shared
```

- Sun Solaris 9 SPARC with [gcc 2.95.3](#):

```
CC=gcc CFLAGS="-O3 -fno-omit-frame-pointer" CXX=gcc CXXFLAGS="-O3
-fno-omit-frame-pointer -felide-constructors -fno-exceptions
-fno-rtti" ./configure --prefix=/usr/local/mysql
--with-extra-charsets=complex --enable-thread-safe-client
--enable-local-infile --enable-assembler --with-named-curses-libs=-lcurses
--disable-shared
```

- Sun Solaris 9 SPARC with [cc-5.0](#) (Sun Forte 5.0):

```
CC=cc-5.0 CXX=CC ASFLAGS="-xarch=v9" CFLAGS="-Xa -xstrconst -mt
-D_FORTEC_ -xarch=v9" CXXFLAGS="-noex -mt -D_FORTEC_ -xarch=v9"
./configure --prefix=/usr/local/mysql --with-extra-charsets=complex
--enable-thread-safe-client --enable-local-infile --enable-assembler
--with-named-z-libs=no --enable-thread-safe-client --disable-shared
```

- IBM AIX 4.3.2 ppc with [gcc 3.2.3](#):

```
CFLAGS="-O2 -mcpu=powerpc -Wa,-many " CXX=gcc CXXFLAGS="-O2
-mcpu=powerpc -Wa,-many -felide-constructors -fno-exceptions
-fno-rtti" ./configure --prefix=/usr/local/mysql
--with-extra-charsets=complex --enable-thread-safe-client
--enable-local-infile --with-named-z-libs=no --disable-shared
```

- IBM AIX 4.3.3 ppc with [x1C\\_r](#) (IBM Visual Age C/C++ 6.0):

```
CC=x1c_r CFLAGS="-ma -O2 -qstrict -qoptimize=2 -qmaxmem=8192"
CXX=x1c_r CXXFLAGS="-ma -O2 -qstrict -qoptimize=2 -qmaxmem=8192"
./configure --prefix=/usr/local/mysql --localstatedir=/usr/local/mysql/data
--libexecdir=/usr/local/mysql/bin --with-extra-charsets=complex
--enable-thread-safe-client --enable-local-infile --with-named-z-libs=no
--disable-shared --with-innodb
```

- IBM AIX 5.1.0 ppc with [gcc 3.3](#):

```
CFLAGS="-O2 -mcpu=powerpc -Wa,-many" CXX=gcc CXXFLAGS="-O2 -mcpu=powerpc
-Wa,-many -felide-constructors -fno-exceptions -fno-rtti" ./configure
--prefix=/usr/local/mysql --with-extra-charsets=complex
--enable-thread-safe-client --enable-local-infile --with-named-z-libs=no
--disable-shared
```

- IBM AIX 5.2.0 ppc with [x1C\\_r](#) (IBM Visual Age C/C++ 6.0):

```
CC=x1c_r CFLAGS="-ma -O2 -qstrict -qoptimize=2 -qmaxmem=8192"
CXX=x1c_r CXXFLAGS="-ma -O2 -qstrict -qoptimize=2 -qmaxmem=8192"
./configure --prefix=/usr/local/mysql --localstatedir=/usr/local/mysql/data
--libexecdir=/usr/local/mysql/bin --with-extra-charsets=complex
--enable-thread-safe-client --enable-local-infile --with-named-z-libs=no
--disable-shared --with-embedded-server --with-innodb
```

- HP-UX 10.20 pa-risc1.1 with [gcc 3.1](#):

```
CFLAGS="-DHPUX -I/opt/dce/include -O3 -fPIC" CXX=gcc CXXFLAGS="-DHPUX
-I/opt/dce/include -felide-constructors -fno-exceptions -fno-rtti
-O3 -fPIC" ./configure --prefix=/usr/local/mysql
--with-extra-charsets=complex --enable-thread-safe-client
--enable-local-infile --with-pthread --with-named-thread-libs=ldce
--with-lib-ccflags=-fPIC --disable-shared
```

- HP-UX 11.00 pa-risc with [aCC](#) (HP ANSI C++ B3910B A.03.50):

```
CC=cc CXX=aCC CFLAGS="+DAportable" CXXFLAGS="+DAportable" ./configure
--prefix=/usr/local/mysql --localstatedir=/usr/local/mysql/data
--libexecdir=/usr/local/mysql/bin --with-extra-charsets=complex
--enable-thread-safe-client --enable-local-infile --disable-shared
--with-embedded-server --with-innodb
```

- HP-UX 11.11 pa-risc2.0 64bit with [aCC](#) (HP ANSI C++ B3910B A.03.33):

```
CC=cc CXX=aCC CFLAGS="+DD64" CXXFLAGS="+DD64" ./configure
--prefix=/usr/local/mysql --with-extra-charsets=complex
--enable-thread-safe-client --enable-local-infile --disable-shared
```

- HP-UX 11.11 pa-risc2.0 32bit with [aCC](#) (HP ANSI C++ B3910B A.03.33):

```
CC=cc CXX=aCC CFLAGS="+DAportable" CXXFLAGS="+DAportable" ./configure
--prefix=/usr/local/mysql --localstatedir=/usr/local/mysql/data
--libexecdir=/usr/local/mysql/bin --with-extra-charsets=complex
--enable-thread-safe-client --enable-local-infile --disable-shared
--with-innodb
```

- HP-UX 11.22 ia64 64bit with [aCC](#) (HP aC++/ANSI C B3910B A.05.50):

```
CC=cc CXX=aCC CFLAGS="+DD64 +DSitanium2" CXXFLAGS="+DD64 +DSitanium2"
./configure --prefix=/usr/local/mysql --localstatedir=/usr/local/mysql/data
--libexecdir=/usr/local/mysql/bin --with-extra-charsets=complex
--enable-thread-safe-client --enable-local-infile --disable-shared
--with-embedded-server --with-innodb
```

- Apple Mac OS X 10.2 powerpc with [gcc 3.1](#):

```
CC=gcc CFLAGS="-O3 -fno-omit-frame-pointer" CXX=gcc CXXFLAGS="-O3
-fno-omit-frame-pointer -felide-constructors -fno-exceptions
-fno-rtti" ./configure --prefix=/usr/local/mysql
--with-extra-charsets=complex --enable-thread-safe-client
--enable-local-infile --disable-shared
```

- FreeBSD 4.7 i386 with [gcc 2.95.4](#):

```
CFLAGS=-DHAVE_BROKEN_REALPATH ./configure --prefix=/usr/local/mysql
--with-extra-charsets=complex --enable-thread-safe-client
--enable-local-infile --enable-asmsembler --with-named-z-libs=not-used
--disable-shared
```

- FreeBSD 4.7 i386 using LinuxThreads with [gcc 2.95.4](#):

```
CFLAGS="-DHAVE_BROKEN_REALPATH -D__USE_UNIX98 -D_REENTRANT
-D_THREAD_SAFE -I/usr/local/include/pthread/linuxthreads"
CXXFLAGS="-DHAVE_BROKEN_REALPATH -D__USE_UNIX98 -D_REENTRANT
-D_THREAD_SAFE -I/usr/local/include/pthread/linuxthreads" ./configure
--prefix=/usr/local/mysql --localstatedir=/usr/local/mysql/data
--libexecdir=/usr/local/mysql/bin --enable-thread-safe-client
--enable-local-infile --enable-asmsembler
--with-named-thread-libs="-DHAVE_GLIBC2_STYLE_GETHOSTBYNAME_R
-D_THREAD_SAFE -I /usr/local/include/pthread/linuxthreads
-L/usr/local/lib -llthread -llgcc_r" --disable-shared
--with-embedded-server --with-innodb
```

- QNX Neutrino 6.2.1 i386 with [gcc 2.95.3qnx-nto 20010315](#):

```
CC=gcc CFLAGS="-O3 -fno-omit-frame-pointer" CXX=gcc CXXFLAGS="-O3
-fno-omit-frame-pointer -felide-constructors -fno-exceptions
-fno-rtti" ./configure --prefix=/usr/local/mysql
--with-extra-charsets=complex --enable-thread-safe-client
--enable-local-infile --disable-shared
```

The following binaries are built on third-party systems kindly provided to MySQL AB by other users. These are provided only as a courtesy; MySQL AB does not have full control over these systems, so we can provide only limited support for the binaries built on them.

- SCO Unix 3.2v5.0.7 i386 with [gcc 2.95.3](#):

```
CFLAGS="-O3 -mpentium" LDFLAGS=-static CXX=gcc CXXFLAGS="-O3 -mpentium
-felide-constructors" ./configure --prefix=/usr/local/mysql
--with-extra-charsets=complex --enable-thread-safe-client
--enable-local-infile --with-named-z-libs=no --enable-thread-safe-client
--disable-shared
```

- SCO UnixWare 7.1.4 i386 with [CC 3.2](#):

```
CC=cc CFLAGS="-O" CXX=CC ./configure --prefix=/usr/local/mysql
--with-extra-charsets=complex --enable-thread-safe-client
--enable-local-infile --with-named-z-libs=no --enable-thread-safe-client
--disable-shared --with-readline
```

- SCO OpenServer 6.0.0 i386 with [CC 3.2](#):

```
CC=cc CFLAGS="-O" CXX=CC ./configure --prefix=/usr/local/mysql
--with-extra-charsets=complex --enable-thread-safe-client
--enable-local-infile --with-named-z-libs=no --enable-thread-safe-client
--disable-shared --with-readline
```

- Compaq Tru64 OSF/1 V5.1 732 alpha with [cc/cxx](#) (Compaq C V6.3-029i / DIGITAL C++ V6.1-027):

```
CC="cc -pthread" CFLAGS="-O4 -ansi_alias -ansi_args -fast -inline
speed -speculate all" CXX="cxx -pthread" CXXFLAGS="-O4 -ansi_alias
-fast -inline speed -speculate all -noexceptions -nortti" ./configure
--prefix=/usr/local/mysql --with-extra-charsets=complex
--enable-thread-safe-client --enable-local-infile
--with-named-thread-libs="-lpthread -lmach -lexc -lc" --disable-shared
--with-mysqld-ldflags=-all-static
```

- SGI Irix 6.5 IP32 with [gcc 3.0.1](#):

```
CC=gcc CFLAGS="-O3 -fno-omit-frame-pointer" CXXFLAGS="-O3
-fno-omit-frame-pointer -felide-constructors -fno-exceptions
-fno-rtti" ./configure --prefix=/usr/local/mysql
--with-extra-charsets=complex --enable-thread-safe-client
--enable-local-infile --disable-shared
```

- FreeBSD/sparc64 5.0 with [gcc 3.2.1](#):

```
CFLAGS=-DHAVE_BROKEN_REALPATH ./configure --prefix=/usr/local/mysql
--localstatedir=/usr/local/mysql/data --libexecdir=/usr/local/mysql/bin
--with-extra-charsets=complex --enable-thread-safe-client
--enable-local-infile --disable-shared --with-innodb
```

The following compile options have been used for binary packages that MySQL AB provided in the past. These binaries no longer are being updated, but the compile options are listed here for reference purposes.

- Linux 2.2.xx SPARC with [egcs 1.1.2](#):

```
CC=gcc CFLAGS="-O3 -fno-omit-frame-pointer" CXX=gcc CXXFLAGS="-O3
-fno-omit-frame-pointer -felide-constructors -fno-exceptions
-fno-rtti" ./configure --prefix=/usr/local/mysql
--with-extra-charsets=complex --enable-thread-safe-client
--enable-local-infile --enable-assembler --disable-shared
```

- Linux 2.2.x with x86 with [gcc 2.95.2](#):

```
CFLAGS="-O3 -mpentiumpro" CXX=gcc CXXFLAGS="-O3 -mpentiumpro
-felide-constructors -fno-exceptions -fno-rtti" ./configure
--prefix=/usr/local/mysql --enable-assembler
--with-mysqld-ldflags=-all-static --disable-shared
--with-extra-charsets=complex
```

- SunOS 4.1.4 2 sun4c with [gcc 2.7.2.1](#):

```
CC=gcc CXX=gcc CXXFLAGS="-O3 -felide-constructors" ./configure
--prefix=/usr/local/mysql --disable-shared --with-extra-charsets=complex
--enable-assembler
```

- SunOS 5.5.1 (and above) sun4u with [egcs 1.0.3a](#) or [2.90.27](#) or [gcc 2.95.2](#) and newer:

```
CC=gcc CFLAGS="-O3" CXX=gcc CXXFLAGS="-O3 -felide-constructors
```

```
-fno-exceptions -fno-rtti" ./configure --prefix=/usr/local/mysql
--with-low-memory --with-extra-charsets=complex --enable-asm
```

- SunOS 5.6 i86pc with [gcc 2.8.1](#):

```
CC=gcc CXX=gcc CXXFLAGS=-O3 ./configure --prefix=/usr/local/mysql
--with-low-memory --with-extra-charsets=complex
```

- BSDI BSD/OS 3.1 i386 with [gcc 2.7.2.1](#):

```
CC=gcc CXX=gcc CXXFLAGS=-O ./configure --prefix=/usr/local/mysql
--with-extra-charsets=complex
```

- BSDI BSD/OS 2.1 i386 with [gcc 2.7.2](#):

```
CC=gcc CXX=gcc CXXFLAGS=-O3 ./configure --prefix=/usr/local/mysql
--with-extra-charsets=complex
```

- AIX 4.2 with [gcc 2.7.2.2](#):

```
CC=gcc CXX=gcc CXXFLAGS=-O3 ./configure --prefix=/usr/local/mysql
--with-extra-charsets=complex
```

### 2.1.3. How to Get MySQL

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

Our main mirror is located at <http://mirrors.sunsite.dk/mysql/>.

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

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

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

The following sections describe how to use these methods.

If you notice that the MD5 checksum or GPG signatures do not match, first try to download the respective package one more time, perhaps from another mirror site. If you repeatedly cannot successfully verify the integrity of the package, please notify us about such incidents, including the full package name and the download site you have been using, at [<webmaster@mysql.com>](mailto:webmaster@mysql.com) or [<build@mysql.com>](mailto:build@mysql.com). Do not report downloading problems using the bug-reporting system.

### 2.1.4.1. Verifying the MD5 Checksum

After you have downloaded a MySQL package, you should make sure that its MD5 checksum matches the one provided on the MySQL download pages. Each package has an individual checksum that you can verify with the following command, where *package\_name* is the name of the package you downloaded:

```
shell> md5sum package_name
```

Example:

```
shell> md5sum mysql-standard-4.0.17-pc-linux-i686.tar.gz
60f5fe969d61c8f82e4f7f62657e1f06  mysql-standard-4.0.17-pc-linux-i686.tar.gz
```

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

**Note:** Make sure to verify the checksum of the *archive file* (for example, the *.zip* or *.tar.gz* file) and not of the files that are contained inside of the archive.

Note that not all operating systems support the `md5sum` command. On some, it is simply called `md5`, and others do not ship it at all. On Linux, it is part of the **GNU Text Utilities** package, which is available for a wide range of platforms. You can download the source code from <http://www.gnu.org/software/textutils/> as well. If you have OpenSSL installed, you can use the command `openssl md5 package_name` instead. A Windows implementation of the `md5` command line utility is available from <http://www.fourmilab.ch/md5/>. `winMd5Sum` is a graphical MD5 checking tool that can be obtained from <http://www.nullriver.com/index/products/winmd5sum>.

### 2.1.4.2. Signature Checking Using GnuPG

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

Beginning with MySQL 4.0.10 (February 2003), MySQL AB started signing downloadable packages with **GnuPG** (**GNU Privacy Guard**). **GnuPG** is an Open Source alternative to the very well-known **Pretty Good Privacy** (**PGP**) by Phil Zimmermann. See <http://www.gnupg.org/> for more information about **GnuPG** and how to obtain and install it on your system. Most Linux distributions ship with **GnuPG** installed by default. For more information about **OpenPGP**, see <http://www.openpgp.org/>.

To verify the signature for a specific package, you first need to obtain a copy of MySQL AB's public GPG build key, which you can download from <http://www.keyserver.net/>. The key that you want to obtain is named `build@mysql.com`. Alternatively, you can cut and paste the key directly from the following text:

```
Key ID:
pub 1024D/5072E1F5 2003-02-03
    MySQL Package signing key (www.mysql.com) <build@mysql.com>
Fingerprint: A4A9 4068 76FC BD3C 4567 70C8 8C71 8D3B 5072 E1F5

Public Key (ASCII-armored):

-----BEGIN PGP PUBLIC KEY BLOCK-----
Version: GnuPG v1.0.6 (GNU/Linux)
Comment: For info see http://www.gnupg.org

mQGIBD4+owwRBAC14GIfufCyEDSIEpVew3SAFUdJBtoQHH/nJKZyQT7h9bP1UWC3
RODjQReyCITRrdwyrKUGku2FmeVGwn2u2WmDMNABLnpPrWpkBdCk96+OmSLN9brZ
fw2vOUgCmYv2hW0hyDHuvY1QA/BThQoAdGj8AW6/0Lo7V1W9/8VuHP0gQwCgvzV3
BqOxRznNCRCRxAuAuVztHRcEAJooQK1+iSiunZMYD1WufeXfshc57S/+yeJkegNW
hxwR9pRWArNYJdDRT+rf2RUe3vpquKNQU/hnEIHJRQqYHo8gTxvxXNqc7fJYLV
K2HtkrPbP72vwsEKMYhhr0eKCbtLGf1s9krjJ6sBgACyP/Vb7hiPwxh6rDZ7ITnE
kYpXBACmWpP8NJTkamEnPCia2ZoOHODANwpUkP43I7jsDmgtobZX9qnrAXw+uNDI
QJEXM6FSbi0LLtZciNLYsafwAPEOMDKpMqAK6IyisNtPvaLd8lH0bPAnWqcyefep
```

```
rv0sxxqUEMcM3o7wvgn83P0kDasDbs3p jwPhxvzh6 //62zQJ7Q7TXlTUUwgUGFj
a2FnZSBzaWduaW5nIGtleSAod3d3Lm15c3FsLmNvbSkpPGJlAWxkQG15c3FsLmNv
bT6IXQQTfEQIAHQUCPj6 jDAUJCWYBgAULBwoDBAMVAwIDFgIBAheAAoJEIxxjTtQ
cuHlcY4AnilUwTXn8MatQOIG0a/bPxrVk/gCAJ4oInSNZRYTnblChwFaazt7PF3q
zIhMBBMRAGAMBQI+PgPRBYMJZgc7AAoJEElQ4SqycpHyJOEAnlmxHijft00bKXvu
cSo/pECUmpPiAJ41M9MRVj5VcdH/KN/KjRtW6tHFPYhMBBMRAGAMBQI+QoIDBYMJ
YiKJAAoJELblzU3GuiQ/lpEAoIhpp6BozKI8p6eaabzF5MlJH58pAKCu/ROofK8J
Eg2aLos+5zEYrB/LsrkCDQQ+PqMdEAgA7+GJfxbMdY4ws1PnjH9rF4N2qfWsen/l
xaZoJYc3a6M02WCnH16ahT2/tBK2w1QI4YFteR47gCvtgb6QlJHffOo2HfLmRDRR
rjd1DTCHqeyX7CHhcghj/dNRlW2Z015QFEcmV9U0Vhp3aFfWC4Ujfs3LU+hkAWzE
7zaD5cH9J7yv/6xuZVw41lx0h4UqsTcWmu0iMlBzELqXlDY7LwoPEb/O9Rkbf4fm
LellEzIaCa4PqARXQZc4dhSinMt6K3X4BrRsKTfOzBu74F47D8Ilbf5vSYHbuE5p
/loIDznkg/p8kW+3FxuWrycciqFTcNz215yyX39LXFnlLzKUb/F5GwADBQf+Lwqq
a8CGrRfsOAJxim63CHfty5mUc5rUSnTslGYEIOCRlBeQauyPZbPdsDD9Mz1ZaSaf
anFvwFG6Llx9xkU7tzq+vKLoWkm4u5xf3vn55VjnSd1aQ9eQnUcXiL4cnBG0TbOW
I39Ecyzgs1zBdC+MPjCQTcA7p6JUVsP6oAB3FQWg54tuUo0Ec8bsM8b3Ev42Lmu
QT5NdKHGwHsXTPt10klk4bQk40a jHsiy1BMahpT27jWjJlMiJc+IwJ0mghkKht92
6s/ymfdf5HkdQ1cyvsz5tryVI3F78XeSYfQvuuwqp2H139pXGekg0n6KdUOetdZ
Whe70YGNPwlyjWJTlIhMBBgRAGAMBQI+PgMdBQkZgGAAoJEIxxjTtQcuH17p4A
n3r1QpVC9yhnW2cSAjq+kr72GX0eAJ4295k16NxyEuFApmr1+0uUq/SlsQ==
=YJkx
-----END PGP PUBLIC KEY BLOCK-----
```

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

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

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

Distribution file	<code>mysql-standard-4.0.17-pc-linux-i686.tar.gz</code>
Signature file	<code>mysql-standard-4.0.17-pc-linux-i686.tar.gz.asc</code>

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

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

Example:

```
shell> gpg --verify mysql-standard-4.0.17-pc-linux-i686.tar.gz.asc
gpg: Warning: using insecure memory!
gpg: Signature made Mon 03 Feb 2003 08:50:39 PM MET
using DSA key ID 5072E1F5
gpg: Good signature from
"MySQL Package signing key (www.mysql.com) <build@mysql.com>"
```

The **Good signature** message indicates that everything is all right. You can ignore any **insecure memory** warning you might obtain.

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

### 2.1.4.3. Signature Checking Using RPM

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

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

Example:



```
shell> rpm --checksig MySQL-server-4.0.10-0.i386.rpm
MySQL-server-4.0.10-0.i386.rpm: md5 gpg OK
```

**Note:** If you are using RPM 4.1 and it complains about (GPG) NOT OK (MISSING KEYS: GPG#5072e1f5), even though you have imported the MySQL public build key into your own GPG keyring, you need to import the key into the RPM keyring first. RPM 4.1 no longer uses your personal GPG keyring (or GPG itself). Rather, it maintains its own keyring because it is a system-wide application and a user's GPG public keyring is a user-specific file. To import the MySQL public key into the RPM keyring, first obtain the key as described in [Section 2.1.4.2, “Signature Checking Using GnuPG”](#). Then use `rpm --import` to import the key. For example, if you have saved the public key in a file named `mysql_pubkey.asc`, import it using this command:

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

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

## 2.1.5. Installation Layouts

This section describes the default layout of the directories created by installing binary or source distributions provided by MySQL AB. A distribution provided by another vendor might use a layout different from those shown here.

On Windows, the default installation directory is `C:\mysql`. With MySQL version 4.1.5 and higher, this has changed to `C:\Program Files\MySQL\MySQL Server 4.1`, where 4.1 is the major version of the installation. The folder has the following subdirectories:

Directory	Contents of Directory
<code>bin</code>	Client programs and the <code>mysqld</code> server
<code>data</code>	Log files, databases
<code>Docs</code>	Documentation
<code>examples</code>	Example programs and scripts
<code>include</code>	Include (header) files
<code>lib</code>	Libraries
<code>scripts</code>	Utility scripts
<code>share</code>	Error message files

Installations created from MySQL AB's Linux RPM distributions result in files under the following system directories:

Directory	Contents of Directory
<code>/usr/bin</code>	Client programs and scripts
<code>/usr/sbin</code>	The <code>mysqld</code> server
<code>/var/lib/mysql</code>	Log files, databases
<code>/usr/share/doc/packages</code>	Documentation
<code>/usr/include/mysql</code>	Include (header) files
<code>/usr/lib/mysql</code>	Libraries
<code>/usr/share/mysql</code>	Error message and character set files
<code>/usr/share/sql-bench</code>	Benchmarks

On Unix, a `tar` file binary distribution is installed by unpacking it at the installation location you choose (typically `/usr/local/mysql`) and creates the following directories in that location:

Directory	Contents of Directory
<code>bin</code>	Client programs and the <code>mysqld</code> server
<code>data</code>	Log files, databases
<code>docs</code>	Documentation, ChangeLog
<code>include</code>	Include (header) files
<code>lib</code>	Libraries
<code>scripts</code>	<code>mysql_install_db</code>
<code>share/mysql</code>	Error message files
<code>sql-bench</code>	Benchmarks

A source distribution is installed after you configure and compile it. By default, the installation step installs files under `/usr/local`, in the following subdirectories:

Directory	Contents of Directory
<code>bin</code>	Client programs and scripts
<code>include/mysql</code>	Include (header) files
<code>info</code>	Documentation in Info format
<code>lib/mysql</code>	Libraries
<code>libexec</code>	The <code>mysqld</code> server
<code>share/mysql</code>	Error message files
<code>sql-bench</code>	Benchmarks and <code>crash-me</code> test
<code>var</code>	Databases and log files

Within its installation directory, the layout of a source installation differs from that of a binary installation in the following ways:

- The `mysqld` server is installed in the `libexec` directory rather than in the `bin` directory.
- The data directory is `var` rather than `data`.
- `mysql_install_db` is installed in the `bin` directory rather than in the `scripts` directory.
- The header file and library directories are `include/mysql` and `lib/mysql` rather than `include` and `lib`.

You can create your own binary installation from a compiled source distribution by executing the `scripts/make_binary_distribution` script from the top directory of the source distribution.

## 2.2. Standard MySQL Installation Using a Binary Distribution

The next several sections cover the installation of MySQL on platforms where we offer packages using the native packaging format of the respective platform. (This is also known as performing a “binary in-

stall.”) However, binary distributions of MySQL are available for many other platforms as well. See [Section 2.8, “Installing MySQL on Other Unix-Like Systems”](#), for generic installation instructions for these packages that apply to all platforms.

See [Section 2.1, “General Installation Issues”](#), for more information on what other binary distributions are available and how to obtain them.

## 2.3. Installing MySQL on Windows

A native Windows distribution of MySQL has been available from MySQL AB since version 3.21 and represents a sizable percentage of the daily downloads of MySQL. This section describes the process for installing MySQL on Windows.

With the release of MySQL 4.1.5, MySQL AB has introduced a new installer for the Windows version of MySQL, combined with a new GUI Configuration Wizard. This combination automatically installs MySQL, creates an option file, starts the server, and secures the default user accounts.

**Note:** If you are upgrading MySQL from an existing installation older than MySQL 4.1.5, you must first perform the the procedure described in [Section 2.3.14, “Upgrading MySQL on Windows”](#).

To run MySQL on Windows, you need the following:

- A 32-bit Windows operating system such as 9x, Me, NT, 2000, XP, or Windows Server 2003.

A Windows NT-based operating system (NT, 2000, XP, 2003) permits you to run the MySQL server as a service. The use of a Windows NT-based operating system is strongly recommended. See [Section 2.3.11, “Starting MySQL as a Windows Service”](#).

Generally, you should install MySQL on Windows using an account that has administrator rights. Otherwise, you may encounter problems with certain operations such as editing the `PATH` environment variable or accessing the [Service Control Manager](#).

- TCP/IP protocol support.
- Enough space on the hard drive to unpack, install, and create the databases in accordance with your requirements (generally a minimum of 200 megabytes is recommended.)

There may also be other requirements, depending on how you plan to use MySQL:

- If you plan to connect to the MySQL server via ODBC, you need a Connector/ODBC driver. See [Chapter 18, \*Connectors\*](#).
- If you need tables with a size larger than 4GB, install MySQL on an NTFS or newer filesystem. Don't forget to use `MAX_ROWS` and `AVG_ROW_LENGTH` when you create tables. See [Section 13.1.5, “CREATE TABLE Syntax”](#).

MySQL for Windows is available in several distribution formats:

- Binary distributions are available that contain a setup program that installs everything you need so that you can start the server immediately. Another binary distribution format contains an archive that you simply unpack in the installation location and then configure yourself. For details, see [Section 2.3.1, “Choosing An Installation Package”](#).
- The source distribution contains all the code and support files for building the executables using the

Visual Studio 7.1 compiler system.

Generally speaking, you should use a binary distribution that includes an installer. It is simpler to use than the others, and you need no additional tools to get MySQL up and running. The installer for the Windows version of MySQL, combined with a GUI Configuration Wizard, automatically installs MySQL, creates an option file, starts the server, and secures the default user accounts.

The following section describes how to install MySQL on Windows using a binary distribution. To use an installation package that does not include an installer, follow the procedure described in [Section 2.3.5, “Installing MySQL from a Noinstall Zip Archive”](#). To install using a source distribution, see [Section 2.9.6, “Installing MySQL from Source on Windows”](#).

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

### 2.3.1. Choosing An Installation Package

Starting with MySQL version 4.1.5, there are three install packages to choose from when installing MySQL on Windows. The Packages are as follows:

- **The Essentials Package:** This package has a filename similar to `mysql-essential-4.1.13a-win32.msi` and contains the minimum set of files needed to install MySQL on Windows, including the Configuration Wizard. This package does not include optional components such as the embedded server and benchmark suite.
- **The Complete Package:** This package has a filename similar to `mysql-4.1.13a-win32.zip` and contains all files needed for a complete Windows installation, including the Configuration Wizard. This package includes optional components such as the embedded server and benchmark suite.
- **The Noinstall Archive:** This package has a filename similar to `mysql-noinstall-4.1.13a-win32.zip` and contains all the files found in the Complete install package, with the exception of the Configuration Wizard. This package does not include an automated installer, and must be manually installed and configured.

The Essentials package is recommended for most users. It is provided as an `.msi` file for use with the Windows Installer. The Complete and Noinstall distributions are packaged as Zip archives. To use them, you must have a tool that can unpack `.zip` files.

Your choice of install package affects the installation process you must follow. If you choose to install either the Essentials or Complete install packages, see [Section 2.3.2, “Installing MySQL with the Automated Installer”](#). If you choose to install MySQL from the Noinstall archive, see [Section 2.3.5, “Installing MySQL from a Noinstall Zip Archive”](#).

### 2.3.2. Installing MySQL with the Automated Installer

Starting with MySQL 4.1.5, users can use the new MySQL Installation Wizard and MySQL Configuration Wizard to install MySQL on Windows. The MySQL Installation Wizard and MySQL Configuration Wizard are designed to install and configure MySQL in such a way that new users can immediately get started using MySQL.

The MySQL Installation Wizard and MySQL Configuration Wizard are available in the Essentials and Complete install packages. They are recommended for most standard MySQL installations. Exceptions include users who need to install multiple instances of MySQL on a single server host and advanced users who want complete control of server configuration.

If you are installing a version of MySQL prior to MySQL 4.1.5, please follow the instructions for installing MySQL from the Noinstall installation package. See [Section 2.3.5, “Installing MySQL from a Noinstall Zip Archive”](#).

## 2.3.3. Using the MySQL Installation Wizard

### 2.3.3.1. Introduction to the Installation Wizard

MySQL Installation Wizard is an installer for the MySQL server that uses the latest installer technologies for Microsoft Windows. The MySQL Installation Wizard, in combination with the MySQL Configuration Wizard, allows a user to install and configure a MySQL server that is ready for use immediately after installation.

The MySQL Installation Wizard is the standard installer for all MySQL server distributions, version 4.1.5 and higher. Users of previous versions of MySQL need to shut down and remove their existing MySQL installations manually before installing MySQL with the MySQL Installation Wizard. See [Section 2.3.3.7, “Upgrading MySQL with the Installation Wizard”](#), for more information on upgrading from a previous version.

Microsoft has included an improved version of their Microsoft Windows Installer (MSI) in the recent versions of Windows. MSI has become the de-facto standard for application installations on Windows 2000, Windows XP, and Windows Server 2003. The MySQL Installation Wizard makes use of this technology to provide a smoother and more flexible installation process.

The Microsoft Windows Installer Engine was updated with the release of Windows XP; those using a previous version of Windows can reference [this Microsoft Knowledge Base article](http://support.microsoft.com/default.aspx?scid=kb;EN-US;292539) [http://support.microsoft.com/default.aspx?scid=kb;EN-US;292539] for information on upgrading to the latest version of the Windows Installer Engine.

In addition, Microsoft has introduced the WiX (Windows Installer XML) toolkit recently. This is the first highly acknowledged Open Source project from Microsoft. We have switched to WiX because it is an Open Source project and it allows us to handle the complete Windows installation process in a flexible manner using scripts.

Improving the MySQL Installation Wizard depends on the support and feedback of users like you. If you find that the MySQL Installation Wizard is lacking some feature important to you, or if you discover a bug, please report it in our bugs database using the instructions given in [Section 1.8, “How to Report Bugs or Problems”](#).

### 2.3.3.2. Downloading and Starting the MySQL Installation Wizard

The MySQL installation packages can be downloaded from <http://dev.mysql.com/downloads/>. If the package you download is contained within a Zip archive, you need to extract the archive first.

The process for starting the wizard depends on the contents of the installation package you download. If there is a `setup.exe` file present, double-click it to start the installation process. If there is an `.msi` file present, double-click it to start the installation process.

### 2.3.3.3. Choosing an Install Type

There are three installation types available: **Typical**, **Complete**, and **Custom**.

The **Typical** installation type installs the MySQL server, the `mysql` command-line client, and the command-line utilities. The command-line clients and utilities include `mysqldump`, `myisamchk`, and several other tools to help you manage the MySQL server.

The **Complete** installation type installs all components included in the installation package. The full in-

stallation package includes components such as the embedded server library, the benchmark suite, support scripts, and documentation.

The **Custom** installation type gives you complete control over which packages you wish to install and the installation path that is used. See [Section 2.3.3.4, “The Custom Install Dialog”](#), for more information on performing a custom install.

If you choose the **Typical** or **Complete** installation types and click the [Next](#) button, you advance to the confirmation screen to verify your choices and begin the installation. If you choose the **Custom** installation type and click the [Next](#) button, you advance to the custom installation dialog, described in [Section 2.3.3.4, “The Custom Install Dialog”](#).

### 2.3.3.4. The Custom Install Dialog

If you wish to change the installation path or the specific components that are installed by the MySQL Installation Wizard, choose the **Custom** installation type.

A tree view on the left side of the custom install dialog lists all available components. Components that are not installed have a red X icon; components that are installed have a gray icon. To change whether a component is installed, click on that component's icon and choose a new option from the drop-down list that appears.

You can change the default installation path by clicking the [Change . . .](#) button to the right of the displayed installation path.

After choosing your installation components and installation path, click the [Next](#) button to advance to the confirmation dialog.

### 2.3.3.5. The Confirmation Dialog

Once you choose an installation type and optionally choose your installation components, you advance to the confirmation dialog. Your installation type and installation path are displayed for you to review.

To install MySQL if you are satisfied with your settings, click the [Install](#) button. To change your settings, click the [Back](#) button. To exit the MySQL Installation Wizard without installing MySQL, click the [Cancel](#) button.

After installation is complete, you have the option of registering with the MySQL web site. Registration gives you access to post in the MySQL forums at [forums.mysql.com](http://forums.mysql.com) [<http://forums.mysql.com>], along with the ability to report bugs at [bugs.mysql.com](http://bugs.mysql.com) [<http://bugs.mysql.com>] and to subscribe to our newsletter. The final screen of the installer provides a summary of the installation and gives you the option to launch the MySQL Configuration Wizard, which you can use to create a configuration file, install the MySQL service, and configure security settings.

### 2.3.3.6. Changes Made by MySQL Installation Wizard

Once you click the [Install](#) button, the MySQL Installation Wizard begins the installation process and makes certain changes to your system which are described in the sections that follow.

#### Changes to the Registry

The MySQL Installation Wizard creates one Windows registry key in a typical install situation, located in `HKEY_LOCAL_MACHINE\SOFTWARE\MySQL AB`.

The MySQL Installation Wizard creates a key named after the major version of the server that is being installed, such as `MySQL Server 4.1`. It contains two string values, `Location` and `Version`. The `Location` string contains the path to the installation directory. In a default installation it contains `C:\Program Files\MySQL\MySQL Server 4.1\`. The `Version` string contains the release

number. For example, for an installation of MySQL Server 4.1.5, the key contains a value of `4.1.5`.

These registry keys are used to help external tools identify the installed location of the MySQL server, preventing a complete scan of the hard-disk to determine the installation path of the MySQL server. The registry keys are not required to run the server, and if you install MySQL using the `noinstall` Zip archive, the registry keys are not created.

### Changes to the Start Menu

The MySQL Installation Wizard creates a new entry in the Windows Start menu under a common MySQL menu heading named after the major version of MySQL that you have installed. For example, if you install MySQL 4.1, the MySQL Installation Wizard creates a MySQL Server 4.1 section in the start menu.

The following entries are created within the new Start menu section:

- **MySQL Command Line Client:** This is a shortcut to the `mysql` command-line client and is configured to connect as the `root` user. The shortcut prompts for a `root` user password when you connect.
- **MySQL Server Instance Config Wizard:** This is a shortcut to the MySQL Configuration Wizard. Use this shortcut to configure a newly installed server, or to reconfigure an existing server.
- **MySQL Documentation:** This is a link to the MySQL server documentation that is stored locally in the MySQL server installation directory. This option is not available when the MySQL server is installed from the Essentials installation package.

### Changes to the File System

The MySQL Installation Wizard by default installs the MySQL server to `C:\Program Files\MySQL\MySQL Server 4.1`, where *Program Files* is the default location for applications in your system, and `4.1` is the major version of your MySQL server. This is the new location for the MySQL server, replacing the former default location of `c:\mysql`.

By default, all MySQL applications are stored in a common directory at `C:\Program Files\MySQL`, where *Program Files* is the default location for applications in your Windows installation. A typical MySQL installation on a developer machine might look like this:

```
C:\Program Files\MySQL\MySQL Server 4.1
C:\Program Files\MySQL\MySQL Administrator 1.0
C:\Program Files\MySQL\MySQL Query Browser 1.0
```

This approach makes it easier to manage and maintain all MySQL applications installed on a particular system.

## 2.3.3.7. Upgrading MySQL with the Installation Wizard

From MySQL version 4.1.5, the new MySQL Installation Wizard can perform server upgrades automatically using the upgrade capabilities of MSI. That means you do not need to remove a previous installation manually before installing a new release. The installer automatically shuts down and removes the previous MySQL service before installing the new version.

Automatic upgrades are available only when upgrading between installations that have the same major and minor version numbers. For example, you can upgrade automatically from MySQL 4.1.5 to MySQL 4.1.6, but not from MySQL 4.1 to MySQL 5.0.

*If you are upgrading MySQL version 4.1.4 or earlier to version 4.1.5 or later, you must first manually shut down and remove the older installation before upgrading. Be sure to back up your databases before*

*performing such an upgrade, so that you can restore the databases after the upgrade is completed. It is always recommended that you back up your data before performing any upgrades.*

See [Section 2.3.14, “Upgrading MySQL on Windows”](#).

## 2.3.4. Using the Configuration Wizard

### 2.3.4.1. Introduction to the Configuration Wizard

The MySQL Configuration Wizard helps automate the process of configuring your server under Windows. The MySQL Configuration Wizard creates a custom `my.ini` file by asking you a series of questions and then applying your responses to a template to generate a `my.ini` file that is tuned to your installation.

The MySQL Configuration Wizard is included with the MySQL server starting with MySQL version 4.1.5, but is designed to work with MySQL servers versions 4.1 and higher. The MySQL Configuration Wizard is currently available for Windows users only.

The MySQL Configuration Wizard is to a large extent the result of feedback that MySQL AB has received from many users over a period of several years. However, if you find that it lacks some feature important to you, please report it in our bugs database using the instructions given in [Section 1.8, “How to Report Bugs or Problems”](#).

### 2.3.4.2. Starting the MySQL Configuration Wizard

The MySQL Configuration Wizard is typically launched from the MySQL Installation Wizard, as the MySQL Installation Wizard exits. You can also launch the MySQL Configuration Wizard by clicking the MySQL Server Instance Config Wizard entry in the MySQL section of the Windows Start menu.

Alternatively, you can navigate to the `bin` directory of your MySQL installation and launch the `MySQLInstanceConfig.exe` file directly.

### 2.3.4.3. Choosing a Maintenance Option

If the MySQL Configuration Wizard detects an existing `my.ini` file, you have the option of either reconfiguring your existing server, or removing the server instance by deleting the `my.ini` file and stopping and removing the MySQL service.

To reconfigure an existing server, choose the Re-configure Instance option and click the `Next` button. Your existing `my.ini` file is renamed to `mytimestamp.ini.bak`, where `timestamp` is the date and time at which the existing `my.ini` file was created. To remove the existing server instance, choose the Remove Instance option and click the `Next` button.

If you choose the Remove Instance option, you advance to a confirmation window. Click the `Execute` button. The MySQL Configuration Wizard stops and removes the MySQL service, and then deletes the `my.ini` file. The server installation and its `data` folder are not removed.

If you choose the Re-configure Instance option, you advance to the Configuration Type dialog where you can choose the type of installation that you wish to configure.

### 2.3.4.4. Choosing a Configuration Type

When you start the MySQL Configuration Wizard for a new MySQL installation, or choose the Re-configure Instance option for an existing installation, you advance to the Configuration Type dialog.

There are two configuration types available: Detailed Configuration and Standard Configuration. The Standard Configuration option is intended for new users who want to get started with MySQL quickly without having to make a lot of decisions about server configuration. The [Detailed Configura-](#)



tion option is intended for advanced users who want more fine-grained control of server configuration.

If you are new to MySQL and need a server configured as a single-user developer machine, the Standard Configuration should suit your needs. Choosing the Standard Configuration option causes the MySQL Configuration Wizard to automatically set all configuration options with the exception of the [Service Options](#) and [Security Options](#).

The [Standard Configuration](#) sets options that may be incompatible with systems where there are existing MySQL installations. If you have an existing MySQL installation on your system in addition to the installation you wish to configure, the [Detailed Configuration](#) option is recommended.

To complete the Standard Configuration, please refer to the sections on Service Options and Security Options in [Section 2.3.4.11, “The Service Options Dialog”](#), and [Section 2.3.4.12, “The Security Options Dialog”](#), respectively.

### 2.3.4.5. The Server Type Dialog

There are three different server types available to choose from. The server type that you choose affects the decisions that the MySQL Configuration Wizard makes with regard to memory, disk, and processor usage.

- **Developer Machine:** Choose this option for a typical desktop workstation where MySQL is intended only for personal use. It is assumed that many other desktop applications are running. The MySQL server is configured to use minimal system resources.
- **Server Machine:** Choose this option for a server machine where the MySQL server is running alongside other server applications such as FTP, email, and Web servers. The MySQL server is configured to use a medium portion of the system resources.
- **Dedicated MySQL Server Machine:** Choose this option for a server machine that is intended to run only the MySQL server. It is assumed that no other applications are running. The MySQL server is configured to use all available system resources.

### 2.3.4.6. The Database Usage Dialog

The Database Usage dialog allows you to indicate the storage engines that you expect to use when creating MySQL tables. The option you choose determines whether the [InnoDB](#) storage engine is available and what percentage of the server resources are available to [InnoDB](#).

- **Multifunctional Database:** This option enables both the [InnoDB](#) and [MyISAM](#) storage engines, and divides resources evenly between the two. This option is recommended for users who use both table handlers on a regular basis.
- **Transactional Database Only:** This option enables both the [InnoDB](#) and [MyISAM](#) storage engines but dedicates most server resources to the [InnoDB](#) storage engine. This option is recommended for users who use [InnoDB](#) almost exclusively and make only minimal use of [MyISAM](#).
- **Non-Transactional Database Only:** This option disables the [InnoDB](#) storage engine completely, and dedicates all server resources to the [MyISAM](#) storage engine. This option is recommended for users who do not use [InnoDB](#).

### 2.3.4.7. The InnoDB Tablespace Dialog

Some users may want to locate the [InnoDB](#) tablespace files in a location other than the MySQL server

data directory. Placing the tablespace files in a separate location can be desirable if your system has available a storage device available with higher capacity or higher performance, such as a RAID storage system.

To change the default location for the [InnoDB](#) tablespace files, choose a new drive from the drop-down list of drive letters and choose a new path from the drop-down list of paths. To create a custom path, click the `...` button.

If you are modifying the configuration of an existing server, you must click the [Modify](#) button before you change the path. In this situation you must move existing tablespace files to the new location manually before starting the server.

### 2.3.4.8. The Concurrent Connections Dialog

To prevent the server from running out of resources, it is important to limit the number of concurrent connections to the MySQL server that can be established. The Concurrent Connections dialog allows you to choose the expected usage of your server, and sets the limit for concurrent connections accordingly. It is also possible to manually set the concurrent connection limit.

- **Decision Support (DSS)/OLAP:** Choose this option if the server does not require a large number of concurrent connections. The maximum number of connections is set at 100, with an average of 20 concurrent connections assumed.
- **Online Transaction Processing (OLTP):** Choose this option if the server requires a large number of concurrent connections. The maximum number of connections is set at 500.
- **Manual Setting:** Choose this option to set the maximum number of concurrent connections to the server manually. Choose the number of concurrent connections from the drop-down box provided, or type the maximum number of connections into the drop-down box if the number you desire is not listed.

### 2.3.4.9. The Networking and Strict Mode Options Dialog

Use the Networking Options dialog to enable or disable TCP/IP networking and to configure the port number that is used to connect to the MySQL server.

TCP/IP networking is enabled by default. To disable TCP/IP networking, uncheck the box next to the Enable TCP/IP Networking option.

Port 3306 is used by default. To change the port used to access MySQL, choose a new port number from the drop-down box or type a new port number directly into the drop-down box. If the port number you choose is in use, you are prompted to confirm your choice of port number.

Set the Server SQL Mode to either enable or disable strict mode. Enabling strict mode (default) makes MySQL behave more like other database management systems. *If you run applications that rely on MySQL's old "forgiving" behavior, make sure to either adapt those applications or to disable strict mode.* For more information about strict mode, see [Section 5.2.5, "SQL Modes"](#).

### 2.3.4.10. The Character Set Dialog

The MySQL server supports multiple character sets and it is possible to set a default server character set that is applied to all tables, columns, and databases unless overridden. Use the Character Set dialog to change the default character set of the MySQL server.

- **Standard Character Set:** Choose this option if you want to use [latin1](#) as the default server character set. [latin1](#) is used for English and many Western European languages.

- **Best Support For Multilingualism:** Choose this option if you want to use `utf8` as the default server character set. This is a Unicode character set that can store characters from many different languages.
- **Manually Selected Default Character Set / Collation:** Choose this option if you want to pick the server's default character set manually. Choose the desired character set from the provided drop-down list.

### 2.3.4.11. The Service Options Dialog

On Windows NT-based platforms, the MySQL server can be installed as a Windows service. When installed this way, the MySQL server can be started automatically during system startup, and even restarted automatically by Windows in the event of a service failure.

The MySQL Configuration Wizard installs the MySQL server as a service by default, using the service name `MySQL`. If you do not wish to install the service, un-check the box next to the `Install As Windows Service` option. You can change the service name by picking a new service name from the drop-down box provided or by typing a new service name into the drop-down box.

To install the MySQL server as a service but not have it started automatically at startup, un-check the box next to the `Launch the MySQL Server automatically` option.

### 2.3.4.12. The Security Options Dialog

*It is strongly recommended that you set a `root` password for your MySQL server, and the MySQL Configuration Wizard requires by default that you do so. If you do not wish to set a `root` password, un-check the box next to the `Modify Security Settings` option.*

To set the `root` password, enter the desired password into both the `New root password` and `Confirm` boxes. If you are reconfiguring an existing server, you need to enter the existing `root` password into the `Current root password` box.

To prevent `root` logins from across the network, check the box next to the `Root may only connect from localhost` option. This increases the security of your `root` account.

To create an anonymous user account, check the box next to the `Create An Anonymous Account` option. Creating an anonymous account can decrease server security and cause login and permission difficulties and is not recommended.

### 2.3.4.13. The Confirmation Dialog

The final dialog in the MySQL Configuration Wizard is the Confirmation Dialog. To start the configuration process, click the `Execute` button. To return to a previous dialog, click the `Back` button. To exit the MySQL Configuration Wizard without configuring the server, click the `Cancel` button.

After you click the `Execute` button, the MySQL Configuration Wizard performs a series of tasks and displays the progress onscreen as the tasks are performed.

The MySQL Configuration Wizard firsts determines various configuration file options based on your choices using a template prepared by MySQL AB developers and engineers. This template is named `my-template.ini` and is located in your server installation directory.

The MySQL Configuration Wizard then writes these options to a `my.ini` file. The final location of the `my.ini` file is displayed next to the `Write configuration file` task.

If you chose to create a service for the MySQL server, the MySQL Configuration Wizard creates and starts the service. If you are reconfiguring an existing service, the MySQL Configuration Wizard restarts

the service to apply your configuration changes.

If you chose to set a `root` password, the MySQL Configuration Wizard connects to the server, sets your new `root` password and applies any other security settings you may have selected.

After the MySQL Configuration Wizard has completed its tasks, it displays a summary. Click the Finish button to exit the MySQL Configuration Wizard.

### 2.3.4.14. The Location of the `my.ini` File

In MySQL installations prior to version 4.1.5 it was customary to name the server configuration file `my.cnf` or `my.ini` and locate the file either at `c:\my.cnf` or `c:\Windows\my.ini`.

The new MySQL Configuration Wizard places the `my.ini` file in the installation directory of the MySQL server. This helps associate configuration files with particular server instances.

To ensure that the MySQL server knows where to look for the `my.ini` file, an argument similar to this is passed to the MySQL server as part of the service installation:

```
--defaults-file="C:\Program Files\MySQL\MySQL Server 4.1\my.ini"
```

Here, `C:\Program Files\MySQL\MySQL Server 4.1` is replaced with the installation path to the MySQL Server. The `--defaults-file` option instructs the MySQL server to read the specified file for configuration options when it starts.

### 2.3.4.15. Editing the `my.ini` File

To modify the `my.ini` file, open it with a text editor and make any necessary changes. You can also modify the server configuration with the [MySQL Administrator](http://www.mysql.com/products/administrator/) [http://www.mysql.com/products/administrator/] utility.

MySQL clients and utilities such as the `mysql` and `mysqldump` command-line clients are not able to locate the `my.ini` file located in the server installation directory. To configure the client and utility applications, create a new `my.ini` file in the `C:\WINDOWS` or `C:\WINNT` directory (whichever is applicable to your Windows version).

## 2.3.5. Installing MySQL from a Noinstall Zip Archive

Users who are installing from the Noinstall package, or who are installing a version of MySQL prior to 4.1.5 can use the instructions in this section to manually install MySQL. If you are installing a version prior to 4.1.5 with an install package that includes a Setup program, substitute running the Setup program for extracting the archive.

The process for installing MySQL from a Zip archive is as follows:

1. Extract the archive to the desired install directory
2. Create an option file
3. Choose a MySQL server type
4. Start the MySQL server
5. Secure the default user accounts

This process is described in the sections that follow.

## 2.3.6. Extracting the Install Archive

To install MySQL manually, do the following:

1. If you are upgrading from a previous version please refer to [Section 2.3.14, “Upgrading MySQL on Windows”](#), before beginning the upgrade process.
2. If you are using a Windows NT-based operating system such as Windows NT, Windows 2000, Windows XP, or Windows Server 2003, make sure that you are logged in as a user with administrator privileges.
3. Choose an installation location. Traditionally, the MySQL server has been installed at `C:\mysql`. The MySQL Installation Wizard installs MySQL under `C:\Program Files\MySQL`. If you do not install MySQL in `C:\mysql`, you must specify the path to the install directory during startup or in an option file. See [Section 2.3.7, “Creating an Option File”](#).
4. Extract the install archive to the chosen installation location using your preferred Zip archive tool. Some tools may extract the archive to a folder within your chosen installation location. If this occurs, you can move the contents of the subfolder into the chosen installation location.

## 2.3.7. Creating an Option File

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

- The installation or data directory locations differ from the default locations (`C:\mysql` and `C:\mysql\data`).
- You need to tune the server settings. For example, to use the [InnoDB](#) transactional tables in MySQL 3.23, you must manually add some extra lines to the option file, as described in [Section 14.2.4, “InnoDB Configuration”](#). (As of MySQL 4.0, [InnoDB](#) creates its data files and log files in the data directory by default. This means you need not configure [InnoDB](#) explicitly. You may still do so if you wish, and an option file is also useful in this case.)

When the MySQL server starts on Windows, it looks for options in two files: the `my.ini` file in the Windows directory, and the `C:\my.cnf` file. The Windows directory typically is named something like `C:\WINDOWS` or `C:\WINNT`. You can determine its exact location from the value of the `WINDIR` environment variable using the following command:

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

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

You can also make use of the example option files included with your MySQL distribution; see [Section 4.3.2.1, “Preconfigured Option Files”](#).

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

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

Note that Windows pathnames are specified in option files using forward slashes rather than backslashes. If you do use backslashes, you must double them:

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

On Windows, the MySQL installer places the data directory directly under the directory where you install MySQL. If you would like to use a data directory in a different location, you should copy the entire contents of the `data` directory to the new location. For example, by default, the installer places MySQL in `C:\mysql`, and the data directory in `C:\mysql\data`. If you want to use `E:\mydata` as the data directory, you must do two things:

- Move the data directory from `C:\mysql\data` to `E:\mydata`.
- Use a `--datadir` option to specify the new data directory location each time you start the server.

## 2.3.8. Selecting a MySQL Server type

Starting with MySQL 3.23.38, the Windows distribution includes both the normal and the MySQL-Max server binaries.

Up through the early releases of MySQL 4.1, the servers included in Windows distributions are named like this:

Binary	Description
<code>mysqld</code>	Compiled with full debugging and automatic memory allocation checking, and <code>InnoDB</code> and <code>BDB</code> tables.
<code>mysqld-opt</code>	Optimized binary. From version 4.0 on, <code>InnoDB</code> is enabled. Before 4.0, this server includes no transactional table support.
<code>mysqld-nt</code>	Optimized binary for Windows NT, 2000, and XP with support for named pipes.
<code>mysqld-max</code>	Optimized binary with <code>InnoDB</code> and <code>BDB</code> support.
<code>mysqld-max-nt</code>	Like <code>mysqld-max</code> , but compiled with support for named pipes.

We have found that the server with the most generic name (`mysqld`) is the one that many users are likely to choose by default. However, that is also the server that results in the highest memory and CPU use due to the inclusion of full debugging support. The server named `mysqld-opt` is a better general-use server choice to make instead if you do not need debugging support and do not want the maximal feature set offered by the `-max` servers or named pipe support offered by the `-nt` servers.

To make it less likely that the debugging server would be chosen inadvertently, some name changes

were made from MySQL 4.1.2 to 4.1.4: `mysqld` has been renamed to `mysqld-debug` and `mysqld-opt` has been renamed to `mysqld`. Thus, the server that includes debugging support indicates that in its name, and the server named `mysqld` is an efficient default choice. The other servers still have their same names. The resulting servers are named like this:

Binary	Description
<code>mysqld-debug</code>	Compiled with full debugging and automatic memory allocation checking, and InnoDB and BDB tables.
<code>mysqld</code>	Optimized binary with InnoDB support.
<code>mysqld-nt</code>	Optimized binary for Windows NT, 2000, and XP with support for named pipes.
<code>mysqld-max</code>	Optimized binary with support for InnoDB and BDB tables.
<code>mysqld-max-nt</code>	Like <code>mysqld-max</code> , but compiled with support for named pipes.

The name changes were not both instituted at the same time. If you have MySQL 4.1.2 or 4.1.3, it might be that you have a server named `mysqld-debug` but not one named `mysqld`. In this case, you should have a server `mysqld-opt`, which you should choose as your default server unless you need maximal features, named pipes, or debugging support.

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

As of MySQL 4.0, all Windows servers have support for symbolic linking of database directories. Before MySQL 4.0, only the debugging and Max server versions include this feature.

MySQL supports TCP/IP on all Windows platforms. The `mysqld-nt` and `mysqld-max-nt` servers support named pipes on Windows NT, 2000, XP, and 2003. However, the default is to use TCP/IP regardless of the platform. (Named pipes are slower than TCP/IP in many Windows configurations.)

Use of named pipes is subject to these conditions:

- Starting from MySQL 3.23.50, named pipes are enabled only if you start the server with the `--enable-named-pipe` option. It is necessary to use this option explicitly because some users have experienced problems shutting down the MySQL server when named pipes were used.
- Named-pipe connections are allowed only by the `mysqld-nt` or `mysqld-max-nt` servers, and only if the server is run on a version of Windows that supports named pipes (NT, 2000, XP, 2003).
- These servers can be run on Windows 98 or Me, but only if TCP/IP is installed; named-pipe connections cannot be used.
- These servers cannot be run on Windows 95.

**Note:** Most of the examples in this manual use `mysqld` as the server name. If you choose to use a different server, such as `mysqld-nt`, make the appropriate substitutions in the commands that are shown in the examples.

## 2.3.9. Starting the Server for the First Time

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

The information here applies primarily if you installed MySQL using the `Noinstall` version, or if

you wish to configure and test MySQL manually rather than with the GUI tools.

On Windows 95, 98, or Me, MySQL clients always connect to the server using TCP/IP. (This allows any machine on your network to connect to your MySQL server.) Because of this, you must make sure that TCP/IP support is installed on your machine before starting MySQL. You can find TCP/IP on your Windows CD-ROM.

Note that if you are using an old Windows 95 release (for example, OSR2), it is likely that you have an old Winsock package; MySQL requires Winsock 2. You can get the newest Winsock from <http://www.microsoft.com/>. Windows 98 has the new Winsock 2 library, so it is unnecessary to update the library.

On NT-based systems such as Windows NT, 2000, XP, or 2003, clients have two options. They can use TCP/IP, or they can use a named pipe if the server supports named-pipe connections. For MySQL to work with TCP/IP on Windows NT 4, you must install service pack 3 (or newer).

In MySQL versions 4.1 and higher, Windows servers also support shared-memory connections if the server is started with the `--shared-memory` option. Clients can connect through shared memory by using the `--protocol=memory` option.

For information about which server binary to run, see [Section 2.3.8, “Selecting a MySQL Server type”](#).

The examples in these sections assume that MySQL is installed under the default location of `C:\mysql`. Adjust the pathnames shown in the examples if you have MySQL installed in a different location.

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

To start the server, enter this command:

```
C:\> C:\mysql\bin\mysqld --console
```

For a server that includes [InnoDB](#) support, you should see the messages similar to those following as it starts (the pathnames and sizes may differ):

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

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

```
mysqld: ready for connections
Version: '4.0.14-log' socket: '' port: 3306
```

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

If you omit the `--console` option, the server writes diagnostic output to the error log in the data dir-



ectory (`C:\mysql\data` by default). The error log is the file with the `.err` extension.

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

## 2.3.10. Starting MySQL from the Windows Command Line

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

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

```
C:\> "C:\Program Files\MySQL\MySQL Server 4.1\bin\mysqld"
```

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

On non-NT versions of Windows, this command starts `mysqld` in the background. That is, after the server starts, you should see another command prompt. If you start the server this way on Windows NT, 2000, XP, or 2003, the server runs in the foreground and no command prompt appears until the server exits. Because of this, you should open another console window to run client programs while the server is running.

You can stop the MySQL server by executing this command:

```
C:\> "C:\Program Files\MySQL\MySQL Server 4.1\bin\mysqladmin" -u root shutdown
```

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

This command invokes the MySQL administrative utility `mysqladmin` to connect to the server and tell it to shut down. The command connects as the MySQL `root` user, which is the default administrative account in the MySQL grant system. Note that users in the MySQL grant system are wholly independent from any login users under Windows.

If `mysqld` does not start, check the error log to see whether the server wrote any messages there to indicate the cause of the problem. The error log is located in the `C:\mysql\data` directory. It is the file with a suffix of `.err`. You can also try to start the server as `mysqld --console`; in this case, you may get some useful information on the screen that may help solve the problem.

The last option is to start `mysqld` with the `--standalone` and `--debug` options. In this case, `mysqld` writes a log file `C:\mysqld.trace` that should contain the reason why `mysqld` doesn't start. See [Section E.1.2](#), “Creating Trace Files”.

Use `mysqld --verbose --help` to display all the options that `mysqld` understands. (Prior to MySQL 4.1, omit the `--verbose` option.)

## 2.3.11. Starting MySQL as a Windows Service

On the NT family (Windows NT, 2000, XP, 2003), the recommended way to run MySQL is to install it as a Windows service. With the MySQL server installed as a service, Windows starts and stops it server automatically when Windows starts and stops. A MySQL server installed as a service can also be controlled from the command line using `NET` commands, or with the graphical `Services` utility.

The `Services` utility (the Windows `Service Control Manager`) can be found in the Windows Control Panel (under Administrative Tools on Windows 2000, XP, and Server 2003). To avoid conflicts, it is advisable to close the `Services` utility while performing server installation or removal operations

from the command line.

Before installing MySQL as a Windows service, you should first stop the current server if it is running by using the following command:

```
C:\> C:\mysql\bin\mysqladmin -u root shutdown
```

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

This command invokes the MySQL administrative utility `mysqladmin` to connect to the server and tell it to shut down. The command connects as the MySQL `root` user, which is the default administrative account in the MySQL grant system. Note that users in the MySQL grant system are wholly independent from any login users under Windows.

Install the server as a service using this command:

```
C:\> C:\mysql\bin\mysqld --install
```

The service-installation command does not start the server. Instructions for that are given later in this section.

Before MySQL 4.0.2, no command-line arguments can be given following the `--install` option. MySQL 4.0.2 and up offers limited support for additional arguments:

- You can specify a service name immediately following the `--install` option. The default service name is `MySQL`.
- As of MySQL 4.0.3, if a service name is given, it can be followed by a single option. By convention, this should be `--defaults-file=file_name` to specify the name of an option file from which the server should read options when it starts.

It is possible to use a single option other than `--defaults-file`, but this is discouraged. `--defaults-file` is more flexible because it enables you to specify multiple startup options for the server by placing them in the named option file. Also, in MySQL 5.0, use of an option different from `--defaults-file` is not supported until 5.0.3.

- As of MySQL 5.0.1, you can also specify a `--local-service` option following the service name. This causes the server to run using the `LocalService` Windows account that has limited system privileges. This account is available only for Windows XP or newer. If both `--defaults-file` and `--local-service` are given following the service name, they can be in any order.

For a MySQL server that is installed as a Windows service, the following rules determine the service name and option files that the server uses:

- If the service-installation command specifies no service name or the default service name (`MySQL`) following the `--install` option, the server uses the a service name of `MySQL` and reads options from the `[mysqld]` group in the standard option files.
- If the service-installation command specifies a service name other than `MySQL` following the `--install` option, the server uses that service name. It reads options from the group that has the same name as the service, and reads options from the standard option files.

As of MySQL 4.0.17, the server also reads options from the `[mysqld]` group from the standard option files. This allows you to use the `[mysqld]` group for options that should be used by all

MySQL services, and an option group with the same name as a service for use by the server installed with that service name.

- If the service-installation command specifies a `--defaults-file` option after the service name, the server reads options only from the `[mysqld]` group of the named file and ignores the standard option files.

As a more complex example, consider the following command:

```
C:\> C:\mysql\bin\mysqld --install MySQL --defaults-file=C:\my-opts.cnf
```

Here, the default service name (`MySQL`) is given after the `--install` option. If no `--defaults-file` option had been given, this command would have the effect of causing the server to read the `[mysqld]` group from the standard option files. However, because the `--defaults-file` option is present, the server reads options from the `[mysqld]` option group, and only from the named file.

You can also specify options in Start parameters in the Windows `Services` utility before you start the MySQL service.

**Note:** Prior to MySQL 4.0.17, a server installed as a Windows service has problems starting if its path-name or the service name contains spaces. For this reason, with older versions, avoid installing MySQL in a directory such as `C:\Program Files` or using a service name containing spaces.

Once a MySQL server has been installed as a service, Windows starts the service automatically whenever Windows starts. The service also can be started immediately from the `Services` utility, or by using a `NET START MySQL` command. The `NET` command is not case sensitive.

When run as a service, `mysqld` has no access to a console window, so no messages can be seen there. If `mysqld` does not start, check the error log to see whether the server wrote any messages there to indicate the cause of the problem. The error log is located in the MySQL data directory (for example, `C:\mysql\data`). It is the file with a suffix of `.err`.

When a MySQL server has been installed as a service, and the service is running, Windows stops the service automatically when Windows shuts down. The server also can be stopped manually by using the `Services` utility, the `NET STOP MySQL` command, or the `mysqladmin shutdown` command.

From MySQL 3.23.44 on, you have the choice of installing the server as a `Manual` service if you do not wish the service to be started automatically during the boot process. To do this, use the `--install-manual` option rather than the `--install` option:

```
C:\> C:\mysql\bin\mysqld --install-manual
```

To remove a server that is installed as a service, first stop it if it is running by executing `NET STOP MySQL`. Then use the `--remove` option to remove it:

```
C:\> C:\mysql\bin\mysqld --remove
```

For MySQL versions older than 3.23.49, one problem with automatic MySQL service shutdown is that Windows waited only for a few seconds for the shutdown to complete, and then killed the database server process if the time limit was exceeded. This had the potential to cause problems. (For example, the `InnoDB` storage engine would have to perform crash recovery at the next startup.) Starting from MySQL 3.23.49, Windows waits longer for the MySQL server shutdown to complete. If you notice this still is not enough for your installation, it is safest not to run the MySQL server as a service. Instead, start it from the command-line prompt, and stop it with `mysqladmin shutdown`.

This change to tell Windows to wait longer when stopping the MySQL server works for Windows 2000

and XP. It does not work for Windows NT, where Windows waits only 20 seconds for a service to shut down, and after that kills the service process. You can increase this default by opening the [Registry Editor](#) (`\winnt\system32\regedt32.exe`) and editing the value of `WaitToKillServiceTimeout` at `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control` in the Registry tree. Specify the new larger value in milliseconds. For example, the value `120000` tells Windows NT to wait up to 120 seconds.

If `mysqld` is not running as a service, you can start it from the command line. For instructions, see [Section 2.3.10, “Starting MySQL from the Windows Command Line”](#).

Please see [Section 2.3.13, “Troubleshooting a MySQL Installation Under Windows”](#), if you encounter difficulties during installation.

## 2.3.12. Testing The MySQL Installation

You can test whether the MySQL server is working by executing any of the following commands:

```
C:\> C:\mysql\bin\mysqlshow
C:\> C:\mysql\bin\mysqlshow -u root mysql
C:\> C:\mysql\bin\mysqladmin version status proc
C:\> C:\mysql\bin\mysql test
```

If `mysqld` is slow to respond to TCP/IP connections from client programs on Windows 9x/Me, there is probably a problem with your DNS. In this case, start `mysqld` with the `--skip-name-resolve` option and use only `localhost` and IP numbers in the `Host` column of the MySQL grant tables.

You can force a MySQL client to use a named-pipe connection rather than TCP/IP by specifying the `-pipe` option or by specifying `.` (period) as the host name. Use the `--socket` option to specify the name of the pipe if you do not want to use the default pipe name. As of MySQL 4.1, you can use the `-protocol=PIPE` option instead.

Note that if you have set a password for the `root` account, deleted the anonymous account, or created a new user account, then you must use the appropriate `-u` and `-p` options with the commands shown above in order to connect with the MySQL Server. See [Section 5.7.4, “Connecting to the MySQL Server”](#).

There are two versions of the MySQL command-line tool on Windows:

Binary	Description
<code>mysql</code>	Compiled on native Windows, offering limited text editing capabilities.
<code>mysqlc</code>	Compiled with the Cygnus GNU compiler and libraries, which offers <code>readline</code> editing. <code>mysqlc</code> was intended for use primarily with Windows 9x/Me. It does not support the updated authentication protocol used beginning with MySQL 4.1, and is not supported in MySQL 4.1 and above. Beginning with MySQL 4.1.8, it is no longer included in MySQL Windows distributions.

To use `mysqlc`, you must have a copy of the `cygwinb19.dll` library installed somewhere that `mysqlc` can find it. If your distribution does not have the `cygwinb19.dll` library in the `bin` directory under the base directory of your MySQL installation, look for it in the `lib` directory and copy it to your Windows system directory (`\Windows\system` or a similar place).

For more information about `mysqlshow`, see [Section 8.15, “mysqlshow — Display Database, Table, and Column Information”](#).

## 2.3.13. Troubleshooting a MySQL Installation Under Windows

When installing and running MySQL for the first time, you may encounter certain errors that prevent the MySQL server from starting. The purpose of this section is to help you diagnose and correct some of these errors.

Your first resource when troubleshooting server issues is the error log. The MySQL server uses the error log to record information relevant to the error that prevents the server from starting. The error log is located in the data directory specified in your `my.ini` file. The default data directory location is `C:\mysql\data`. See [Section 5.11.1, “The Error Log”](#).

Another source of information regarding possible errors is the console messages displayed when the MySQL service is starting. Use the `NET START MySQL` command from the command line after installing `mysqld` as a service to see any error messages regarding the starting of the MySQL server as a service. See [Section 2.3.11, “Starting MySQL as a Windows Service”](#).

The following examples show other common error messages you may encounter when installing MySQL and starting the server for the first time:

- If the MySQL server cannot find the `mysql` privileges database or other critical files, you may see these messages:

```
System error 1067 has occurred.
Fatal error: Can't open privilege tables: Table 'mysql.host' does not exist
```

These messages often occur when the MySQL base or data directories are installed in different locations than the default locations (`C:\mysql` and `C:\mysql\data`, respectively).

This situation may occur when MySQL is upgraded and installed to a new location, but the configuration file is not updated to reflect the new location. In addition, there may be old and new configuration files that conflict. Be sure to delete or rename any old configuration files when upgrading MySQL.

If you have installed MySQL to a directory other than `C:\mysql`, you need to ensure that the MySQL server is aware of this through the use of a configuration (`my.ini`) file. The `my.ini` file needs to be located in your Windows directory, typically `C:\WINDOWS` or `C:\WINNT`. You can determine its exact location from the value of the `WINDIR` environment variable by issuing the following command from the command prompt:

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

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

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

Note that Windows pathnames are specified in option files using forward slashes rather than backslashes. If you do use backslashes, you must double them:

```
[mysqld]
# set basedir to your installation path
basedir=C:\\Program Files\\mysql
# set datadir to the location of your data directory
datadir=D:\\MySQLdata
```

If you change the `datadir` value in your MySQL configuration file, you must move the contents of the existing MySQL data directory before restarting the MySQL server.

See [Section 2.3.7, “Creating an Option File”](#).

- If you reinstall or upgrade MySQL without first stopping and removing the existing MySQL service and install MySQL using the MySQL Configuration Wizard, you may see this error:

```
Error: Cannot create Windows service for MySql. Error: 0
```

This occurs when the Configuration Wizard tries to install the service and finds an existing service with the same name.

One solution to this problem is to choose a service name other than `mysql` when using the configuration wizard. This allows the new service to be installed correctly, but leaves the outdated service in place. Although this is harmless, it is best to remove old services that are no longer in use.

To permanently remove the old `mysql` service, execute the following command as a user with administrative privileges, on the command-line:

```
C:\> sc delete mysql
[SC] DeleteService SUCCESS
```

If the `sc` utility is not available for your version of Windows, download the `delsrv` utility from <http://www.microsoft.com/windows2000/techinfo/reskit/tools/existing/delsrv-o.asp> and use the `delsrv mysql` syntax.

## 2.3.14. Upgrading MySQL on Windows

This section lists some of the steps you should take when upgrading MySQL on Windows.

1. Review [Section 2.11, “Upgrading MySQL”](#), for additional information on upgrading MySQL that is not specific to Windows.
2. You should always back up your current MySQL installation before performing an upgrade. See [Section 5.9.1, “Database Backups”](#).
3. Download the latest Windows distribution of MySQL from <http://dev.mysql.com>.
4. Before upgrading MySQL, you must stop the server. If the server is installed as a service, stop the service with the following command from the command prompt:

```
C:\> NET STOP MySQL
```

If you are not running the MySQL server as a service, use the following command to stop it:

```
C:\> C:\mysql\bin\mysqladmin -u root shutdown
```

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

5. When upgrading to MySQL 4.1.5 or higher from a previous version, or when upgrading from a version of MySQL installed from a Zip archive to a version of MySQL installed with the MySQL Installation Wizard, you must manually remove the previous installation and MySQL service (if the

server is installed as a service).

To remove the MySQL service, use the following command:

```
C:\> C:\mysql\bin\mysqld --remove
```

**If you do not remove the existing service, the MySQL Installation Wizard may fail to properly install the new MySQL service.**

6. If you are using the MySQL Installation Wizard, start the wizard as described in [Section 2.3.3, “Using the MySQL Installation Wizard”](#).
7. If you are installing MySQL from a Zip archive, extract the archive. You may either overwrite your existing MySQL installation (usually located at `C:\mysql`), or install it into a different directory, such as `C:\mysql5`. Overwriting the existing installation is recommended.
8. If you were running MySQL as a Windows service and you had to remove the service earlier in this procedure, reinstall the service. (See [Section 2.3.11, “Starting MySQL as a Windows Service”](#).)
9. Restart the server. For example, use `NET START MySQL` if you run MySQL as a service, or invoke `mysqld` directly otherwise.
10. If you encounter errors, see [Section 2.3.13, “Troubleshooting a MySQL Installation Under Windows”](#).

## 2.3.15. MySQL on Windows Compared to MySQL on Unix

MySQL for Windows has proven itself to be very stable. The Windows version of MySQL has the same features as the corresponding Unix version, with the following exceptions:

- **Windows 95 and threads**

Windows 95 leaks about 200 bytes of main memory for each thread creation. Each connection in MySQL creates a new thread, so you should not run `mysqld` for an extended time on Windows 95 if your server handles many connections! Newer versions of Windows don't suffer from this bug.

- **Limited number of ports**

Windows systems have about 4,000 ports available for client connections, and after a connection on a port closes, it takes two to four minutes before the port can be reused. In situations where clients connect to and disconnect from the server at a high rate, it is possible for all available ports to be used up before closed ports become available again. If this happens, the MySQL server appears to be unresponsive even though it is running. Note that ports may be used by other applications running on the machine as well, in which case the number of ports available to MySQL is lower.

For more information about this problem, see <http://support.microsoft.com/default.aspx?scid=kb;en-us;196271>.

- **Concurrent reads**

MySQL depends on the `pread()` and `pwrite()` calls to be able to mix `INSERT` and `SELECT`. Currently, we use mutexes to emulate `pread()/pwrite()`. We will, in the long run, replace the file level interface with a virtual interface so that we can use the `readfile()/writefile()` interface on NT, 2000, and XP to get more speed. The current implementation limits the number of open files that MySQL can use to 2,048 (1,024 before MySQL 4.0.19), which means that you cannot run as many concurrent threads on NT, 2000, XP, and 2003 as on Unix.

- **Blocking read**

MySQL uses a blocking read for each connection. That has the following implications if named-pipe connections are enabled:

- A connection is not disconnected automatically after eight hours, as happens with the Unix version of MySQL.
- If a connection hangs, it is impossible to break it without killing MySQL.
- `mysqladmin kill` does not work on a sleeping connection.
- `mysqladmin shutdown` cannot abort as long as there are sleeping connections.

We plan to fix this problem when our Windows developers have figured out a workaround.

- **ALTER TABLE**

While you are executing an `ALTER TABLE` statement, the table is locked from being used by other threads. This has to do with the fact that on Windows, you cannot delete a file that is in use by another thread. In the future, we may find some way to work around this problem.

- **DROP TABLE**

`DROP TABLE` on a table that is in use by a `MERGE` table does not work on Windows because the `MERGE` handler does the table mapping hidden from the upper layer of MySQL. Because Windows does not allow you to drop files that are open, you first must flush all `MERGE` tables (with `FLUSH TABLES`) or drop the `MERGE` table before dropping the table. We will fix this at the same time we introduce views.

- **DATA DIRECTORY and INDEX DIRECTORY**

The `DATA DIRECTORY` and `INDEX DIRECTORY` options for `CREATE TABLE` are ignored on Windows, because Windows does not support symbolic links. These options also are ignored on systems that have a non-functional `realpath()` call.

- **DROP DATABASE**

You cannot drop a database that is in use by a thread.

- **Killing MySQL from the Task Manager**

On Windows 95, you cannot kill MySQL from the Task Manager or with the shutdown utility. You must stop it with `mysqladmin shutdown` or the `NET STOP ...` command.

- **Case-insensitive names**

Filenames are not case sensitive on Windows, so MySQL database and table names are also not case sensitive on Windows. The only restriction is that database and table names must be specified using the same case throughout a given statement. See [Section 9.2.2, “Identifier Case Sensitivity”](#).

- **The ‘\’ pathname separator character**

Pathname components in Windows are separated by the ‘\’ character, which is also the escape character in MySQL. If you are using `LOAD DATA INFILE` or `SELECT ... INTO OUTFILE`, use Unix-style filenames with ‘/’ characters:

```
mysql> LOAD DATA INFILE 'C:/tmp/skr.txt' INTO TABLE skr;  
mysql> SELECT * INTO OUTFILE 'C:/tmp/skr.txt' FROM skr;
```



Alternatively, you must double the ‘\’ character:

```
mysql> LOAD DATA INFILE 'C:\\tmp\\skr.txt' INTO TABLE skr;  
mysql> SELECT * INTO OUTFILE 'C:\\tmp\\skr.txt' FROM skr;
```

- **Problems with pipes**

Pipes do not work reliably from the Windows command-line prompt. If the pipe includes the character `^Z / CHAR(24)`, Windows thinks that it has encountered end-of-file and aborts the program.

This is a problem mainly when you try to apply a binary log as follows:

```
C:\> mysqlbinlog binary_log_file | mysql --user=root
```

If you have a problem applying the log and suspect that it is because of a `^Z / CHAR(24)` character, you can use the following workaround:

```
C:\> mysqlbinlog binary_log_file --result-file=/tmp/bin.sql  
C:\> mysql --user=root --execute "source /tmp/bin.sql"
```

The latter command also can be used to reliably read in any SQL file that may contain binary data.

- **Access denied for user error**

If MySQL cannot resolve your hostname properly, you may get the following error when you attempt to run a MySQL client program to connect to a server running on the same machine:

```
Access denied for user 'some_user'@'unknown'  
to database 'mysql'
```

To fix this problem, you should create a file named `\windows\hosts` containing the following information:

```
127.0.0.1 localhost
```

Here are some open issues for anyone who might want to help us improve MySQL on Windows:

- Add macros to use the faster thread-safe increment/decrement methods provided by Windows.

## 2.4. Installing MySQL on Linux

The recommended way to install MySQL on Linux is by using the RPM packages. The MySQL RPMs are currently built on a SuSE Linux 7.3 system, but should work on most versions of Linux that support `rpm` and use `glibc`. To obtain RPM packages, see [Section 2.1.3, “How to Get MySQL”](#).

MySQL AB does provide some platform-specific RPMs; the difference between a platform-specific RPM and a generic RPM is that a platform-specific RPM is built on the targeted platform and is linked dynamically whereas a generic RPM is linked statically with LinuxThreads.

**Note:** RPM distributions of MySQL often are provided by other vendors. Be aware that they may differ in features and capabilities from those built by MySQL AB, and that the instructions in this manual do not necessarily apply to installing them. The vendor's instructions should be consulted instead.

If you have problems with an RPM file (for example, if you receive the error `Sorry, the host 'xxxx' could not be looked up`), see [Section 2.13.1.2, “Linux Binary Distribution Notes”](#).

In most cases, you need to install only the `MySQL-server` and `MySQL-client` packages to get a functional MySQL installation. The other packages are not required for a standard installation. If you want to run a MySQL-Max server that has additional capabilities, you should also install the `MySQL-Max` RPM. However, you should do so only *after* installing the `MySQL-server` RPM. See [Section 5.3, “The `mysqld-max` Extended MySQL Server”](#).

If you get a dependency failure when trying to install the MySQL 4.0 packages (for example, `error: removing these packages would break dependencies: libmysqlclient.so.10 is needed by ...`), you should also install the `MySQL-shared-compat` package, which includes both the shared libraries for backward compatibility (`libmysqlclient.so.12` for MySQL 4.0 and `libmysqlclient.so.10` for MySQL 3.23).

Some Linux distributions still ship with MySQL 3.23 and they usually link applications dynamically to save disk space. If these shared libraries are in a separate package (for example, `MySQL-shared`), it is sufficient to simply leave this package installed and just upgrade the MySQL server and client packages (which are statically linked and do not depend on the shared libraries). For distributions that include the shared libraries in the same package as the MySQL server (for example, Red Hat Linux), you could either install our 3.23 `MySQL-shared` RPM, or use the `MySQL-shared-compat` package instead. (Do not install both.)

The following RPM packages are available:

- `MySQL-server-VERSION.i386.rpm`

The MySQL server. You need this unless you only want to connect to a MySQL server running on another machine.

Note: Server RPM files were called `MySQL-VERSION.i386.rpm` before MySQL 4.0.10. That is, they did not have `-server` in the name.

- `MySQL-Max-VERSION.i386.rpm`

The MySQL-Max server. This server has additional capabilities that the one provided in the `MySQL-server` RPM does not. You must install the `MySQL-server` RPM first, because the `MySQL-Max` RPM depends on it.

- `MySQL-client-VERSION.i386.rpm`

The standard MySQL client programs. You probably always want to install this package.

- `MySQL-bench-VERSION.i386.rpm`

Tests and benchmarks. Requires Perl and the `DBI` and `DBD:mysql` modules.

- `MySQL-devel-VERSION.i386.rpm`

The libraries and include files that are needed if you want to compile other MySQL clients, such as the Perl modules.

- `MySQL-shared-VERSION.i386.rpm`

This package contains the shared libraries (`libmysqlclient.so*`) that certain languages and applications need to dynamically load and use MySQL. It contains single-threaded and thread-safe libraries. If you install this package, do not install the `MySQL-shared-compat` package.

- `MySQL-shared-compat-VERSION.i386.rpm`

This package includes the shared libraries for MySQL 3.23, 4.0, and 4.1. It contains single-threaded and thread-safe libraries. Install this package instead of `MySQL-shared` if you have applications installed that are dynamically linked against older versions of MySQL but you want to upgrade to the current version without breaking the library dependencies. This package has been available since MySQL 4.0.13.

- `MySQL-embedded-VERSION.i386.rpm`

The embedded MySQL server library (available as of MySQL 4.0).

- `MySQL-VERSION.src.rpm`

This contains the source code for all of the previous packages. It can also be used to rebuild the RPMs on other architectures (for example, Alpha or SPARC).

To see all files in an RPM package (for example, a `MySQL-server` RPM), run a command like this:

```
shell> rpm -qpl MySQL-server-VERSION.i386.rpm
```

To perform a standard minimal installation, install the server and client RPMs:

```
shell> rpm -i MySQL-server-VERSION.i386.rpm
shell> rpm -i MySQL-client-VERSION.i386.rpm
```

To install only the client programs, install just the client RPM:

```
shell> rpm -i MySQL-client-VERSION.i386.rpm
```

RPM provides a feature to verify the integrity and authenticity of packages before installing them. If you would like to learn more about this feature, see [Section 2.1.4, “Verifying Package Integrity Using MD5 Checksums or GnuPG”](#).

The server RPM places data under the `/var/lib/mysql` directory. The RPM also creates a login account for a user named `mysql` (if one does not exist) to use for running the MySQL server, and creates the appropriate entries in `/etc/init.d/` to start the server automatically at boot time. (This means that if you have performed a previous installation and have made changes to its startup script, you may want to make a copy of the script so that you do not lose it when you install a newer RPM.) See [Section 2.10.2.2, “Starting and Stopping MySQL Automatically”](#), for more information on how MySQL can be started automatically on system startup.

If you want to install the MySQL RPM on older Linux distributions that do not support initialization scripts in `/etc/init.d` (directly or via a symlink), you should create a symbolic link that points to the location where your initialization scripts actually are installed. For example, if that location is `/etc/rc.d/init.d`, use these commands before installing the RPM to create `/etc/init.d` as a symbolic link that points there:

```
shell> cd /etc
shell> ln -s rc.d/init.d .
```

However, all current major Linux distributions should support the new directory layout that uses `/etc/init.d`, because it is required for LSB (Linux Standard Base) compliance.

If the RPM files that you install include `MySQL-server`, the `mysqld` server should be up and running after installation. You should be able to start using MySQL.

If something goes wrong, you can find more information in the binary installation section. See [Section 2.8, “Installing MySQL on Other Unix-Like Systems”](#).

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

## 2.5. Installing MySQL on Mac OS X

Beginning with MySQL 4.0.11, you can install MySQL on Mac OS X 10.3.x (“Panther”) or newer using a Mac OS X binary package in PKG format instead of the binary tarball distribution. Please note that older versions of Mac OS X (for example, 10.1.x or 10.2.x) are not supported by this package.

The package is located inside a disk image (`.dmg`) file that you first need to mount by double-clicking its icon in the Finder. It should then mount the image and display its contents.

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

**Note:** Before proceeding with the installation, be sure to shut down all running MySQL server instances by either using the MySQL Manager Application (on Mac OS X Server) or via `mysqladmin shutdown` on the command line.

To install the MySQL PKG file, double-click on the package icon. This launches the Mac OS X Package Installer, which guides you through the installation of MySQL.

Due to a bug in the Mac OS X package installer, you may see this error message in the destination disk selection dialog:

```
You cannot install this software on this disk. (null)
```

If this error occurs, simply click the [Go Back](#) button once to return to the previous screen. Then click [Continue](#) to advance to the destination disk selection again, and you should be able to choose the destination disk correctly. We have reported this bug to Apple and it is investigating this problem.

The Mac OS X PKG of MySQL installs itself into `/usr/local/mysql-VERSION` and also installs a symbolic link, `/usr/local/mysql`, that points to the new location. If a directory named `/usr/local/mysql` exists, it is renamed to `/usr/local/mysql.bak` first. In addition, the installer creates the grant tables in the `mysql` database by executing `mysql_install_db`.

The installation layout is similar to that of a `tar` file binary distribution; all MySQL binaries are located in the directory `/usr/local/mysql/bin`. The MySQL socket file is created as `/tmp/mysql.sock` by default. See [Section 2.1.5, “Installation Layouts”](#).

MySQL installation requires a Mac OS X user account named `mysql`. A user account with this name should exist by default on Mac OS X 10.2 and up.

If you are running Mac OS X Server, a version of MySQL should already be installed. The following table shows the versions of MySQL that ship with Mac OS X Server versions.

Mac OS X Server Version	MySQL Version
10.2-10.2.2	3.23.51
10.2.3-10.2.6	3.23.53
10.3	4.0.14
10.3.2	4.0.16
10.4.0	4.1.10a

This manual section covers the installation of the official MySQL Mac OS X PKG only. Make sure to read Apple's help information about installing MySQL: Run the [Help View](#) application, select Mac OS X Server help, search for “MySQL”, and read the item entitled “Installing MySQL”.

For pre-installed versions of MySQL on Mac OS X Server, note especially that you should start `mysqld` with `safe_mysqld` instead of `mysqld_safe` if MySQL is older than version 4.0.

If you previously used Marc Liyanage's MySQL packages for Mac OS X from <http://www.entropy.ch>, you can simply follow the update instructions for packages using the binary installation layout as given on his pages.

If you are upgrading from Marc's 3.23.x versions or from the Mac OS X Server version of MySQL to the official MySQL PKG, you also need to convert the existing MySQL privilege tables to the current format, because some new security privileges have been added. See [Section 5.5.1](#), “`mysql_fix_privilege_tables` — Upgrade MySQL System Tables”.

If you want MySQL to start automatically during system startup, you also need to install the MySQL Startup Item. Starting with MySQL 4.0.15, it is part of the Mac OS X installation disk images as a separate installation package. Simply double-click the `MySQLStartupItem.pkg` icon and follow the instructions to install it. The Startup Item need be installed only once. There is no need to install it each time you upgrade the MySQL package later.

The Startup Item for MySQL is installed into `/Library/StartupItems/MySQLCOM`. (Before MySQL 4.1.2, the location was `/Library/StartupItems/MySQL`, but that collided with the MySQL Startup Item installed by Mac OS X Server.) Startup Item installation adds a variable `MYSQLCOM=-YES-` to the system configuration file `/etc/hostconfig`. If you want to disable the automatic startup of MySQL, simply change this variable to `MYSQLCOM=-NO-`.

On Mac OS X Server, the default MySQL installation uses the variable `MYSQL` in the `/etc/hostconfig` file. The MySQL AB Startup Item installer disables this variable by setting it to `MYSQL=-NO-`. This avoids boot time conflicts with the `MYSQLCOM` variable used by the MySQL AB Startup Item. However, it does not shut down a running MySQL server. You should do that yourself.

After the installation, you can start up MySQL by running the following commands in a terminal window. You must have administrator privileges to perform this task.

If you have installed the Startup Item, use this command:

```
shell> sudo /Library/StartupItems/MySQLCOM/MySQLCOM start
(Enter your password, if necessary)
(Press Control-D or enter "exit" to exit the shell)
```

For versions of MySQL older than 4.1.3, substitute `/Library/StartupItems/MySQLCOM/MySQLCOM` with `/Library/StartupItems/MySQL/MySQL` above.

If you do not use the Startup Item, enter the following command sequence:

```
shell> cd /usr/local/mysql
shell> sudo ./bin/mysqld_safe
(Enter your password, if necessary)
(Press Control-Z)
shell> bg
(Press Control-D or enter "exit" to exit the shell)
```

You should be able to connect to the MySQL server, for example, by running `/usr/local/mysql/bin/mysql`.

**Note:** The accounts that are listed in the MySQL grant tables initially have no passwords. After starting the server, you should set up passwords for them using the instructions in [Section 2.10](#).

### “Post-Installation Setup and Testing”.

You might want to add aliases to your shell's resource file to make it easier to access commonly used programs such as `mysql` and `mysqladmin` from the command line. The syntax for `bash` is:

```
alias mysql=/usr/local/mysql/bin/mysql
alias mysqladmin=/usr/local/mysql/bin/mysqladmin
```

For `tcsh`, use:

```
alias mysql /usr/local/mysql/bin/mysql
alias mysqladmin /usr/local/mysql/bin/mysqladmin
```

Even better, add `/usr/local/mysql/bin` to your `PATH` environment variable. For example, add the following line to your `$HOME/.bashrc` file if your shell is `bash`:

```
PATH=${PATH}:/usr/local/mysql/bin
```

Add the following line to your `$HOME/.tcshrc` file if your shell is `tcsh`:

```
setenv PATH ${PATH}:/usr/local/mysql/bin
```

If no `.bashrc` or `.tcshrc` file exists in your home directory, create it with a text editor.

If you are upgrading an existing installation, note that installing a new MySQL PKG does not remove the directory of an older installation. Unfortunately, the Mac OS X Installer does not yet offer the functionality required to properly upgrade previously installed packages.

To use your existing databases with the new installation, you will need to copy the contents of the old data directory to the new data directory. Make sure that neither the old server nor the new one is running when you do this. After you have copied over the MySQL database files from the previous installation and have successfully started the new server, you should consider removing the old installation files to save disk space. Additionally, you should also remove older versions of the Package Receipt directories located in `/Library/Receipts/mysql-VERSION.pkg`.

## 2.6. Installing MySQL on Solaris

If you install MySQL using a binary tarball distribution on Solaris, you may run into trouble even before you get the MySQL distribution unpacked, as the Solaris `tar` cannot handle long filenames. This means that you may see errors when you try to unpack MySQL.

If this occurs, you must use GNU `tar` (`gtar`) to unpack the distribution. You can find a precompiled copy for Solaris at <http://dev.mysql.com/downloads/os-solaris.html>.

You can install MySQL on Solaris using a binary package in PKG format instead of the binary tarball distribution. Some basic PKG-handling commands follow:

- To add a package:

```
pkgadd -d package_name.pkg
```

- To remove a package:

```
pkgrm package_name
```

- To get a full list of installed packages:

```
pkginfo
```

- To get detailed information for a package:

```
pkginfo -l package_name
```

- To list the files belonging to a package:

```
pkgchk -v package_name
```

- To get packaging information for an arbitrary file:

```
pkgchk -l -p file_name
```

For additional information about installing MySQL on Solaris, see [Section 2.13.3, “Solaris Notes”](#).

## 2.7. Installing MySQL on NetWare

Porting MySQL to NetWare was an effort spearheaded by Novell. Novell customers should be pleased to note that NetWare 6.5 ships with bundled MySQL binaries, complete with an automatic commercial use license for all servers running that version of NetWare.

MySQL for NetWare is compiled using a combination of [Metrowerks CodeWarrior for NetWare](#) and special cross-compilation versions of the GNU autotools.

The latest binary packages for NetWare can be obtained at <http://dev.mysql.com/downloads/>. See [Section 2.1.3, “How to Get MySQL”](#).

To host MySQL, the NetWare server must meet these requirements:

- The latest Support Pack of [NetWare 6.5](#) [<http://support.novell.com/filefinder/18197/index.html>] must be installed.
- The system must meet Novell's minimum requirements to run the respective version of NetWare.
- MySQL data and the program binaries must be installed on an NSS volume; traditional volumes are not supported.

To install MySQL for NetWare, use the following procedure:

1. If you are upgrading from a prior installation, stop the MySQL server. This is done from the server console, using the following command:

```
SERVER: mysqladmin -u root shutdown
```

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

2. Log on to the target server from a client machine with access to the location where you are installing MySQL.
3. Extract the binary package Zip file onto the server. Be sure to allow the paths in the Zip file to be

used. It is safe to simply extract the file to `SYS:\`.

If you are upgrading from a prior installation, you may need to copy the data directory (for example, `SYS:MYSQL\DATA`), as well as `my.cnf`, if you have customized it. You can then delete the old copy of MySQL.

4. You might want to rename the directory to something more consistent and easy to use. The examples in this manual use `SYS:MYSQL` to refer to the installation directory.

Note that MySQL installation on NetWare does not detect if a version of MySQL is already installed outside the NetWare release. Therefore, if you have installed the latest MySQL version from the Web (for example, MySQL 4.1 or later) in `SYS:MYSQL`, you must rename the folder before upgrading the NetWare server; otherwise, files in `SYS:MYSQL` are overwritten by the MySQL version present in NetWare Support Pack.

5. At the server console, add a search path for the directory containing the MySQL NLMs. For example:

```
SERVER: SEARCH ADD SYS:MYSQL\BIN
```

6. Initialize the data directory and the grant tables, if necessary, by executing `mysql_install_db` at the server console.
7. Start the MySQL server using `mysqld_safe` at the server console.
8. To finish the installation, you should also add the following commands to `autoexec.ncf`. For example, if your MySQL installation is in `SYS:MYSQL` and you want MySQL to start automatically, you could add these lines:

```
#Starts the MySQL 4.0.x database server
SEARCH ADD SYS:MYSQL\BIN
MYSQLD_SAFE
```

If you are running MySQL on NetWare 6.0, we strongly suggest that you use the `-skip-external-locking` option on the command line:

```
#Starts the MySQL 4.0.x database server
SEARCH ADD SYS:MYSQL\BIN
MYSQLD_SAFE --skip-external-locking
```

It is also necessary to use `CHECK TABLE` and `REPAIR TABLE` instead of `myisamchk`, because `myisamchk` makes use of external locking. External locking is known to have problems on NetWare 6.0; the problem has been eliminated in NetWare 6.5. Note that the use of MySQL on NetWare 6.0 is not officially supported.

`mysqld_safe` on NetWare provides a screen presence. When you unload (shut down) the `mysqld_safe` NLM, the screen does not go away by default. Instead, it prompts for user input:

```
*<NLM has terminated; Press any key to close the screen>*
```

If you want NetWare to close the screen automatically instead, use the `--autoclose` option to `mysqld_safe`. For example:

```
#Starts the MySQL 4.0.x database server
SEARCH ADD SYS:MYSQL\BIN
MYSQLD_SAFE --autoclose
```

The behavior of `mysqld_safe` on NetWare is described further in [Section 5.4.1](#),



`mysqld_safe` — MySQL Server Startup Script”.

9. When installing MySQL version 4.1.x or later, either for the first time or upgrading the 4.0.x version to 4.1.x or later, download and install the latest and appropriate Perl module and PHP extension for NetWare:
  - Perl: <http://forge.novell.com/modules/xfcontent/downloads.php/perl/Modules/>
  - PHP: <http://forge.novell.com/modules/xfcontent/downloads.php/php/Modules/>

If there was an existing installation of MySQL on the NetWare server, be sure to check for existing MySQL startup commands in `autoexec.ncf`, and edit or delete them as necessary.

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

## 2.8. Installing MySQL on Other Unix-Like Systems

This section covers the installation of MySQL binary distributions that are provided for various platforms in the form of compressed `tar` files (files with a `.tar.gz` extension). See [Section 2.1.2.5](#), “MySQL Binaries Compiled by MySQL AB”, for a detailed list.

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

MySQL `tar` file binary distributions have names of the form `mysql-VERSION-OS.tar.gz`, where `VERSION` is a number (for example, `4.0.17`), and `OS` indicates the type of operating system for which the distribution is intended (for example, `pc-linux-i686`).

In addition to these generic packages, we also offer binaries in platform-specific package formats for selected platforms. See [Section 2.2](#), “Standard MySQL Installation Using a Binary Distribution”, for more information on how to install these.

You need the following tools to install a MySQL `tar` file binary distribution:

- GNU `gunzip` to uncompress the distribution.
- A reasonable `tar` to unpack the distribution. GNU `tar` is known to work. Some operating systems come with a pre-installed version of `tar` that is known to have problems. For example, Mac OS X `tar` and Sun `tar` are known to have problems with long filenames. On Mac OS X, you can use the pre-installed `gnutar` program. On other systems with a deficient `tar`, you should install GNU `tar` first.

If you run into problems and need to file a bug report, please use the instructions in [Section 1.8](#), “How to Report Bugs or Problems”.

The basic commands that you must execute to install and use a MySQL binary distribution are:

```
shell> groupadd mysql
shell> useradd -g mysql mysql
shell> cd /usr/local
shell> gunzip < /path/to/mysql-VERSION-OS.tar.gz | tar xvf -
shell> ln -s full-path-to-mysql-VERSION-OS mysql
shell> cd mysql
shell> scripts/mysql_install_db --user=mysql
shell> chown -R root .
shell> chown -R mysql data
shell> chgrp -R mysql .
shell> bin/mysqld_safe --user=mysql &
```

For versions of MySQL older than 4.0, substitute `bin/safe_mysqld` for `bin/mysqld_safe` in the final command.

**Note:** This procedure does not set up any passwords for MySQL accounts. After following the procedure, proceed to [Section 2.10, “Post-Installation Setup and Testing”](#).

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

1. Add a login user and group for `mysqld` to run as:

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

These commands add the `mysql` group and the `mysql` user. The syntax for `useradd` and `groupadd` may differ slightly on different versions of Unix, or they may have different names such as `adduser` and `addgroup`.

You might want to call the user and group something else instead of `mysql`. If so, substitute the appropriate name in the following steps.

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

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

3. Obtain a distribution file using the instructions in [Section 2.1.3, “How to Get MySQL”](#). For a given release, binary distributions for all platforms are built from the same MySQL source distribution.
4. Unpack the distribution, which creates the installation directory. Then create a symbolic link to that directory:

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

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

With GNU `tar`, no separate invocation of `gunzip` is necessary. You can replace the first line with the following alternative command to uncompress and extract the distribution:

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

5. Change location into the installation directory:

```
shell> cd mysql
```

You will find several files and subdirectories in the `mysql` directory. The most important for installation purposes are the `bin` and `scripts` subdirectories:

- The `bin` directory contains client programs and the server. You should add the full pathname of this directory to your `PATH` environment variable so that your shell finds the MySQL programs properly. See [Appendix F, \*Environment Variables\*](#).

- The `scripts` directory contains the `mysql_install_db` script used to initialize the `mysql` database containing the grant tables that store the server access permissions.
6. If you have not installed MySQL before, you must create the MySQL grant tables:

```
shell> scripts/mysql_install_db --user=mysql
```

If you run the command as `root`, you must use the `--user` option as shown. The value of the option should be the name of the login account that you created in the first step to use for running the server. If you run the command while logged in as that user, you can omit the `--user` option.

Note that for MySQL versions older than 3.22.10, `mysql_install_db` left the server running after creating the grant tables. This is no longer true; you need to restart the server after performing the remaining steps in this procedure.

7. Change the ownership of program binaries to `root` and ownership of the data directory to the user that you run `mysqld` as. Assuming that you are located in the installation directory (`/usr/local/mysql`), the commands look like this:

```
shell> chown -R root .
shell> chown -R mysql data
shell> chgrp -R mysql .
```

The first command changes the owner attribute of the files to the `root` user. The second changes the owner attribute of the data directory to the `mysql` user. The third changes the group attribute to the `mysql` group.

8. If you want MySQL to start automatically when you boot your machine, you can copy `support-files/mysql.server` to the location where your system has its startup files. More information can be found in the `mysql.server` script itself, and in [Section 2.10.2.2, “Starting and Stopping MySQL Automatically”](#).
9. You can set up new accounts using the `bin/mysql_setpermission` script if you install the `DBI` and `DBD: :mysql` Perl modules. For instructions, see [Section 2.14, “Perl Installation Notes”](#).
10. If you would like to use `mysqlaccess` and have the MySQL distribution in some non-standard location, you must change the location where `mysqlaccess` expects to find the `mysql` client. Edit the `bin/mysqlaccess` script at approximately line 18. Search for a line that looks like this:

```
$MYSQL = '/usr/local/bin/mysql'; # path to mysql executable
```

Change the path to reflect the location where `mysql` actually is stored on your system. If you do not do this, a `Broken pipe` error will occur when you run `mysqlaccess`.

After everything has been unpacked and installed, you should test your distribution. To start the MySQL server, use the following command:

```
shell> bin/mysqld_safe --user=mysql &
```

If that command fails immediately and prints `mysqld ended`, you can find some information in the `host_name.err` file in the data directory.

For versions of MySQL older than 4.0, substitute `bin/safe_mysqld` for `bin/mysqld_safe` in the command.

More information about `mysqld_safe` is given in [Section 5.4.1, “mysqld\\_safe — MySQL Server](#)

Startup Script”.

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

## 2.9. MySQL Installation Using a Source Distribution

Before you proceed with an installation from source, first check whether our binary is available for your platform and whether it works for you. We put a great deal of effort into ensuring that our binaries are built with the best possible options.

To obtain a source distribution for MySQL, [Section 2.1.3](#), “How to Get MySQL”.

MySQL source distributions are provided as compressed `tar` archives and have names of the form `mysql-VERSION.tar.gz`, where `VERSION` is a number like `4.1.13`.

You need the following tools to build and install MySQL from source:

- GNU `gunzip` to uncompress the distribution.
- A reasonable `tar` to unpack the distribution. GNU `tar` is known to work. Some operating systems come with a pre-installed version of `tar` that is known to have problems. For example, the `tar` provided with early versions of Mac OS X `tar`, SunOS 4.x and Solaris 8 and earlier are known to have problems with long filenames. On Mac OS X, you can use the pre-installed `gnutar` program. On other systems with a deficient `tar`, you should install GNU `tar` first.
- A working ANSI C++ compiler. `gcc` 2.95.2 or later, `egcs` 1.0.2 or later or `egcs 2.91.66`, SGI C++, and SunPro C++ are some of the compilers that are known to work. `libg++` is not needed when using `gcc`. `gcc 2.7.x` has a bug that makes it impossible to compile some perfectly legal C++ files, such as `sql/sql_base.cc`. If you have only `gcc 2.7.x`, you must upgrade your `gcc` to be able to compile MySQL. `gcc 2.8.1` is also known to have problems on some platforms, so it should be avoided if a new compiler exists for the platform.

`gcc 2.95.2` or later is recommended when compiling MySQL 3.23.x.

- A good `make` program. GNU `make` is always recommended and is sometimes required. If you have problems, we recommend GNU `make 3.75` or newer.

If you are using a version of `gcc` recent enough to understand the `-fno-exceptions` option, it is *very important* that you use this option. Otherwise, you may compile a binary that crashes randomly. We also recommend that you also use `-felide-constructors` and `-fno-rtti` as well. When in doubt, do the following:

```
CFLAGS="-O3" CXX=gcc CXXFLAGS="-O3 -felide-constructors \
-fno-exceptions -fno-rtti" ./configure \
--prefix=/usr/local/mysql --enable-assembly \
--with-mysqld-ldflags=-all-static
```

On most systems, this gives you a fast and stable binary.

If you run into problems and need to file a bug report, please use the instructions in [Section 1.8](#), “How to Report Bugs or Problems”.

### 2.9.1. Source Installation Overview

The basic commands that you must execute to install a MySQL source distribution are:

```
shell> groupadd mysql
shell> useradd -g mysql mysql
shell> gunzip < mysql-VERSION.tar.gz | tar -xvf -
shell> cd mysql-VERSION
shell> ./configure --prefix=/usr/local/mysql
shell> make
shell> make install
shell> cp support-files/my-medium.cnf /etc/my.cnf
shell> cd /usr/local/mysql
shell> bin/mysql_install_db --user=mysql
shell> chown -R root .
shell> chown -R mysql var
shell> chgrp -R mysql .
shell> bin/mysqld_safe --user=mysql &
```

For versions of MySQL older than 4.0, substitute `bin/safe_mysqld` for `bin/mysqld_safe` in the final command.

If you start from a source RPM, do the following:

```
shell> rpmbuild --rebuild --clean MySQL-VERSION.src.rpm
```

This makes a binary RPM that you can install. For older versions of RPM, you may have to replace the command `rpmbuild` with `rpm` instead.

**Note:** This procedure does not set up any passwords for MySQL accounts. After following the procedure, proceed to [Section 2.10, “Post-Installation Setup and Testing”](#), for post-installation setup and testing.

A more detailed version of the preceding description for installing MySQL from a source distribution follows:

1. Add a login user and group for `mysqld` to run as:

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

These commands add the `mysql` group and the `mysql` user. The syntax for `useradd` and `groupadd` may differ slightly on different versions of Unix, or they may have different names such as `adduser` and `addgroup`.

You might want to call the user and group something other than `mysql`. If so, substitute the appropriate name in the following steps.

2. Pick the directory under which you want to unpack the distribution and change location into it.
3. Obtain a distribution file using the instructions in [Section 2.1.3, “How to Get MySQL”](#).
4. Unpack the distribution into the current directory:

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

This command creates a directory named `mysql-VERSION`.

With GNU `tar`, no separate invocation of `gunzip` is necessary. You can use the following alternative command to uncompress and extract the distribution:

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

5. Change location into the top-level directory of the unpacked distribution:

```
shell> cd mysql-VERSION
```

Note that currently you must configure and build MySQL from this top-level directory. You cannot build it in a different directory.

6. Configure the release and compile everything:

```
shell> ./configure --prefix=/usr/local/mysql
shell> make
```

When you run `configure`, you might want to specify other options. Run `./configure --help` for a list of options. [Section 2.9.2, “Typical configure Options”](#), discusses some of the more useful options.

If `configure` fails and you are going to send mail to a MySQL mailing list to ask for assistance, please include any lines from `config.log` that you think can help solve the problem. Also include the last couple of lines of output from `configure`. To file a bug report, please use the instructions in [Section 1.8, “How to Report Bugs or Problems”](#).

If the compile fails, see [Section 2.9.4, “Dealing with Problems Compiling MySQL”](#), for help.

7. Install the distribution:

```
shell> make install
```

If you want to set up an option file, use one of those present in the `support-files` directory as a template. For example:

```
shell> cp support-files/my-medium.cnf /etc/my.cnf
```

You might need to run these commands as `root`.

If you want to configure support for InnoDB tables, you should edit the `/etc/my.cnf` file, removing the `#` character before the option lines that start with `innodb_...`, and modify the option values to be what you want. See [Section 4.3.2, “Using Option Files”](#), and [Section 14.2.4, “InnoDB Configuration”](#).

8. Change location into the installation directory:

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

9. If you have not installed MySQL before, you must create the MySQL grant tables:

```
shell> bin/mysql_install_db --user=mysql
```

If you run the command as `root`, you should use the `--user` option as shown. The value of the option should be the name of the login account that you created in the first step to use for running the server. If you run the command while logged in as that user, you can omit the `--user` option.

Note that for MySQL versions older than 3.22.10, `mysql_install_db` left the server running after creating the grant tables. This is no longer true; you need to restart the server after performing the remaining steps in this procedure.

10. Change the ownership of program binaries to `root` and ownership of the data directory to the user that you run `mysqld` as. Assuming that you are in the installation directory (`/usr/local/mysql`), the commands look like this:

```
shell> chown -R root .
shell> chown -R mysql var
shell> chgrp -R mysql .
```

The first command changes the owner attribute of the files to the `root` user. The second changes the owner attribute of the data directory to the `mysql` user. The third changes the group attribute to the `mysql` group.

11. If you want MySQL to start automatically when you boot your machine, you can copy `support-files/mysql.server` to the location where your system has its startup files. More information can be found in the script itself, and in [Section 2.10.2.2, “Starting and Stopping MySQL Automatically”](#).
12. You can set up new accounts using the `bin/mysql_setpermission` script if you install the `DBI` and `DBD: :mysql` Perl modules. For instructions, see [Section 2.14, “Perl Installation Notes”](#).

After everything has been installed, you should test your distribution. To start the MySQL server, use the following command:

```
shell> /usr/local/mysql/bin/mysqld_safe --user=mysql &
```

For versions of MySQL older than 4.0, substitute `safe_mysqld` for `mysqld_safe` in the command.

If that command fails immediately and prints `mysqld ended`, you can find some information in the `host_name.err` file in the data directory.

More information about `mysqld_safe` is given in [Section 5.4.1, “mysqld\\_safe — MySQL Server Startup Script”](#).

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

## 2.9.2. Typical `configure` Options

The `configure` script gives you a great deal of control over how you configure a MySQL source distribution. Typically you do this using options on the `configure` command line. You can also affect `configure` using certain environment variables. See [Appendix F, \*Environment Variables\*](#). For a list of options supported by `configure`, run this command:

```
shell> ./configure --help
```

Some of the more commonly used `configure` options are described here:

- To compile just the MySQL client libraries and client programs and not the server, use the `--without-server` option:

```
shell> ./configure --without-server
```

If you have no C++ compiler, some client programs such as `mysql` cannot be compiled because they require C++. In this case, you can remove the code in `configure` that tests for the C++ compiler and then run `./configure` with the `--without-server` option. The compile step should still try to build all clients, but you can ignore any warnings about files such as `mysql.cc`. (If `make` stops, try `make -k` to tell it to continue with the rest of the build even if errors occur.)

- If you want to build the embedded MySQL library (`libmysqld.a`), use the `--with-embedded-server` option.
- If you do not want your log files and database directories located under `/usr/local/var`, use a `configure` command something like one of these:

```
shell> ./configure --prefix=/usr/local/mysql
shell> ./configure --prefix=/usr/local \
               --localstatedir=/usr/local/mysql/data
```

The first command changes the installation prefix so that everything is installed under `/usr/local/mysql` rather than the default of `/usr/local`. The second command preserves the default installation prefix, but overrides the default location for database directories (normally `/usr/local/var`) and changes it to `/usr/local/mysql/data`.

You can also specify the installation directory and data directory locations at server startup time by using the `--basedir` and `--datadir` options. These can be given on the command line or in an MySQL option file, although it is more common to use an option file. See [Section 4.3.2, “Using Option Files”](#).

- If you are using Unix and you want the MySQL socket file location to be somewhere other than the default location (normally in the directory `/tmp` or `/var/run`), use a `configure` command like this:

```
shell> ./configure \
               --with-unix-socket-path=/usr/local/mysql/tmp/mysql.sock
```

The socket filename must be an absolute pathname. You can also change the location of `mysql.sock` at server startup by using a MySQL option file. See [Section A.4.5, “How to Protect or Change the MySQL Unix Socket File”](#).

- If you want to compile statically linked programs (for example, to make a binary distribution, to get better performance, or to work around problems with some Red Hat Linux distributions), run `configure` like this:

```
shell> ./configure --with-client-ldflags=-all-static \
               --with-mysqld-ldflags=-all-static
```

- If you are using `gcc` and do not have `libg++` or `libstdc++` installed, you can tell `configure` to use `gcc` as your C++ compiler:

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

When you use `gcc` as your C++ compiler, it does not attempt to link in `libg++` or `libstdc++`. This may be a good thing to do even if you have those libraries installed. Some versions of them have caused strange problems for MySQL users in the past.

The following list indicates some compilers and environment variable settings that are commonly used with each one.

- `gcc` 2.7.2:

```
CC=gcc CXX=gcc CXXFLAGS="-O3 -felide-constructors"
```

- `egcs` 1.0.3a:

```
CC=gcc CXX=gcc CXXFLAGS="-O3 -felide-constructors \
-fno-exceptions -fno-rtti"
```



- `gcc` 2.95.2:

```
CFLAGS="-O3 -mpentiumpro" CXX=gcc CXXFLAGS="-O3 -mpentiumpro \  
-felide-constructors -fno-exceptions -fno-rtti"
```

- `pgcc` 2.90.29 or newer:

```
CFLAGS="-O3 -mpentiumpro -mstack-align-double" CXX=gcc \  
CXXFLAGS="-O3 -mpentiumpro -mstack-align-double \  
-felide-constructors -fno-exceptions -fno-rtti"
```

In most cases, you can get a reasonably optimized MySQL binary by using the options from the preceding list and adding the following options to the `configure` line:

```
--prefix=/usr/local/mysql --enable-asm --enable-asm-asm \  
--with-mysqld-ldflags=-all-static
```

The full `configure` line would, in other words, be something like the following for all recent `gcc` versions:

```
CFLAGS="-O3 -mpentiumpro" CXX=gcc CXXFLAGS="-O3 -mpentiumpro \  
-felide-constructors -fno-exceptions -fno-rtti" ./configure \  
--prefix=/usr/local/mysql --enable-asm --enable-asm-asm \  
--with-mysqld-ldflags=-all-static
```

The binaries we provide on the MySQL Web site at <http://dev.mysql.com/downloads/> are all compiled with full optimization and should be perfect for most users. See [Section 2.1.2.5, “MySQL Binaries Compiled by MySQL AB”](#). There are some configuration settings you can tweak to build an even faster binary, but these are only for advanced users. See [Section 7.5.3, “How Compiling and Linking Affects the Speed of MySQL”](#).

If the build fails and produces errors about your compiler or linker not being able to create the shared library `libmysqlclient.so.N` (where `N` is a version number), you can work around this problem by giving the `--disable-shared` option to `configure`. In this case, `configure` does not build a shared `libmysqlclient.so.N` library.

- By default, MySQL uses the `latin1` (cp1252 West European) character set. To change the default set, use the `--with-charset` option:

```
shell> ./configure --with-charset=CHARSET
```

`CHARSET` may be one of `big5`, `cp1251`, `cp1257`, `czech`, `danish`, `dec8`, `dos`, `euc_kr`, `gb2312`, `gbk`, `german1`, `hebrew`, `hp8`, `hungarian`, `koi8_ru`, `koi8_ukr`, `latin1`, `latin2`, `sjis`, `swe7`, `tis620`, `ujis`, `usa7`, or `win1251ukr`. See [Section 5.10.1, “The Character Set Used for Data and Sorting”](#). (Additional character sets might be available. Check the output from `./configure --help` for the current list.)

As of MySQL 4.1.1, the default collation may also be specified. MySQL uses the `latin1_swedish_ci` collation. To change this, use the `--with-collation` option:

```
shell> ./configure --with-collation=COLLATION
```

To change both the character set and the collation, use both the `--with-charset` and `--with-collation` options. The collation must be a legal collation for the character set. (Use the `SHOW COLLATION` statement to determine which collations are available for each character set.)

**Warning:** If you change character sets after having created any tables, you have to run `myisamchk`

`-r -q --set-collation=collation_name` on every MyISAM table. Your indexes may be sorted incorrectly otherwise. This can happen if you install MySQL, create some tables, and then reconfigure MySQL to use a different character set and reinstall it. (Use `--set-character-set` before MySQL 4.1.1.)

With the `configure` option `--with-extra-charsets=LIST`, you can define which additional character sets should be compiled into the server. *LIST* is one of the following:

- A list of character set names separated by spaces
- `complex` to include all character sets that can't be dynamically loaded
- `all` to include all character sets into the binaries

Clients that want to convert characters between the server and the client should use the `SET NAMES` statement. See [Section 13.5.3, “SET Syntax”](#), and [Section 10.4, “Connection Character Sets and Collations”](#).

- To configure MySQL with debugging code, use the `--with-debug` option:

```
shell> ./configure --with-debug
```

This causes a safe memory allocator to be included that can find some errors and that provides output about what is happening. See [Section E.1, “Debugging a MySQL Server”](#).

- If your client programs are using threads, you must compile a thread-safe version of the MySQL client library with the `--enable-thread-safe-client` configure option. This creates a `libmysqlclient_r` library with which you should link your threaded applications. See [Section 17.2.15, “How to Make a Threaded Client”](#).
- It is now possible to build MySQL with big table support using the `--with-big-tables` option, beginning with the following MySQL versions:
  - **4.0 series:** 4.0.25
  - **4.1 series:** 4.1.11

This option causes the variables that store table row counts to be declared as `unsigned long long` rather than `unsigned long`. This enables tables to hold up to approximately  $1.844\text{E}+19$  ( $(2^{32})^2$ ) rows rather than  $2^{32}$  ( $\sim 4.295\text{E}+09$ ) rows. Previously it was necessary to pass `-DBIG_TABLES` to the compiler manually in order to enable this feature.

- See [Section 2.13, “Operating System-Specific Notes”](#), for options that pertain to particular operating systems.
- See [Section 5.8.7.2, “Using SSL Connections”](#), for options that pertain to configuring MySQL to support secure (encrypted) connections.

### 2.9.3. Installing from the Development Source Tree

**Caution:** You should read this section only if you are interested in helping us test new code. If you just want to get MySQL up and running on your system, you should use a standard release distribution (either a binary or source distribution).

To obtain our most recent development source tree, first download and install the BitKeeper free client if you do not have it. The client can be obtained from <http://www.bitmover.com/bk-client.shar>.

To install the BitKeeper client on Unix, use these commands:

```
shell> sh bk-client.shar
shell> cd bk_client-1.1
shell> make all
shell> PATH=$PWD:$PATH
```

To install the BitKeeper client on Windows, use these instructions:

1. Download and install Cygwin from <http://cygwin.com> [<http://cygwin.com/>].
2. Make sure `gcc` and `make` have been installed under Cygwin. You can test this by issuing `which gcc` and `which make` commands. If either one is not installed, run Cygwin's package manager, select `gcc`, `make`, or both, and install them.
3. Under Cygwin, execute these commands:

```
shell> sh bk-client.shar
shell> cd bk_client-1.1
```

Then edit the `Makefile` and change the line that reads `$(CC) $(CFLAGS) -o sfio -lz sfio.c` to this:

```
$(CC) $(CFLAGS) -o sfio sfio.c -lz
```

Now run the `make` command and set the path:

```
shell> make all
shell> PATH=$PWD:$PATH
```

The BitKeeper free client is shipped with its source code. The only documentation available for the free client is the source code itself.

After you have installed the BitKeeper client, you can access the MySQL development source tree:

1. Change location to the directory you want to work from, and then use one of the following commands to make a local copy of the MySQL version branch of your choice:

To copy the 3.23 branch, use this command:

```
shell> sfioball -r+ bk://mysql.bkbits.net/mysql-3.23 mysql-3.23
```

To copy the 4.0 branch, use this command:

```
shell> sfioball -r+ bk://mysql.bkbits.net/mysql-4.0 mysql-4.0
```

To copy the 4.1 branch, use this command:

```
shell> sfioball -r+ bk://mysql.bkbits.net/mysql-4.1 mysql-4.1
```

In the preceding examples, the source tree is set up in the `mysql-3.23/`, `mysql-4.0/`, or `mysql-4.1/` subdirectory of your current directory.

The initial download of the source tree may take a while, depending on the speed of your connection. Please be patient.

2. You need GNU `make`, `autoconf` 2.58 (or newer), `automake` 1.8, `libtool` 1.5, and `m4` to run the next set of commands. Even though many operating systems come with their own implementation of `make`, chances are high that the compilation fails with strange error messages. Therefore, it is highly recommended that you use GNU `make` (sometimes named `gmake`) instead.

Fortunately, a large number of operating systems ship with the GNU toolchain preinstalled or supply installable packages of these. They can also be downloaded from the following locations:

- <http://www.gnu.org/software/autoconf/>
- <http://www.gnu.org/software/automake/>
- <http://www.gnu.org/software/libtool/>
- <http://www.gnu.org/software/m4/>
- <http://www.gnu.org/software/make/>

If you are trying to configure MySQL 4.1, you also need GNU `bison` 1.75 or later. Older versions of `bison` may report this error:

```
sql_yacc.yy:#####: fatal error: maximum table size (32767) exceeded
```

Note: The maximum table size is not actually exceeded; the error is caused by bugs in older versions of `bison`.

Versions of MySQL before version 4.1 may also compile with other `yacc` implementations (for example, BSD `yacc` 91.7.30). For later versions, GNU `bison` is required.

The following example shows the typical commands required to configure a source tree. The first `cd` command changes location into the top-level directory of the tree; replace `mysql-5.0` with the appropriate directory name.

```
shell> cd mysql-5.0
shell> (cd bdb/dist; sh s_all)
shell> (cd innobase; autoreconf --force --install)
shell> autoreconf --force --install
shell> ./configure # Add your favorite options here
shell> make
```

Or you can use `BUILD/autorun.sh` as a shortcut for the following sequence of commands:

```
shell> alocal; autoheader
shell> libtoolize --automake --force
shell> automake --force --add-missing; autoconf
shell> (cd innobase; alocal; autoheader; autoconf; automake)
shell> (cd bdb/dist; sh s_all)
```

The command lines that change directory into the `innobase` and `bdb/dist` directories are used to configure the `InnoDB` and Berkeley DB (`BDB`) storage engines. You can omit these command lines if you do not require `InnoDB` or `BDB` support.

If you get some strange errors during this stage, verify that you really have `libtool` installed.

A collection of our standard configuration scripts is located in the `BUILD` subdirectory. You may find it more convenient to use the `BUILD/compile-pentium-debug` script than the preceding set of shell commands. To compile on a different architecture, modify the script by removing flags that are Pentium-specific.

3. When the build is done, run `make install`. Be careful with this on a production machine; the

command may overwrite your live release installation. If you have another installation of MySQL, we recommend that you run `./configure` with different values for the `--prefix`, `--with-tcp-port`, and `--unix-socket-path` options than those used for your production server.

4. Play hard with your new installation and try to make the new features crash. Start by running `make test`. See [Section 19.1.2, “MySQL Test Suite”](#).
5. If you have gotten to the `make` stage, but the distribution does not compile, please enter the problem into our bugs database using the instructions given in [Section 1.8, “How to Report Bugs or Problems”](#). If you have installed the latest versions of the required GNU tools, and they crash trying to process our configuration files, please report that also. However, if you execute `aclocal` and get a `command not found` error or a similar problem, do not report it. Instead, make sure that all the necessary tools are installed and that your `PATH` variable is set correctly so that your shell can find them.
6. After initially copying the repository with `sfioball` to obtain the source tree, you should use `update` periodically to update your local copy. To do this any time after you have set up the repository, use this command: (the example is for updating the 4.1 branch):

```
shell> update bk://mysql.bkbits.net/mysql-4.1
```

7. You can examine the change history for the tree with all the diffs by viewing the [BK/ChangeLog](#) file in the source tree and looking at the [ChangeSet](#) descriptions listed there. To examine a particular changeset, you would have to use the `sfioball` command to extract two particular revisions of the source tree, and then use an external `diff` command to compare them. If you see diffs or code that you have a question about, do not hesitate to send email to the MySQL [internals](#) mailing list. See [Section 1.7.1, “MySQL Mailing Lists”](#). If you think you have a better idea on how to do something, send an email message to the same address with a patch.

You can also browse changesets, comments, and source code online. For example, to browse this information for MySQL 4.1, go to <http://mysql.bkbits.net:8080/mysql-4.1>.

## 2.9.4. Dealing with Problems Compiling MySQL

All MySQL programs compile cleanly for us with no warnings on Solaris or Linux using `gcc`. On other systems, warnings may occur due to differences in system include files. See [Section 2.9.5, “MIT-pthreads Notes”](#), for warnings that may occur when using MIT-pthreads. For other problems, check the following list.

The solution to many problems involves reconfiguring. If you do need to reconfigure, take note of the following:

- If `configure` is run after it has previously been run, it may use information that was gathered during its previous invocation. This information is stored in `config.cache`. When `configure` starts up, it looks for that file and reads its contents if it exists, on the assumption that the information is still correct. That assumption is invalid when you reconfigure.
- Each time you run `configure`, you must run `make` again to recompile. However, you may want to remove old object files from previous builds first because they were compiled using different configuration options.

To prevent old configuration information or object files from being used, run these commands before re-running `configure`:

```
shell> rm config.cache
shell> make clean
```

Alternatively, you can run `make distclean`.

The following list describes some of the problems when compiling MySQL that have been found to occur most often:

- If you get errors such as the ones shown here when compiling `sql_yacc.cc`, you probably have run out of memory or swap space:

```
Internal compiler error: program cclplus got fatal signal 11
Out of virtual memory
Virtual memory exhausted
```

The problem is that `gcc` requires a huge amount of memory to compile `sql_yacc.cc` with inline functions. Try running `configure` with the `--with-low-memory` option:

```
shell> ./configure --with-low-memory
```

This option causes `-fno-inline` to be added to the compile line if you are using `gcc` and `-O0` if you are using something else. You should try the `--with-low-memory` option even if you have so much memory and swap space that you think you can't possibly have run out. This problem has been observed to occur even on systems with generous hardware configurations, and the `--with-low-memory` option usually fixes it.

- By default, `configure` picks `c++` as the compiler name and GNU `c++` links with `-lg++`. If you are using `gcc`, that behavior can cause problems during configuration such as this:

```
configure: error: installation or configuration problem:
C++ compiler cannot create executables.
```

You might also observe problems during compilation related to `g++`, `libg++`, or `libstdc++`.

One cause of these problems is that you may not have `g++`, or you may have `g++` but not `libg++`, or `libstdc++`. Take a look at the `config.log` file. It should contain the exact reason why your C++ compiler did not work. To work around these problems, you can use `gcc` as your C++ compiler. Try setting the environment variable `CXX` to `"gcc -O3"`. For example:

```
shell> CXX="gcc -O3" ./configure
```

This works because `gcc` compiles C++ source files as well as `g++` does, but does not link in `libg++` or `libstdc++` by default.

Another way to fix these problems is to install `g++`, `libg++`, and `libstdc++`. However, we recommend that you not use `libg++` or `libstdc++` with MySQL because this only increases the binary size of `mysqld` without providing any benefits. Some versions of these libraries have also caused strange problems for MySQL users in the past.

Using `gcc` as the C++ compiler is also required if you want to compile MySQL with RAID functionality (see [Section 13.1.5, "CREATE TABLE Syntax"](#), for more info on RAID table type) and you are using GNU `gcc` version 3 and above. If you get errors like those following during the linking stage when you configure MySQL to compile with the option `--with-raid`, try to use `gcc` as your C++ compiler by defining the `CXX` environment variable:

```
gcc -O3 -DDEBUG_OFF -rdynamic -o isamchk isamchk.o sort.o libnisam.a
./mysys/libmysys.a ./debug/libdbug.a ./strings/libmystrings.a
-lpthread -lz -lcrypt -lnsl -lm -lpthread
```

```
./mysys/libmysys.a(raid.o)(.text+0x79): In function
`my_raid_create'::: undefined reference to `operator new(unsigned)'
./mysys/libmysys.a(raid.o)(.text+0xdd): In function
`my_raid_create'::: undefined reference to `operator delete(void*)'
./mysys/libmysys.a(raid.o)(.text+0x129): In function
`my_raid_open'::: undefined reference to `operator new(unsigned)'
./mysys/libmysys.a(raid.o)(.text+0x189): In function
`my_raid_open'::: undefined reference to `operator delete(void*)'
./mysys/libmysys.a(raid.o)(.text+0x64b): In function
`my_raid_close'::: undefined reference to `operator delete(void*)'
collect2: ld returned 1 exit status
```

- If your compilation fails with errors such as any of the following, you must upgrade your version of [make](#) to GNU [make](#):

```
making all in mit-pthreads
make: Fatal error in reader: Makefile, line 18:
Badly formed macro assignment
```

Or:

```
make: file `Makefile' line 18: Must be a separator (:
```

Or:

```
pthread.h: No such file or directory
```

Solaris and FreeBSD are known to have troublesome [make](#) programs.

GNU [make](#) 3.75 is known to work.

- If you want to define flags to be used by your C or C++ compilers, do so by adding the flags to the [CFLAGS](#) and [CXXFLAGS](#) environment variables. You can also specify the compiler names this way using [CC](#) and [CXX](#). For example:

```
shell> CC=gcc
shell> CFLAGS=-O3
shell> CXX=gcc
shell> CXXFLAGS=-O3
shell> export CC CFLAGS CXX CXXFLAGS
```

See [Section 2.1.2.5, “MySQL Binaries Compiled by MySQL AB”](#), for a list of flag definitions that have been found to be useful on various systems.

- If you get errors such as those shown here when compiling [mysqld](#), [configure](#) did not correctly detect the type of the last argument to [accept\(\)](#), [getsockname\(\)](#), or [getpeername\(\)](#):

```
cxx: Error: mysqld.cc, line 645: In this statement, the referenced
      type of the pointer value 'length' is 'unsigned long',
      which is not compatible with 'int'.
new_sock = accept(sock, (struct sockaddr *)&cAddr, &length);
```

To fix this, edit the [config.h](#) file (which is generated by [configure](#)). Look for these lines:

```
/* Define as the base type of the last arg to accept */
#define SOCKET_SIZE_TYPE XXX
```

Change [XXX](#) to [size\\_t](#) or [int](#), depending on your operating system. (You must do this each time you run [configure](#) because [configure](#) regenerates [config.h](#).)

- The [sql\\_yacc.cc](#) file is generated from [sql\\_yacc.yy](#). Normally, the build process does not

need to create `sql_yacc.cc` because MySQL comes with a pre-generated copy. However, if you do need to re-create it, you might encounter this error:

```
"sql_yacc.yy", line xxx fatal: default action causes potential...
```

This is a sign that your version of `yacc` is deficient. You probably need to install `bison` (the GNU version of `yacc`) and use that instead.

- On Debian Linux 3.0, you need to install `gawk` instead of the default `mawk` if you want to compile MySQL 4.1 or higher with Berkeley DB support.
- If you need to debug `mysqld` or a MySQL client, run `configure` with the `--with-debug` option, and then recompile and link your clients with the new client library. See [Section E.2, “Debugging a MySQL Client”](#).
- If you get a compilation error on Linux (for example, SuSE Linux 8.1 or Red Hat Linux 7.3) similar to the following one, you probably do not have `g++` installed:

```
libmysql.c:1329: warning: passing arg 5 of `gethostbyname_r' from
incompatible pointer type
libmysql.c:1329: too few arguments to function `gethostbyname_r'
libmysql.c:1329: warning: assignment makes pointer from integer
without a cast
make[2]: *** [libmysql.lo] Error 1
```

By default, the `configure` script attempts to determine the correct number of arguments by using `g++` (the GNU C++ compiler). This test yields incorrect results if `g++` is not installed. There are two ways to work around this problem:

- Make sure that the GNU C++ `g++` is installed. On some Linux distributions, the required package is called `gpp`; on others, it is named `gcc-c++`.
- Use `gcc` as your C++ compiler by setting the `CXX` environment variable to `gcc`:

```
export CXX="gcc"
```

You must run `configure` again after making either of those changes.

## 2.9.5. MIT-pthreads Notes

This section describes some of the issues involved in using MIT-pthreads.

On Linux, you should *not* use MIT-pthreads. Use the installed LinuxThreads implementation instead. See [Section 2.13.1, “Linux Notes”](#).

If your system does not provide native thread support, you should build MySQL using the MIT-pthreads package. This includes older FreeBSD systems, SunOS 4.x, Solaris 2.4 and earlier, and some others. See [Section 2.1.1, “Operating Systems Supported by MySQL”](#).

Beginning with MySQL 4.0.2, MIT-pthreads is no longer part of the source distribution. If you require this package, you need to download it separately from [http://dev.mysql.com/Downloads/Contrib/pthreads-1\\_60\\_beta6-mysql.tar.gz](http://dev.mysql.com/Downloads/Contrib/pthreads-1_60_beta6-mysql.tar.gz)

After downloading, extract this source archive into the top level of the MySQL source directory. It creates a new subdirectory named `mit-pthreads`.

- On most systems, you can force MIT-pthreads to be used by running `configure` with the `-`



`-with-mit-threads` option:

```
shell> ./configure --with-mit-threads
```

Building in a non-source directory is not supported when using MIT-pthreads because we want to minimize our changes to this code.

- The checks that determine whether to use MIT-pthreads occur only during the part of the configuration process that deals with the server code. If you have configured the distribution using `-without-server` to build only the client code, clients do not know whether MIT-pthreads is being used and use Unix socket file connections by default. Because Unix socket files do not work under MIT-pthreads on some platforms, this means you need to use `-h` or `--host` with a value other than `localhost` when you run client programs.
- When MySQL is compiled using MIT-pthreads, system locking is disabled by default for performance reasons. You can tell the server to use system locking with the `--external-locking` option. This is needed only if you want to be able to run two MySQL servers against the same data files, but that is not recommended, anyway.
- Sometimes the pthread `bind()` command fails to bind to a socket without any error message (at least on Solaris). The result is that all connections to the server fail. For example:

```
shell> mysqladmin version
mysqladmin: connect to server at '' failed;
error: 'Can't connect to mysql server on localhost (146)'
```

The solution to this problem is to kill the `mysqld` server and restart it. This has happened to us only when we forced the server to shut down and then restarted it immediately.

- With MIT-pthreads, the `sleep()` system call isn't interruptible with `SIGINT` (break). This is noticeable only when you run `mysqladmin --sleep`. You must wait for the `sleep()` call to terminate before the interrupt is served and the process stops.
- When linking, you might receive warning messages like these (at least on Solaris); they can be ignored:

```
ld: warning: symbol `__iob' has differing sizes:
(file /my/local/pthreads/lib/libpthread.a(findfp.o) value=0x4;
file /usr/lib/libc.so value=0x140);
/my/local/pthreads/lib/libpthread.a(findfp.o) definition taken
ld: warning: symbol `__iob' has differing sizes:
(file /my/local/pthreads/lib/libpthread.a(findfp.o) value=0x4;
file /usr/lib/libc.so value=0x140);
/my/local/pthreads/lib/libpthread.a(findfp.o) definition taken
```

- Some other warnings also can be ignored:

```
implicit declaration of function `int strtoll(...)'
implicit declaration of function `int strtoul(...)'
```

- We have not gotten `readline` to work with MIT-pthreads. (This is not needed, but may be interesting for someone.)

## 2.9.6. Installing MySQL from Source on Windows

These instructions describe how to build binaries from source for MySQL 4.1 and up on Windows. Instructions are provided for building binaries from a standard source distribution or from the BitKeeper tree that contains the latest development source.

**Note:** The instructions here are strictly for users who want to test MySQL on Windows from the latest source distribution or from the BitKeeper tree. For production use, MySQL AB does not advise using a MySQL server built by yourself from source. Normally, it is best to use precompiled binary distributions of MySQL that are built specifically for optimal performance on Windows by MySQL AB. Instructions for installing a binary distributions are available in [Section 2.3](#), “Installing MySQL on Windows”.

To build MySQL on Windows from source, you need the following compiler and resources available on your Windows system:

- Visual Studio .Net 2003 (7.1) compiler system
- Between 3GB and 5GB disk space.
- Windows XP, Windows 2000 or higher.

The exact system requirements can be found here: <http://msdn.microsoft.com/vstudio/Previous/2003/sysreqs/default.aspx>

You also need a MySQL source distribution for Windows. For MySQL 4.1 and up, there are two ways to obtain a source distribution:

1. Obtain a Windows source distribution packaged by MySQL AB for the particular version of MySQL in which you are interested. These are available from <http://dev.mysql.com/downloads/>.
2. You can package a source distribution yourself from the latest BitKeeper developer source tree. If you plan to do this, you must create the package on a Unix system and then transfer it to your Windows system. (Some of the configuration and build steps require tools that work only on Unix.) The BitKeeper approach thus requires:
  - A system running Unix, or a Unix-like system such as Linux.
  - BitKeeper installed on that system. See [Section 2.9.3](#), “Installing from the Development Source Tree”, for instructions how to download and install BitKeeper.

If you are using a Windows source distribution, you can go directly to [Section 2.9.6.1](#), “Building MySQL Using VC++”. To build from the BitKeeper tree, proceed to [Section 2.9.6.2](#), “Creating a Windows Source Package from the Latest Development Source”.

If you find something not working as expected, or you have suggestions about ways to improve the current build process on Windows, please send a message to the [win32](#) mailing list. See [Section 1.7.1](#), “MySQL Mailing Lists”.

### 2.9.6.1. Building MySQL Using VC++

**Note:** VC++ workspace files for MySQL 4.1 and above are compatible with Microsoft Visual Studio 7.1 and tested by MySQL AB staff before each release.

Follow this procedure to build MySQL:

1. Create a work directory (for example, `C:\workdir`).
2. Unpack the source distribution in the aforementioned directory using [WinZip](#) or other Windows tool that can read `.zip` files.

3. Start Visual Studio .Net 2003 (7.1).
4. From the File menu, select Open Solution....
5. Open the `mysql.sln` solution you find in the work directory.
6. From the Build menu, select Configuration Manager....
7. In the Active Solution Configuration pop-up menu, select the configuration to use. You likely want to use one of `nt` (normal server, not for Windows 98/ME), `Max nt` (more engines and features, not for 98/ME) or `Debug` configuration.
8. From the Build menu, select Build Solution.
9. Debug versions of the programs and libraries are placed in the `client_debug` and `lib_debug` directories. Release versions of the programs and libraries are placed in the `client_release` and `lib_release` directories.
10. Test the server. The server built using the preceding instructions expects that the MySQL base directory and data directory are `C:\mysql` and `C:\mysql\data` by default. If you want to test your server using the source tree root directory and its data directory as the base directory and data directory, you need to tell the server their pathnames. You can either do this on the command line with the `--basedir` and `--datadir` options, or by placing appropriate options in an option file. (See [Section 4.3.2, “Using Option Files”](#).) If you have an existing data directory elsewhere that you want to use, you can specify its pathname instead.
11. Start your server from the `client_release` or `client_debug` directory, depending on which server you built or want to use. The general server startup instructions are in [Section 2.3, “Installing MySQL on Windows”](#). You must adapt the instructions appropriately if you want to use a different base directory or data directory.
12. When the server is running in standalone fashion or as a service based on your configuration, try to connect to it from the `mysql` interactive command-line utility that exists in your `client_release` or `client_debug` directory.

When you are satisfied that the programs you have built are working correctly, stop the server. Then install MySQL as follows:

1. Create the directories where you want to install MySQL. For example, to install into `C:\mysql`, use these commands:

```
C:\> mkdir C:\mysql
C:\> mkdir C:\mysql\bin
C:\> mkdir C:\mysql\data
C:\> mkdir C:\mysql\share
C:\> mkdir C:\mysql\scripts
```

If you want to compile other clients and link them to MySQL, you should also create several additional directories:

```
C:\> mkdir C:\mysql\include
C:\> mkdir C:\mysql\lib
C:\> mkdir C:\mysql\lib\debug
C:\> mkdir C:\mysql\lib\opt
```

If you want to benchmark MySQL, create this directory:

```
C:\> mkdir C:\mysql\sql-bench
```

Benchmarking requires Perl support. See [Section 2.14, “Perl Installation Notes”](#).

- From the `workdir` directory, copy into the `C:\mysql` directory the following directories:

```
C:\> cd \workdir
C:\workdir> copy client_release\*.exe C:\mysql\bin
C:\workdir> copy client_debug\mysqld.exe C:\mysql\bin\mysqld-debug.exe
C:\workdir> xcopy scripts\*. * C:\mysql\scripts /E
C:\workdir> xcopy share\*. * C:\mysql\share /E
```

If you want to compile other clients and link them to MySQL, you should also copy several libraries and header files:

```
C:\workdir> copy lib_debug\mysqlclient.lib C:\mysql\lib\debug
C:\workdir> copy lib_debug\libmysql.* C:\mysql\lib\debug
C:\workdir> copy lib_debug\zlib.* C:\mysql\lib\debug
C:\workdir> copy lib_release\mysqlclient.lib C:\mysql\lib\opt
C:\workdir> copy lib_release\libmysql.* C:\mysql\lib\opt
C:\workdir> copy lib_release\zlib.* C:\mysql\lib\opt
C:\workdir> copy include\*.h C:\mysql\include
C:\workdir> copy libmysql\libmysql.def C:\mysql\include
```

If you want to benchmark MySQL, you should also do this:

```
C:\workdir> xcopy sql-bench\*. * C:\mysql\bench /E
```

Set up and start the server in the same way as for the binary Windows distribution. See [Section 2.3, “Installing MySQL on Windows”](#).

### 2.9.6.2. Creating a Windows Source Package from the Latest Development Source

To create a Windows source package from the current BitKeeper source tree, use the instructions here. This procedure must be performed on a system running a Unix or Unix-like operating system because some of the configuration and build steps require tools that work only on Unix. For example, the following procedure is known to work well on Linux.

- Copy the BitKeeper source tree for MySQL 4.1 (or above, as desired). For instructions on how to do this, see [Section 2.9.3, “Installing from the Development Source Tree”](#).
- Configure and build the distribution so that you have a server binary to work with. One way to do this is to run the following command in the top-level directory of your source tree:

```
shell> ./BUILD/compile-pentium-max
```

- After making sure that the build process completed successfully, run the following utility script from top-level directory of your source tree:

```
shell> ./scripts/make_win_src_distribution
```

This script creates a Windows source package to be used on your Windows system. You can supply different options to the script based on your needs. It accepts the following options:

- `--help`

Display a help message.

- `--debug`  
Print information about script operations, do not create package.
- `--tmp`  
Specify the temporary location.
- `--suffix`  
The suffix name for the package.
- `--dirname`  
Directory name to copy files (intermediate).
- `--silent`  
Do not print verbose list of files processed.
- `--tar`  
Create `tar.gz` package instead of `.zip` package.

By default, `make_win_src_distribution` creates a Zip-format archive with the name `mysql-VERSION-win-src.zip`, where `VERSION` represents the version of your MySQL source tree.

4. Copy or upload the Windows source package that you have just created to your Windows machine. To compile it, use the instructions in [Section 2.9.6.1, “Building MySQL Using VC++”](#).

## 2.9.7. Compiling MySQL Clients on Windows

In your source files, you should include `my_global.h` before `mysql.h`:

```
#include <my_global.h>
#include <mysql.h>
```

`my_global.h` includes any other files needed for Windows compatibility (such as `windows.h`) if you compile your program on Windows.

You can either link your code with the dynamic `libmysql.lib` library, which is just a wrapper to load in `libmysql.dll` on demand, or link with the static `mysqlclient.lib` library.

The MySQL client libraries are compiled as threaded libraries, so you should also compile your code to be multi-threaded.

## 2.10. Post-Installation Setup and Testing

After installing MySQL, there are some issues that you should address. For example, on Unix, you should initialize the data directory and create the MySQL grant tables. On all platforms, an important security concern is that the initial accounts in the grant tables have no passwords. You should assign passwords to prevent unauthorized access to the MySQL server. Optionally, for MySQL 4.1.3 and up, you can create time zone tables to enable recognition of named time zones.

The following sections include post-installation procedures that are specific to Windows systems and to

Unix systems. Another section, [Section 2.10.2.3, “Starting and Troubleshooting the MySQL Server”](#), applies to all platforms; it describes what to do if you have trouble getting the server to start. [Section 2.10.3, “Securing the Initial MySQL Accounts”](#), also applies to all platforms. You should follow its instructions to make sure that you have properly protected your MySQL accounts by assigning passwords to them.

When you are ready to create additional user accounts, you can find information on the MySQL access control system and account management in [Section 5.7, “The MySQL Access Privilege System”](#), and [Section 5.8, “MySQL User Account Management”](#).

## 2.10.1. Windows Post-Installation Procedures

On Windows, the data directory and the grant tables do not have to be created. MySQL Windows distributions include the grant tables with a set of preinitialized accounts in the `mysql` database under the data directory. It is unnecessary to run the `mysql_install_db` script that is used on Unix. Regarding passwords, if you installed MySQL using the Windows Installation Wizard, you may have already assigned passwords to the accounts. (See [Section 2.3.3, “Using the MySQL Installation Wizard”](#).) Otherwise, use the password-assignment procedure given in [Section 2.10.3, “Securing the Initial MySQL Accounts”](#).

Before setting up passwords, you might want to try running some client programs to make sure that you can connect to the server and that it is operating properly. Make sure that the server is running (see [Section 2.3.9, “Starting the Server for the First Time”](#)), and then issue the following commands to verify that you can retrieve information from the server. The output should be similar to what is shown here:

```
C:\> C:\mysql\bin\mysqlshow
+-----+
| Databases |
+-----+
| mysql    |
| test     |
+-----+

C:\> C:\mysql\bin\mysqlshow mysql
Database: mysql
+-----+
|          Tables          |
+-----+
| columns_priv |
| db           |
| func        |
| help_category |
| help_keyword |
| help_relation |
| help_topic  |
| host        |
| proc       |
| procs_priv  |
| tables_priv |
| time_zone  |
| time_zone_leap_second |
| time_zone_name |
| time_zone_transition |
| time_zone_transition_type |
| user       |
+-----+

C:\> C:\mysql\bin\mysql -e "SELECT Host,Db,User FROM db" mysql
+-----+
| host | db   | user |
+-----+
| %    | test% |    |
+-----+
```

If you are running a version of Windows that supports services and you want the MySQL server to run automatically when Windows starts, see [Section 2.3.11, “Starting MySQL as a Windows Service”](#).

## 2.10.2. Unix Post-Installation Procedures

After installing MySQL on Unix, you need to initialize the grant tables, start the server, and make sure that the server works satisfactorily. You may also wish to arrange for the server to be started and stopped automatically when your system starts and stops. You should also assign passwords to the accounts in the grant tables.

On Unix, the grant tables are set up by the `mysql_install_db` program. For some installation methods, this program is run for you automatically:

- If you install MySQL on Linux using RPM distributions, the server RPM runs `mysql_install_db`.
- If you install MySQL on Mac OS X using a PKG distribution, the installer runs `mysql_install_db`.

Otherwise, you will need to run `mysql_install_db` yourself.

The following procedure describes how to initialize the grant tables (if that has not previously been done) and then start the server. It also suggests some commands that you can use to test whether the server is accessible and working properly. For information about starting and stopping the server automatically, see [Section 2.10.2.2, “Starting and Stopping MySQL Automatically”](#).

After you complete the procedure and have the server running, you should assign passwords to the accounts created by `mysql_install_db`. Instructions for doing so are given in [Section 2.10.3, “Securing the Initial MySQL Accounts”](#).

In the examples shown here, the server runs under the user ID of the `mysql` login account. This assumes that such an account exists. Either create the account if it does not exist, or substitute the name of a different existing login account that you plan to use for running the server.

1. Change location into the top-level directory of your MySQL installation, represented here by `BASEDIR`:

```
shell> cd BASEDIR
```

`BASEDIR` is likely to be something like `/usr/local/mysql` or `/usr/local`. The following steps assume that you are located in this directory.

2. If necessary, run the `mysql_install_db` program to set up the initial MySQL grant tables containing the privileges that determine how users are allowed to connect to the server. You'll need to do this if you used a distribution type for which the installation procedure doesn't run the program for you.

Typically, `mysql_install_db` needs to be run only the first time you install MySQL, so you can skip this step if you are upgrading an existing installation. However, `mysql_install_db` does not overwrite any existing privilege tables, so it should be safe to run in any circumstances.

To initialize the grant tables, use one of the following commands, depending on whether `mysql_install_db` is located in the `bin` or `scripts` directory:

```
shell> bin/mysql_install_db --user=mysql
shell> scripts/mysql_install_db --user=mysql
```

The `mysql_install_db` script creates the server's data directory. Under the data directory, it creates directories for the `mysql` database that holds all database privileges and the `test` database

that you can use to test MySQL. The script also creates privilege table entries for `root` and anonymous-user accounts. The accounts have no passwords initially. A description of their initial privileges is given in [Section 2.10.3, “Securing the Initial MySQL Accounts”](#). Briefly, these privileges allow the MySQL `root` user to do anything, and allow anybody to create or use databases with a name of `test` or starting with `test_`.

It is important to make sure that the database directories and files are owned by the `mysql` login account so that the server has read and write access to them when you run it later. To ensure this, the `--user` option should be used as shown if you run `mysql_install_db` as `root`. Otherwise, you should execute the script while logged in as `mysql`, in which case you can omit the `--user` option from the command.

`mysql_install_db` creates several tables in the `mysql` database: `user`, `db`, `host`, `tables_priv`, `columns_priv`, `func`, and possibly others depending on your version of MySQL.

If you do not want to have the `test` database, you can remove it with `mysqladmin -u root drop test` after starting the server.

If you have trouble with `mysql_install_db` at this point, see [Section 2.10.2.1, “Problems Running mysql\\_install\\_db”](#).

3. Start the MySQL server:

```
shell> bin/mysqld_safe --user=mysql &
```

For versions of MySQL older than 4.0, substitute `bin/safe_mysqld` for `bin/mysqld_safe` in this command.

It is important that the MySQL server be run using an unprivileged (non-`root`) login account. To ensure this, the `--user` option should be used as shown if you run `mysql_safe` as `root`. Otherwise, you should execute the script while logged in as `mysql`, in which case you can omit the `--user` option from the command.

Further instructions for running MySQL as an unprivileged user are given in [Section 5.6.5, “How to Run MySQL as a Normal User”](#).

If you neglected to create the grant tables before proceeding to this step, the following message appears in the error log file when you start the server:

```
mysqld: Can't find file: 'host.frm'
```

If you have other problems starting the server, see [Section 2.10.2.3, “Starting and Troubleshooting the MySQL Server”](#).

4. Use `mysqladmin` to verify that the server is running. The following commands provide simple tests to check whether the server is up and responding to connections:

```
shell> bin/mysqladmin version
shell> bin/mysqladmin variables
```

The output from `mysqladmin version` varies slightly depending on your platform and version of MySQL, but should be similar to that shown here:

```
shell> bin/mysqladmin version
mysqladmin Ver 14.7 Distrib 4.1.19, for linux on i586
Copyright (C) 2000 MySQL AB & MySQL Finland AB & TCX DataKonsult AB
This software comes with ABSOLUTELY NO WARRANTY. This is free software,
and you are welcome to modify and redistribute it under the GPL license
```



```

Server version      4.1.19-max
Protocol version   10
Connection         Localhost via Unix socket
TCP port           3306
UNIX socket        var/lib/mysql/mysql.sock
Uptime:            5 days 19 hours 19 min 0 sec

Threads: 1  Questions: 163  Slow queries: 0
Opens: 11  Flush tables:1  Open tables: 0  Queries per second avg: 0.007
Threads: 1  Questions: 9   Slow queries: 0
    
```

To see what else you can do with `mysqladmin`, invoke it with the `--help` option.

5. Verify that you can shut down the server:

```
shell> bin/mysqladmin -u root shutdown
```

6. Verify that you can start the server again. Do this by using `mysqld_safe` or by invoking `mysqld` directly. For example:

```
shell> bin/mysqld_safe --user=mysql --log &
```

If `mysqld_safe` fails, see [Section 2.10.2.3, “Starting and Troubleshooting the MySQL Server”](#).

7. Run some simple tests to verify that you can retrieve information from the server. The output should be similar to what is shown here:

```

shell> bin/mysqlshow
+-----+
| Databases |
+-----+
|  mysql   |
|   test   |
+-----+

shell> bin/mysqlshow mysql
Database: mysql
+-----+
| Tables |
+-----+
| columns_priv |
| db           |
| func         |
| host         |
| tables_priv  |
| user         |
+-----+

shell> bin/mysql -e "SELECT Host,Db,User FROM db" mysql
+-----+-----+-----+
| host | db   | user |
+-----+-----+-----+
| %    | test |      |
| %    | test_% |    |
+-----+-----+-----+
    
```

8. There is a benchmark suite in the `sql-bench` directory (under the MySQL installation directory) that you can use to compare how MySQL performs on different platforms. The benchmark suite is written in Perl. It requires the Perl DBI module that provides a database-independent interface to the various databases, and some other additional Perl modules:

```

DBI
DBD::mysql
Data::Dumper
Data::ShowTable
    
```

These modules can be obtained from CPAN (<http://www.cpan.org/>). See [Section 2.14.1, “Installing](#)

Perl on Unix”.

The `sql-bench/Results` directory contains the results from many runs against different databases and platforms. To run all tests, execute these commands:

```
shell> cd sql-bench
shell> perl run-all-tests
```

If you do not have the `sql-bench` directory, you probably installed MySQL using RPM files other than the source RPM. (The source RPM includes the `sql-bench` benchmark directory.) In this case, you must first install the benchmark suite before you can use it. Beginning with MySQL 3.22, there are separate benchmark RPM files named `mysql-bench-VERSION-i386.rpm` that contain benchmark code and data.

If you have a source distribution, there are also tests in its `tests` subdirectory that you can run. For example, to run `auto_increment.tst`, execute this command from the top-level directory of your source distribution:

```
shell> mysql -vfv test < ./tests/auto_increment.tst
```

The expected result of the test can be found in the `./tests/auto_increment.res` file.

- At this point, you should have the server running. However, none of the initial MySQL accounts have a password, so you should assign passwords using the instructions found in [Section 2.10.3, “Securing the Initial MySQL Accounts”](#).

As of MySQL 4.1.3, the installation procedure creates time zone tables in the `mysql` database. However, you must populate the tables manually using the instructions in [Section 5.10.8, “MySQL Server Time Zone Support”](#).

### 2.10.2.1. Problems Running `mysql_install_db`

The purpose of the `mysql_install_db` script is to generate new MySQL privilege tables. It does not overwrite existing MySQL privilege tables, and it does not affect any other data.

If you want to re-create your privilege tables, first stop the `mysqld` server if it is running. Then rename the `mysql` directory under the data directory to save it, and then run `mysql_install_db`. Suppose that your current directory is the MySQL installation directory and that `mysql_install_db` is located in the `bin` directory and the data directory is named `data`. To rename the `mysql` database and re-run `mysql_install_db`, use these commands.

```
shell> mv data/mysql data/mysql.old
shell> bin/mysql_install_db --user=mysql
```

When you run `mysql_install_db`, you might encounter the following problems:

- `mysql_install_db` does not install the grant tables**

You may find that `mysql_install_db` fails to install the grant tables and terminates after displaying the following messages:

```
Starting mysqld daemon with databases from XXXXXX
mysqld ended
```

In this case, you should examine the error log file very carefully. The log should be located in the directory `XXXXXX` named by the error message and should indicate why `mysqld` didn't start. If you

do not understand what happened, include the log when you post a bug report. See [Section 1.8](#), “How to Report Bugs or Problems”.

- **There is a `mysqld` process running**

This indicates that the server is running, in which case the grant tables have probably been created already. If so, there is no need to run `mysql_install_db` at all because it needs to be run only once (when you install MySQL the first time).

- **Installing a second `mysqld` server does not work when one server is running**

This can happen when you have an existing MySQL installation, but want to put a new installation in a different location. For example, you might have a production installation, but you want to create a second installation for testing purposes. Generally the problem that occurs when you try to run a second server is that it tries to use a network interface that is in use by the first server. In this case, you should see one of the following error messages:

```
Can't start server: Bind on TCP/IP port:
Address already in use
Can't start server: Bind on unix socket...
```

For instructions on setting up multiple servers, see [Section 5.12](#), “Running Multiple MySQL Servers on the Same Machine”.

- **You do not have write access to the `/tmp` directory**

If you do not have write access to create temporary files or a Unix socket file in the default location (the `/tmp` directory), an error occurs when you run `mysql_install_db` or the `mysqld` server.

You can specify different locations for the temporary directory and Unix socket file by executing these commands prior to starting `mysql_install_db` or `mysqld`, where `some_tmp_dir` is the full pathname to some directory for which you have write permission:

```
shell> TMPDIR=/some_tmp_dir/
shell> MYSQL_UNIX_PORT=/some_tmp_dir/mysql.sock
shell> export TMPDIR MYSQL_UNIX_PORT
```

Then you should be able to run `mysql_install_db` and start the server with these commands:

```
shell> bin/mysql_install_db --user=mysql
shell> bin/mysqld_safe --user=mysql &
```

If `mysql_install_db` is located in the `scripts` directory, modify the first command to use `scripts/mysql_install_db`.

See [Section A.4.5](#), “How to Protect or Change the MySQL Unix Socket File”, and [Appendix F](#), *Environment Variables*.

There are some alternatives to running the `mysql_install_db` script provided in the MySQL distribution:

- If you want the initial privileges to be different from the standard defaults, you can modify `mysql_install_db` before you run it. However, a preferable technique is to use `GRANT` and `REVOKE` to change the privileges after the grant tables have been set up. In other words, you can run `mysql_install_db`, and then use `mysql -u root mysql` to connect to the server as the MySQL `root` user so that you can issue the `GRANT` and `REVOKE` statements.

If you want to install MySQL on several machines with the same privileges, you can put the `GRANT` and `REVOKE` statements in a file and execute the file as a script using `mysql` after running `mysql_install_db`. For example:

```
shell> bin/mysql_install_db --user=mysql
shell> bin/mysql -u root < your_script_file
```

By doing this, you can avoid having to issue the statements manually on each machine.

- It is possible to re-create the grant tables completely after they have previously been created. You might want to do this if you are just learning how to use `GRANT` and `REVOKE` and have made so many modifications after running `mysql_install_db` that you want to wipe out the tables and start over.

To re-create the grant tables, remove all the `.frm`, `.MYI`, and `.MYD` files in the `mysql` database directory. Then run the `mysql_install_db` script again.

**Note:** For MySQL versions older than 3.22.10, you should not delete the `.frm` files. If you accidentally do this, you should copy them back into the `mysql` directory from your MySQL distribution before running `mysql_install_db`.

- You can start `mysqld` manually using the `--skip-grant-tables` option and add the privilege information yourself using `mysql`:

```
shell> bin/mysqld_safe --user=mysql --skip-grant-tables &
shell> bin/mysql mysql
```

From `mysql`, manually execute the SQL commands contained in `mysql_install_db`. Make sure that you run `mysqladmin flush-privileges` or `mysqladmin reload` afterward to tell the server to reload the grant tables.

Note that by not using `mysql_install_db`, you not only have to populate the grant tables manually, you also have to create them first.

## 2.10.2.2. Starting and Stopping MySQL Automatically

Generally, you start the `mysqld` server in one of these ways:

- By invoking `mysqld` directly. This works on any platform.
- By running the MySQL server as a Windows service. This can be done on versions of Windows that support services (such as NT, 2000, XP, and 2003). The service can be set to start the server automatically when Windows starts, or as a manual service that you start on request. For instructions, see [Section 2.3.11, “Starting MySQL as a Windows Service”](#).
- By invoking `mysqld_safe`, which tries to determine the proper options for `mysqld` and then runs it with those options. This script is used on Unix and Unix-like systems. See [Section 5.4.1, “mysqld\\_safe — MySQL Server Startup Script”](#).
- By invoking `mysql.server`. This script is used primarily at system startup and shutdown on systems that use System V-style run directories, where it usually is installed under the name `mysql`. The `mysql.server` script starts the server by invoking `mysqld_safe`. See [Section 5.4.2, “mysql.server — MySQL Server Startup Script”](#).
- On Mac OS X, you can install a separate MySQL Startup Item package to enable the automatic star-

top of MySQL on system startup. The Startup Item starts the server by invoking `mysql.server`. See [Section 2.5, “Installing MySQL on Mac OS X”](#), for details.

The `mysqld_safe` and `mysql.server` scripts and the Mac OS X Startup Item can be used to start the server manually, or automatically at system startup time. `mysql.server` and the Startup Item also can be used to stop the server.

To start or stop the server manually using the `mysql.server` script, invoke it with `start` or `stop` arguments:

```
shell> mysql.server start
shell> mysql.server stop
```

Before `mysql.server` starts the server, it changes location to the MySQL installation directory, and then invokes `mysqld_safe`. If you want the server to run as some specific user, add an appropriate `user` option to the `[mysqld]` group of the `/etc/my.cnf` option file, as shown later in this section. (It is possible that you will need to edit `mysql.server` if you've installed a binary distribution of MySQL in a non-standard location. Modify it to `cd` into the proper directory before it runs `mysqld_safe`. If you do this, your modified version of `mysql.server` may be overwritten if you upgrade MySQL in the future, so you should make a copy of your edited version that you can reinstall.)

`mysql.server stop` stops the server by sending a signal to it. You can also stop the server manually by executing `mysqladmin shutdown`.

To start and stop MySQL automatically on your server, you need to add start and stop commands to the appropriate places in your `/etc/rc*` files.

If you use the Linux server RPM package (`MySQL-server-VERSION.rpm`), the `mysql.server` script is installed in the `/etc/init.d` directory with the name `mysql`. You need not install it manually. See [Section 2.4, “Installing MySQL on Linux”](#), for more information on the Linux RPM packages.

Some vendors provide RPM packages that install a startup script under a different name such as `mysqld`.

If you install MySQL from a source distribution or using a binary distribution format that does not install `mysql.server` automatically, you can install it manually. The script can be found in the `support-files` directory under the MySQL installation directory or in a MySQL source tree.

To install `mysql.server` manually, copy it to the `/etc/init.d` directory with the name `mysql`, and then make it executable. Do this by changing location into the appropriate directory where `mysql.server` is located and executing these commands:

```
shell> cp mysql.server /etc/init.d/mysql
shell> chmod +x /etc/init.d/mysql
```

Older Red Hat systems use the `/etc/rc.d/init.d` directory rather than `/etc/init.d`. Adjust the preceding commands accordingly. Alternatively, first create `/etc/init.d` as a symbolic link that points to `/etc/rc.d/init.d`:

```
shell> cd /etc
shell> ln -s rc.d/init.d .
```

After installing the script, the commands needed to activate it to run at system startup depend on your operating system. On Linux, you can use `chkconfig`:

```
shell> chkconfig --add mysql
```

On some Linux systems, the following command also seems to be necessary to fully enable the `mysql` script:

```
shell> chkconfig --level 345 mysql on
```

On FreeBSD, startup scripts generally should go in `/usr/local/etc/rc.d/`. The `rc(8)` manual page states that scripts in this directory are executed only if their basename matches the `*.sh` shell file-name pattern. Any other files or directories present within the directory are silently ignored. In other words, on FreeBSD, you should install the `mysql.server` script as `/usr/local/etc/rc.d/mysql.server.sh` to enable automatic startup.

As an alternative to the preceding setup, some operating systems also use `/etc/rc.local` or `/etc/init.d/boot.local` to start additional services on startup. To start up MySQL using this method, you could append a command like the one following to the appropriate startup file:

```
/bin/sh -c 'cd /usr/local/mysql; ./bin/mysqld_safe --user=mysql &'
```

For other systems, consult your operating system documentation to see how to install startup scripts.

You can add options for `mysql.server` in a global `/etc/my.cnf` file. A typical `/etc/my.cnf` file might look like this:

```
[mysqld]
datadir=/usr/local/mysql/var
socket=/var/tmp/mysql.sock
port=3306
user=mysql

[mysql.server]
basedir=/usr/local/mysql
```

The `mysql.server` script understands the following options: `basedir`, `datadir`, and `pid-file`. If specified, they *must* be placed in an option file, not on the command line. `mysql.server` understands only `start` and `stop` as command-line arguments.

The following table shows which option groups the server and each startup script read from option files:

Script	Option Groups
<code>mysqld</code>	<code>[mysqld]</code> , <code>[server]</code> , <code>[mysqld-major_version]</code>
<code>mysqld_safe</code>	<code>[mysqld]</code> , <code>[server]</code> , <code>[mysqld_safe]</code>
<code>mysql.server</code>	<code>[mysqld]</code> , <code>[mysql.server]</code>

`[mysqld-major_version]` means that groups with names like `[mysqld-4.0]`, `[mysqld-4.1]`, and `[mysqld-5.0]` are read by servers having versions 4.0.x, 4.1.x, 5.0.x, and so forth. This feature was added in MySQL 4.0.14. It can be used to specify options that can be read only by servers within a given release series.

For backward compatibility, `mysql.server` also reads the `[mysql_server]` group and `mysqld_safe` also reads the `[safe_mysqld]` group. However, you should update your option files to use the `[mysql.server]` and `[mysqld_safe]` groups instead when you begin using MySQL 4.0 or later.

See [Section 4.3.2, “Using Option Files”](#).

### 2.10.2.3. Starting and Troubleshooting the MySQL Server

This section provides troubleshooting suggestions for problems starting the server on Unix. If you are using Windows, see [Section 2.3.13, “Troubleshooting a MySQL Installation Under Windows”](#).

If you have problems starting the server, here are some things to try:

- Check the error log to see why the server does not start.
- Specify any special options needed by the storage engines you are using.
- Make sure that the server knows where to find the data directory.
- Make sure that the server can access the data directory. The ownership and permissions of the data directory and its contents must be set such that the server can read and modify them.
- Verify that the network interfaces the server wants to use are available.

Some storage engines have options that control their behavior. You can create a `my.cnf` file and specify startup options for the engines that you plan to use. If you are going to use storage engines that support transactional tables ([InnoDB](#), [BDB](#), [NDB](#)), be sure that you have them configured the way you want before starting the server:

- If you are using [InnoDB](#) tables, refer to the [InnoDB-specific startup options](#). In MySQL 3.23, you must configure [InnoDB](#) explicitly or the server fails to start. From MySQL 4.0 on, [InnoDB](#) uses default values for its configuration options if you specify none. See [Section 14.2.4, “InnoDB Configuration”](#).
- If you are using [BDB](#) (Berkeley DB) tables, see [Section 14.5.3, “BDB Startup Options”](#).
- If you are using MySQL Cluster, see [Section 15.4, “MySQL Cluster Configuration”](#).

Storage engines will use default option values if you specify none, but it is recommended that you review the available options and specify explicit values for those for which the defaults are not appropriate for your installation.

When the `mysqld` server starts, it changes location to the data directory. This is where it expects to find databases and where it expects to write log files. The server also writes the pid (process ID) file in the data directory.

The data directory location is hardwired in when the server is compiled. This is where the server looks for the data directory by default. If the data directory is located somewhere else on your system, the server will not work properly. You can determine what the default path settings are by invoking `mysqld` with the `--verbose` and `--help` options. (Prior to MySQL 4.1, omit the `--verbose` option.)

If the default locations don't match the MySQL installation layout on your system, you can override them by specifying options to `mysqld` or `mysqld_safe` on the command line or in an option file.

To specify the location of the data directory explicitly, use the `--datadir` option. However, normally you can tell `mysqld` the location of the base directory under which MySQL is installed and it looks for the data directory there. You can do this with the `--basedir` option.

To check the effect of specifying path options, invoke `mysqld` with those options followed by the `--verbose` and `--help` options. For example, if you change location into the directory where `mysqld` is installed and then run the following command, it shows the effect of starting the server with a base directory of `/usr/local`:

```
shell> ./mysqld --basedir=/usr/local --verbose --help
```

You can specify other options such as `--datadir` as well, but `--verbose` and `--help` must be the last options. (Prior to MySQL 4.1, omit the `--verbose` option.)

Once you determine the path settings you want, start the server without `--verbose` and `--help`.

If `mysqld` is currently running, you can find out what path settings it is using by executing this command:

```
shell> mysqladmin variables
```

Or:

```
shell> mysqladmin -h host_name variables
```

`host_name` is the name of the MySQL server host.

If you get `Errcode 13` (which means `Permission denied`) when starting `mysqld`, this means that the privileges of the data directory or its contents do not allow the server access. In this case, you change the permissions for the involved files and directories so that the server has the right to use them. You can also start the server as `root`, but this raises security issues and should be avoided.

On Unix, change location into the data directory and check the ownership of the data directory and its contents to make sure the server has access. For example, if the data directory is `/usr/local/mysql/var`, use this command:

```
shell> ls -la /usr/local/mysql/var
```

If the data directory or its files or subdirectories are not owned by the login account that you use for running the server, change their ownership to that account. If the account is named `mysql`, use these commands:

```
shell> chown -R mysql /usr/local/mysql/var
shell> chgrp -R mysql /usr/local/mysql/var
```

If the server fails to start up correctly, check the error log. Log files are located in the data directory (typically `C:\mysql\data` on Windows, `/usr/local/mysql/data` for a Unix binary distribution, and `/usr/local/var` for a Unix source distribution). Look in the data directory for files with names of the form `host_name.err` and `host_name.log`, where `host_name` is the name of your server host. (Older servers on Windows use `mysql.err` as the error log name.) Then check the last few lines of these files. On Unix, you can use `tail` to display the last few lines:

```
shell> tail host_name.err
shell> tail host_name.log
```

The error log should contain information that indicates why the server couldn't start. For example, you might see something like this in the log:

```
000729 14:50:10 bdb: Recovery function for LSN 1 27595 failed
000729 14:50:10 bdb: warning: ./test/t1.db: No such file or directory
000729 14:50:10 Can't init databases
```

This means that you did not start `mysqld` with the `--bdb-no-recover` option and Berkeley DB found something wrong with its own log files when it tried to recover your databases. To be able to continue, you should move away the old Berkeley DB log files from the database directory to some other place, where you can later examine them. The BDB log files are named in sequence beginning with



`log.0000000001`, where the number increases over time.

If you are running `mysqld` with BDB table support and `mysqld` dumps core at startup, this could be due to problems with the BDB recovery log. In this case, you can try starting `mysqld` with `--bdb-no-recover`. If that helps, you should remove all BDB log files from the data directory and try starting `mysqld` again without the `--bdb-no-recover` option.

If either of the following errors occur, it means that some other program (perhaps another `mysqld` server) is using the TCP/IP port or Unix socket file that `mysqld` is trying to use:

```
Can't start server: Bind on TCP/IP port: Address already in use
Can't start server: Bind on unix socket...
```

Use `ps` to determine whether you have another `mysqld` server running. If so, shut down the server before starting `mysqld` again. (If another server is running, and you really want to run multiple servers, you can find information about how to do so in [Section 5.12, “Running Multiple MySQL Servers on the Same Machine”](#).)

If no other server is running, try to execute the command `telnet your_host_name tcp_ip_port_number`. (The default MySQL port number is 3306.) Then press Enter a couple of times. If you don't get an error message like `telnet: Unable to connect to remote host: Connection refused`, some other program is using the TCP/IP port that `mysqld` is trying to use. You'll need to track down what program this is and disable it, or else to tell `mysqld` to listen on a different port with the `--port` option. In this case, you also need to specify the port number for client programs when connecting to the server via TCP/IP.

Another reason the port might be inaccessible is that you have a firewall running that blocks connections to it. If so, modify the firewall settings to allow access to the port.

If the server starts but you cannot connect to it, you should make sure that you have an entry in `/etc/hosts` that looks like this:

```
127.0.0.1 localhost
```

This problem occurs only on systems that do not have a working thread library and for which MySQL must be configured to use MIT-pthreads.

If you cannot get `mysqld` to start, you can try to make a trace file to find the problem by using the `--debug` option. See [Section E.1.2, “Creating Trace Files”](#).

### 2.10.3. Securing the Initial MySQL Accounts

Part of the MySQL installation process is to set up the `mysql` database that contains the grant tables:

- Windows distributions contain preinitialized grant tables that are installed automatically.
- On Unix, the grant tables are populated by the `mysql_install_db` program. Some installation methods run this program for you. Others require that you execute it manually. For details, see [Section 2.10.2, “Unix Post-Installation Procedures”](#).

The grant tables define the initial MySQL user accounts and their access privileges. These accounts are set up as follows:

- Accounts are created with the username `root`. These are superuser accounts that can do anything. The initial `root` account passwords are empty, so anyone can connect to the MySQL server as `root` without a password and be granted all privileges.

- On Windows, prior to MySQL 4.1.10, two `root` accounts are created; one of these is for connecting from the local host and the other allows connections from any host. Beginning with MySQL 4.1.10, the Windows installer creates only one `root` account, which can connect from the local machine only. The Windows installer will optionally create an account allowing for connections from any host only if the user selects the `Enable root access from remote machines` option during installation.
- On Unix, both `root` accounts are for connections from the local host. Connections must be made from the local host by specifying a hostname of `localhost` for one of the accounts, or the actual hostname or IP number for the other.
- Two anonymous-user accounts are created, each with an empty username. The anonymous accounts have no password, so anyone can use them to connect to the MySQL server.
  - On Windows, one anonymous account is for connections from the local host. It has all privileges, just like the `root` accounts. The other is for connections from any host and has all privileges for the `test` database and for other databases with names that start with `test`.
  - On Unix, both anonymous accounts are for connections from the local host. Connections must be made from the local host by specifying a hostname of `localhost` for one of the accounts, or the actual hostname or IP number for the other. These accounts have all privileges for the `test` database and for other databases with names that start with `test_`.

As noted, none of the initial accounts have passwords. This means that your MySQL installation is unprotected until you do something about it:

- If you want to prevent clients from connecting as anonymous users without a password, you should either assign a password to each anonymous account or else remove the accounts.
- You should assign a password to each MySQL `root` account.

The following instructions describe how to set up passwords for the initial MySQL accounts, first for the anonymous accounts and then for the `root` accounts. Replace `newpwd` in the examples with the actual password that you want to use. The instructions also cover how to remove the anonymous accounts, should you prefer not to allow anonymous access at all.

You might want to defer setting the passwords until later, so that you don't need to specify them while you perform additional setup or testing. However, be sure to set them before using your installation for production purposes.

### Anonymous Account Password Assignment

To assign passwords to the anonymous accounts, connect to the server as `root` and then use either `SET PASSWORD` or `UPDATE`. In either case, be sure to encrypt the password using the `PASSWORD()` function.

To use `SET PASSWORD` on Windows, do this:

```
shell> mysql -u root
mysql> SET PASSWORD FOR '@'localhost' = PASSWORD('newpwd');
mysql> SET PASSWORD FOR '@%' = PASSWORD('newpwd');
```

To use `SET PASSWORD` on Unix, do this:

```
shell> mysql -u root
mysql> SET PASSWORD FOR '@'localhost' = PASSWORD('newpwd');
```

```
mysql> SET PASSWORD FOR '@'host_name' = PASSWORD('newpwd');
```

In the second `SET PASSWORD` statement, replace `host_name` with the name of the server host. This is the name that is specified in the `Host` column of the non-`localhost` record for `root` in the `user` table. If you do not know what hostname this is, issue the following statement before using `SET PASSWORD`:

```
mysql> SELECT Host, User FROM mysql.user;
```

Look for the record that has `root` in the `User` column and something other than `localhost` in the `Host` column. Then use that `Host` value in the second `SET PASSWORD` statement.

The other way to assign passwords to the anonymous accounts is by using `UPDATE` to modify the `user` table directly. Connect to the server as `root` and issue an `UPDATE` statement that assigns a value to the `Password` column of the appropriate `user` table records. The procedure is the same for Windows and Unix. The following `UPDATE` statement assigns a password to both anonymous accounts at once:

```
shell> mysql -u root
mysql> UPDATE mysql.user SET Password = PASSWORD('newpwd')
-> WHERE User = '';
mysql> FLUSH PRIVILEGES;
```

After you update the passwords in the `user` table directly using `UPDATE`, you must tell the server to re-read the grant tables with `FLUSH PRIVILEGES`. Otherwise, the change goes unnoticed until you restart the server.

### Anonymous Account Removal

If you prefer to remove the anonymous accounts instead, do so as follows:

```
shell> mysql -u root
mysql> DELETE FROM mysql.user WHERE User = '';
mysql> FLUSH PRIVILEGES;
```

The `DELETE` statement applies both to Windows and to Unix. On Windows, if you want to remove only the anonymous account that has the same privileges as `root`, do this instead:

```
shell> mysql -u root
mysql> DELETE FROM mysql.user WHERE Host='localhost' AND User='';
mysql> FLUSH PRIVILEGES;
```

That account allows anonymous access but has full privileges, so removing it improves security.

### root Account Password Assignment

You can assign passwords to the `root` accounts in several ways. The following discussion demonstrates three methods:

- Use the `SET PASSWORD` statement
- Use the `mysqladmin` command-line client program
- Use the `UPDATE` statement

To assign passwords using `SET PASSWORD`, connect to the server as `root` and issue two `SET PASSWORD` statements. Be sure to encrypt the password using the `PASSWORD( )` function.

For Windows, do this:

```
shell> mysql -u root
mysql> SET PASSWORD FOR 'root'@'localhost' = PASSWORD('newpwd');
mysql> SET PASSWORD FOR 'root'@'%' = PASSWORD('newpwd');
```

For Unix, do this:

```
shell> mysql -u root
mysql> SET PASSWORD FOR 'root'@'localhost' = PASSWORD('newpwd');
mysql> SET PASSWORD FOR 'root'@'host_name' = PASSWORD('newpwd');
```

In the second `SET PASSWORD` statement, replace `host_name` with the name of the server host. This is the same hostname that you used when you assigned the anonymous account passwords.

To assign passwords to the `root` accounts using `mysqladmin`, execute the following commands:

```
shell> mysqladmin -u root password "newpwd"
shell> mysqladmin -u root -h host_name password "newpwd"
```

These commands apply both to Windows and to Unix. In the second command, replace `host_name` with the name of the server host. The double quotes around the password are not always necessary, but you should use them if the password contains spaces or other characters that are special to your command interpreter.

If you are using a server from a *very* old version of MySQL, the `mysqladmin` commands to set the password fail with the message `parse error near 'SET password'`. The solution to this problem is to upgrade the server to a newer version of MySQL.

You can also use `UPDATE` to modify the `user` table directly. The following `UPDATE` statement assigns a password to both `root` accounts at once:

```
shell> mysql -u root
mysql> UPDATE mysql.user SET Password = PASSWORD('newpwd')
-> WHERE User = 'root';
mysql> FLUSH PRIVILEGES;
```

The `UPDATE` statement applies both to Windows and to Unix.

After the passwords have been set, you must supply the appropriate password whenever you connect to the server. For example, if you want to use `mysqladmin` to shut down the server, you can do so using this command:

```
shell> mysqladmin -u root -p shutdown
Enter password: (enter root password here)
```

**Note:** If you forget your `root` password after setting it up, [Section A.4.1, “How to Reset the Root Password”](#), covers the procedure for resetting it.

To set up additional accounts, you can use the `GRANT` statement. For instructions, see [Section 5.8.2, “Adding New User Accounts to MySQL”](#).

## 2.11. Upgrading MySQL

As a general rule, we recommend that when upgrading from one release series to another, you should go to the next series rather than skipping a series. For example, if you currently are running MySQL 3.23 and wish to upgrade to a newer series, upgrade to MySQL 4.0 rather than to 4.1 or 5.0.

The following items form a checklist of things that you should do whenever you perform an upgrade:

- Read the upgrading section for the release series to which you are upgrading. Read the change notes as well. These provide information about new features you can use.
- Before you perform an upgrade, back up your databases, including the `mysql` database that contains the grant tables.
- Some releases of MySQL introduce changes to the structure of the grant tables to add new privileges or features. Whenever you update to a new version of MySQL, you should update your grant tables to make sure that they have the current structure so that you can take advantage of any new capabilities. See [Section 5.5.1, “mysql\\_fix\\_privilege\\_tables — Upgrade MySQL System Tables”](#).
- If you are running MySQL Server on Windows, see [Section 2.3.14, “Upgrading MySQL on Windows”](#).
- If you are using replication, see [Section 6.6, “Upgrading a Replication Setup”](#), for information on upgrading your replication setup.
- If you previously installed a MySQL-Max distribution that includes a server named `mysqld-max`, and then upgrade later to a non-Max version of MySQL, `mysqld_safe` still attempts to run the old `mysqld-max` server. If you perform such an upgrade, you should remove the old `mysqld-max` server manually to ensure that `mysqld_safe` runs the new `mysqld` server.

You can always move the MySQL format files and data files between different versions on the same architecture as long as you stay within versions for the same release series of MySQL. If you change the character set when running MySQL, you must run `myisamchk -r -q --set-collation=collation_name` on all `MyISAM` tables. Otherwise, your indexes may not be ordered correctly, because changing the character set may also change the sort order. (Use `--set-character-set` before MySQL 4.1.1.)

Normally, you can upgrade MySQL to a newer MySQL version without having to do any changes to your tables. Please confirm whether the upgrade notes to the particular version you are upgrading to tell you anything about this. If there would be any incompatibilities you can use `mysqldump` to dump your tables before upgrading. After upgrading, reload the dump file using `mysql` or `mysqlimport` to re-create your tables.

If you are cautious about using new versions, you can always rename your old `mysqld` before installing a newer one. For example, if you are using MySQL 4.0.18 and want to upgrade to 4.1.1, rename your current server from `mysqld` to `mysqld-4.0.18`. If your new `mysqld` then does something unexpected, you can simply shut it down and restart with your old `mysqld`.

If, after an upgrade, you experience problems with recompiled client programs, such as [Commands out of sync](#) or unexpected core dumps, you probably have used old header or library files when compiling your programs. In this case, you should check the date for your `mysql.h` file and `libmysqlclient.a` library to verify that they are from the new MySQL distribution. If not, recompile your programs with the new headers and libraries.

If problems occur, such as that the new `mysqld` server does not want to start or that you cannot connect without a password, verify that you do not have some old `my.cnf` file from your previous installation. You can check this with the `--print-defaults` option (for example, `mysqld --print-defaults`). If this command displays anything other than the program name, you have an active `my.cnf` file that affects server or client operation.

It is a good idea to rebuild and reinstall the Perl `DBD:mysql` module whenever you install a new release of MySQL. The same applies to other MySQL interfaces as well, such as the PHP `mysql` and `mysqli` extensions or the Python `MySQLdb` module.

## 2.11.1. Upgrading from MySQL 4.0 to 4.1

**Note:** It is good practice to back up your data before installing any new version of software. Although MySQL works very hard to ensure a high level of quality, you should protect your data by making a backup. MySQL generally recommends that you dump and reload your tables from any previous version to upgrade to 4.1.

In general, you should do the following when upgrading from MySQL 4.0 to 4.1:

- Check the items in the change lists found later in this section to see whether any of them might affect your applications. Note particularly any that are marked **Incompatible change**. These result in incompatibilities with earlier versions of MySQL and you should consider the implications of these incompatibilities *before you upgrade*. Note particularly the items under “Server Changes” that related to changes in character set support.
- Read the 4.1 news items to see what significant new features you can use in 4.1. See [Section D.1, “Changes in release 4.1.x \(Production\)”](#).
- If you are running MySQL Server on Windows, see [Section 2.3.14, “Upgrading MySQL on Windows”](#). You should also be aware that two of the Windows MySQL servers were renamed in MySQL 4.1. See [Section 2.3.8, “Selecting a MySQL Server type”](#).
- After upgrading, update the grant tables to obtain the new longer `Password` column that is needed for more secure handling of passwords. The procedure uses `mysql_fix_privilege_tables` and is described in [Section 5.5.1, “mysql\\_fix\\_privilege\\_tables — Upgrade MySQL System Tables”](#). If you do not do this, MySQL does not use the new more secure protocol to authenticate. Implications of the password-handling change for applications are given later in this section.
- If you are using replication, see [Section 6.6, “Upgrading a Replication Setup”](#), for information on upgrading your replication setup.
- The Berkeley DB table handler is updated to DB 4.1 (from 3.2) which has a new log format. If you have to downgrade back to 4.0 you must use `mysqldump` to dump your BDB tables in text format and delete all `log.XXXXXXXXXX` files before you start MySQL 4.0 and reload the data.
- MySQL 4.1.3 introduces support for per-connection time zones. See [Section 5.10.8, “MySQL Server Time Zone Support”](#). To enable recognition of named time zones, you should create the time zone tables in the `mysql` database. For instructions, see [Section 2.10, “Post-Installation Setup and Testing”](#).
- If you are using an old `DBD-mysql` module (`Msql-MySQL-modules`) you must upgrade to the newer `DBD-mysql` module. Anything above `DBD-mysql 2.xx` should be satisfactory.

If you do not upgrade, some methods (such as `DBI->do( )`) do not notice error conditions correctly.

- The `--defaults-file=option-file-name` option gives an error if the option file does not exist.
- Some notes about upgrading from MySQL 4.0 to MySQL 4.1 on Netware: Make sure to upgrade Perl and PHP versions. Download Perl 5 for Netware from <http://forge.novell.com/modules/xfmod/project/?perl5> and PHP from <http://forge.novell.com/modules/xfmod/project/?php>. Download and install the Perl module for MySQL 4.1 from [http://forge.novell.com/modules/xfmod/project/showfiles.php?group\\_id=1126](http://forge.novell.com/modules/xfmod/project/showfiles.php?group_id=1126) and the PHP Extension for MySQL 4.1 from [http://forge.novell.com/modules/xfmod/project/showfiles.php?group\\_id=1078](http://forge.novell.com/modules/xfmod/project/showfiles.php?group_id=1078).

Several visible behaviors have changed between MySQL 4.0 and MySQL 4.1 to fix some critical bugs and make MySQL more compatible with standard SQL. These changes may affect your applications.

Some of the 4.1 behaviors can be tested in 4.0 before performing a full upgrade to 4.1. We have added to later MySQL 4.0 releases (from 4.0.12 on) a `--new` startup option for `mysqld`. See [Section 5.2.1](#), “Command Options”.

This option gives you the 4.1 behavior for the most critical changes. You can also enable these behaviors for a given client connection with the `SET @@new=1` command, or turn them off if they are on with `SET @@new=0`.

If you believe that some of the 4.1 changes affect you, we recommend that before upgrading to 4.1, you download the latest MySQL 4.0 version and run it with the `--new` option by adding the following to your config file:

```
[mysqld-4.0]
new
```

That way you can test the new behaviors in 4.0 to make sure that your applications work with them. This helps you have a smooth, painless transition when you perform a full upgrade to 4.1 later. Putting the `--new` option in the `[mysqld-4.0]` option group ensures that you do not accidentally later run the 4.1 version with the `--new` option.

The following lists describe changes that may affect applications and that you should watch out for when upgrading to version 4.1.

#### Server Changes:

The most notable change is that character set support has been improved. The server supports multiple character sets, and all tables and non-binary string columns (`CHAR`, `VARCHAR`, and `TEXT`) have a character set. See [Chapter 10, Character Set Support](#). Binary string columns (`BINARY`, `VARBINARY`, and `BLOB`) contain strings of bytes and do not have a character set. **Note:** This change in character set support results in the potential for table damage if you do not upgrade properly, so consider carefully the incompatibilities noted here.

- **Incompatible change:** There are conditions under which you should rebuild tables. In general, to rebuild a table, dump it with `mysqldump` and reload the dump file. Some items in the following list indicate alternatives means for rebuilding.
  - If you have created or used `InnoDB` tables with `TIMESTAMP` columns in MySQL versions 4.1.0 to 4.1.3, you must rebuild those tables when you upgrade to MySQL 4.1.4 or later. The storage format in those MySQL versions for `TIMESTAMP` columns was incorrect. If you upgrade from MySQL 4.0 to 4.1.4 or later, no rebuild of tables with `TIMESTAMP` columns is needed.
  - Starting from MySQL 4.1.3, `InnoDB` uses the same character set comparison functions as MySQL for non-`latin1_swedish_ci` character strings that are not `BINARY`. This changes the sorting order of space and characters with a code < ASCII(32) in those character sets. For `latin1_swedish_ci` character strings and `BINARY` strings, `InnoDB` uses its own pad-spaces-at-end comparison method, which stays unchanged. Note that `latin1_swedish_ci` is the default collation order for `latin1` in 4.0. If you have an `InnoDB` table created with MySQL 4.1.2 or earlier, with an index on a non-`latin1_swedish_ci` character set and collation order column that is not `BINARY` (in the case of 4.1.0 and 4.1.1, with any character set and collation), and that column may contain characters with a code < ASCII(32), you should do `ALTER TABLE` or `OPTIMIZE TABLE` on it to regenerate the index, after upgrading to MySQL 4.1.3 or later. You can also rebuild the table from a dump.

`MyISAM` tables also have to be rebuilt or repaired in these cases. You can use `mysqldump` to dump them in 4.0 and then reload them in 4.1. An alternative is to use `OPTIMIZE TABLE` after upgrading, but this *must* be done before any updates are made in 4.1.

- If you have used column prefix indexes on UTF-8 columns or other multi-byte character set columns in MySQL 4.1.0 to 4.1.5, you must rebuild the tables when you upgrade to MySQL 4.1.6 or later.
- If you have used accent characters (characters with byte values of 128 to 255) in database names, table names, constraint names, or column names in versions of MySQL earlier than 4.1, you cannot upgrade to MySQL 4.1 directly, because 4.1 uses UTF-8 to store metadata names. Use `RENAME TABLE` to overcome this if the accent character is in the table name or the database name, or rebuild the table.
- String comparison works according to SQL standard: Instead of stripping end spaces before comparison, we now extend the shorter string with spaces. The problem with this is that now `'a' > 'a\t'`, which it was not before. If you have any tables where you have an indexed `CHAR`, `VARCHAR` or `TEXT` column in which the last character in the index may be less than `ASCII(32)`, you should use `REPAIR TABLE` or `mysqlcheck` to ensure that the table is correct.
- `MyISAM` tables now use an improved checksum algorithm in MySQL 4.1. If you have `MyISAM` tables with live checksum enabled (you used `CHECKSUM=1` in `CREATE TABLE` or `ALTER TABLE`), these tables appear to be corrupted following an upgrade. Use `REPAIR TABLE` to recalculate the checksum for each such table.
- **Incompatible change:** MySQL interprets length specifications in character column definitions in characters. (Earlier versions interpret them in bytes.) For example, `CHAR(N)` means `N` characters, not `N` bytes.

For single-byte character sets, this change makes no difference. However, if you upgrade to MySQL 4.1 and configure the server to use a multi-byte character set, the apparent length of character columns changes. Suppose that a 4.0 table contains a `CHAR(8)` column used to store `ujis` characters. Eight bytes can store from two to four `ujis` characters. If you upgrade to 4.1 and configure the server to use `ujis` as its default character set, the server interprets character column lengths based on the maximum size of a `ujis` character, which is three bytes. The number of three-byte characters that fit in eight bytes is two. Consequently, if you use `SHOW CREATE TABLE` to view the table definition, MySQL displays `CHAR(2)`. You can retrieve existing data from the table, but you can only store new values containing up to two characters. To correct this issue, use `ALTER TABLE` to change the column definition. For example:

```
ALTER TABLE tbl_name MODIFY col_name CHAR(8);
```

- **Warning: Incompatible change:** As of MySQL 4.1.2, handling of the `FLOAT` and `DOUBLE` floating-point data types is more strict to follow standard SQL. For example, a data type of `FLOAT(3,1)` stores a maximum value of 99.9. Before 4.1.2, the server allowed larger numbers to be stored. That is, it stored a value such as 100.0 as 100.0. As of 4.1.2, the server clips 100.0 to the maximum allowable value of 99.9. If you have tables that were created before MySQL 4.1.2 and that contain floating-point data not strictly legal for the data type, you should alter the data types of those columns. For example:

```
ALTER TABLE tbl_name MODIFY col_name FLOAT(4,1);
```

- **Incompatible change:** In connection with the support for per-connection time zones in MySQL 4.1.3, the `timezone` system variable was renamed to `system_time_zone`.
- **Important note:** MySQL 4.1 stores table names and column names in `utf8`. If you have table names or column names that use characters outside of the standard 7-bit US-ASCII range, you may have to do a `mysqldump` of your tables in MySQL 4.0 and restore them after upgrading to MySQL 4.1. The symptom for this problem is that you get a `table not found` error when trying to access your tables. In this case, you should be able to downgrade back to MySQL 4.0 and access your data.



- **Important note:** If you upgrade to MySQL 4.1.1 or higher, it is difficult to downgrade back to 4.0 or 4.1.0. That is because, for earlier versions, [InnoDB](#) is not aware of multiple tablespaces.
- All tables and non-binary string columns ([CHAR](#), [VARCHAR](#), and [TEXT](#)) have a character set. See [Chapter 10, Character Set Support](#). Binary string columns ([BINARY](#), [VARBINARY](#), and [BLOB](#)) contain strings of bytes and do not have a character set.

Character set information is displayed by `SHOW CREATE TABLE` and `mysqldump`. (MySQL versions 4.0.6 and above can read the new dump files; older versions cannot.) This change should not affect applications that use only one character set.

- If you were using columns with the [CHAR BINARY](#) or [VARCHAR BINARY](#) data types in MySQL 4.0, these were treated as binary strings. To have them treated as binary strings in MySQL 4.1, you should convert them to the [BINARY](#) and [VARBINARY](#) data types, respectively.
- If you have table columns that store character data represented in a character set that the 4.1 server supports directly, you can convert the columns to the proper character set using the instructions in [Section 10.9.2, “Converting 4.0 Character Columns to 4.1 Format”](#). Also, database, table, and column identifiers are stored internally using Unicode (UTF-8) regardless of the default character set. See [Section 9.2, “Database, Table, Index, Column, and Alias Names”](#).
- The table definition format used in `.frm` files has changed slightly in 4.1. MySQL 4.0 versions from 4.0.11 on can read the new `.frm` format directly, but older versions cannot. If you need to move tables from 4.1 to a version earlier than 4.0.11, you should use `mysqldump`. See [Section 8.12, “mysqldump — A Database Backup Program”](#).
- Windows servers support connections from local clients using shared memory if run with the `--shared-memory` option. If you are running multiple servers this way on the same Windows machine, you should use a different `--shared-memory-base-name` option for each server.
- **Incompatible change:** The interface to aggregate user-defined functions changed as of MySQL 4.1.1. You must declare a `xxx_clear()` function for each aggregate function `xxx()`. `xxx_clear()` is used instead of `xxx_reset()`. See [Section 19.2.4.2, “UDF Calling Sequences for Aggregate Functions”](#).
- As of MySQL 4.1.10a, the server by default no longer loads user-defined functions (UDFs) unless they have at least one auxiliary symbol defined in addition to the main function symbol. This behavior can be overridden with the `--allow-suspicious-udfs` option. See [Section 19.2.4.6, “User-Defined Function Security Precautions”](#).

#### Client Changes:

- `mysqldump` has the `--opt` and `--quote-names` options enabled by default. You can turn these off using `--skip-opt` and `--skip-quote-names`.

#### SQL Changes:

- **Incompatible change:** In MySQL 4.1, string comparison works according to the SQL standard: Instead of stripping end spaces before comparison, the shorter string is extended using spaces. This means that `'a' > 'a\t'`, which it was not previously. If you have any tables containing an indexed [CHAR](#), [VARCHAR](#) or [TEXT](#) column in which the last character in the index may be less than `ASCII(32)`, you should use `REPAIR TABLE` or `mysqlcheck` to ensure that the table is correct.
- **Incompatible change:** `TIMESTAMP` is returned in MySQL 4.1 as a string in `'YYYY-MM-DD HH:MM:SS'` format. (See [Section 11.3.1.2, “TIMESTAMP Properties as of MySQL 4.1”](#).) From

4.0.12 on, the `--new` option can be used to make a 4.0 server behave as 4.1 in this respect. The effect of this option is described in [Section 11.3.1.1, “TIMESTAMP Properties Prior to MySQL 4.1”](#).

When running the server with `--new`, if you want to have a `TIMESTAMP` column returned as a number (as MySQL 4.0 does by default), you should add `+0` when you retrieve it:

```
mysql> SELECT ts_col + 0 FROM tbl_name;
```

Display widths for `TIMESTAMP` columns are no longer supported in MySQL 4.1. For example, if you declare a column as `TIMESTAMP(10)`, the `(10)` is ignored.

- **Incompatible change:** Binary values such as `0xFFDF` are assumed to be strings instead of numbers. This fixes some problems with character sets where it is convenient to input a string as a binary value. With this change, you should use `CAST()` if you want to compare binary values numerically as integers:

```
mysql> SELECT CAST(0xFEFF AS UNSIGNED INTEGER)
->      < CAST(0xFF AS UNSIGNED INTEGER);
-> 0
```

If you do not use `CAST()`, a lexical string comparison is made instead:

```
mysql> SELECT 0xFEFF < 0xFF;
-> 1
```

Using binary items in a numeric context or comparing them using the `=` operator should work as before. (The `--new` option can be used from 4.0.13 on to make a 4.0 server behave as 4.1 in this respect.)

- **Incompatible change:** Before MySQL 4.1.13, conversion of `DATETIME` values to numeric form by adding zero produced a result in `YYYYMMDDHHMMSS` format. The result of `DATETIME+0` is now in `YYYYMMDDHHMMSS.000000` format.
- **Incompatible change:** In MySQL 4.1.12, the behavior of `LOAD DATA INFILE` and `SELECT ... INTO OUTFILE` has changed when the `FIELDS TERMINATED BY` and `FIELDS ENCLOSED BY` values both are empty. Formerly, a column was read or written the display width of the column. For example, `INT(4)` was read or written using a field with a width of 4. Now columns are read and written using a field width wide enough to hold all values in the field. However, data files written before this change was made might not be reloaded correctly with `LOAD DATA INFILE` for MySQL 4.1.12 and up. This change also affects data files read by `mysqlimport` and written by `mysqldump --tab`, which use `LOAD DATA INFILE` and `SELECT ... INTO OUTFILE`. For more information, see [Section 13.2.5, “LOAD DATA INFILE Syntax”](#).
- **Incompatible change:** Before MySQL 4.1.1, the statement parser was less strict and its string-to-date conversion would ignore everything up to the first digit. As a result, invalid statements such as the following were accepted:

```
INSERT INTO t (datetime_col) VALUES ('stuff 2005-02-11 10:17:01');
```

As of MySQL 4.1.1, the parser is stricter and treats the string as an invalid date, so the preceding statement results in a warning.

- **Incompatible change:** In MySQL 4.1.2, the `Type` column in the output from `SHOW TABLE STATUS` was renamed to `Engine`. This affects applications that identify output columns by name rather than by position.
- **Incompatible change:** The syntax for multiple-table `DELETE` statements that use table aliases changed between MySQL 4.0 and 4.1. In MySQL 4.0, you should use the true table name to refer to

any table from which rows should be deleted:

```
DELETE test FROM test AS t1, test2 WHERE ...
```

In MySQL 4.1, you must use the alias:

```
DELETE t1 FROM test AS t1, test2 WHERE ...
```

We did not make this change in 4.0 to avoid breaking any old 4.0 applications that were using the old syntax. However, if you use such `DELETE` statements and are using replication, the change in syntax means that a 4.0 master cannot replicate to 4.1 (or higher) slaves.

- Some keywords are reserved in MySQL 4.1 that were not reserved in MySQL 4.0. See [Section 9.5, “Treatment of Reserved Words in MySQL”](#).
- The `LOAD DATA FROM MASTER` and `LOAD TABLE FROM MASTER` statements are deprecated. See [Section 13.6.2.2, “LOAD DATA FROM MASTER Syntax”](#), for recommended alternatives.
- When using multiple-table `DELETE` statements, you should use the alias of the tables from which you want to delete, not the actual table name. For example, instead of doing this:

```
DELETE test FROM test AS t1, test2 WHERE ...
```

Do this:

```
DELETE t1 FROM test AS t1, test2 WHERE ...
```

This corrects a problem that was present in MySQL 4.0.

- For functions that produce a `DATE`, `DATETIME`, or `TIME` value, the result returned to the client is fixed up to have a temporal type. For example, in MySQL 4.1, you obtain the following:

```
mysql> SELECT CAST('2001-1-1' AS DATETIME);  
-> '2001-01-01 00:00:00'
```

In MySQL 4.0, the result of the statement is different:

```
mysql> SELECT CAST('2001-1-1' AS DATETIME);  
-> '2001-01-01'
```

- `DEFAULT` values no longer can be specified for `AUTO_INCREMENT` columns. (In 4.0, a `DEFAULT` value is silently ignored; in 4.1, an error occurs.)
- `LIMIT` no longer accepts negative arguments. Use some large number (maximum 18446744073709551615) instead of -1.
- `SERIALIZE` is no longer a valid mode value for the `sql_mode` variable. You should use `SET TRANSACTION ISOLATION LEVEL SERIALIZABLE` instead. `SERIALIZE` is no longer valid for the `--sql-mode` option for `mysqld`, either. Use `--transaction-isolation=SERIALIZABLE` instead.
- A new startup option named `innodb_table_locks` was added that causes `LOCK TABLE` to also acquire InnoDB table locks. This option is enabled by default. This can cause deadlocks in applications that use `AUTOCOMMIT=1` and `LOCK TABLES`. If your application encounters deadlocks after upgrading, you may need to add `innodb_table_locks=0` to your `my.cnf` file.

**C API Changes:**

- **Incompatible change:** The `mysql_shutdown()` C API function has an extra parameter as of MySQL 4.1.3: `SHUTDOWN-level`. You should convert any `mysql_shutdown(X)` call you have in your application to `mysql_shutdown(X, SHUTDOWN_DEFAULT)`. Any third-party API that links against the C API library must be modified to account for this change or it will not compile.
- Some C API calls such as `mysql_real_query()` return `1` on error, not `-1`. You may have to change some old applications if they use constructs like this:

```
if (mysql_real_query(mysql_object, query, query_length) == -1)
{
    printf("Got error");
}
```

Change the call to test for a non-zero value instead:

```
if (mysql_real_query(mysql_object, query, query_length) != 0)
{
    printf("Got error");
}
```

**Password-Handling Changes:**

The password hashing mechanism changed in 4.1 to provide better security; this may cause compatibility problems if you have clients using the client library from 4.0 or earlier. (It is very likely that you have 4.0 clients in situations where clients connect from remote hosts that have not yet upgraded to 4.1.) The following list indicates some possible upgrade strategies. They represent various tradeoffs between the goals of compatibility with old clients and security.

- Only upgrade the client to use 4.1 client libraries (not the server). No behavior changes (except the return value of some API calls), but you cannot use any of the new features provided by the 4.1 client/server protocol, either. (MySQL 4.1 has an extended client/server protocol that offers such features as prepared statements and multiple result sets.) See [Section 17.2.4, “C API Prepared Statements”](#).
- Upgrade to 4.1 and run the `mysql_fix_privilege_tables` script to widen the `Password` column in the `user` table so that it can hold long password hashes. However — to provide backward compatibility allowing pre-4.1 clients to continue connecting to their short-hash accounts — run the server with the `--old-passwords` option. Eventually, when all your clients are upgraded to 4.1, you can stop using the `--old-passwords` server option. You can also change the passwords for your MySQL accounts to use the new more secure format. A 4.1 installation using only the improved authentication protocol is the most secure one.

Further background on password hashing with respect to client authentication and password-changing operations may be found in [Section 5.7.9, “Password Hashing as of MySQL 4.1”](#), and [Section A.2.3, “Client does not support authentication protocol”](#).

## 2.11.2. Upgrading from MySQL 3.23 to 4.0

In general, you should do the following when upgrading from MySQL 3.23 to 4.0:

- Check the items in the change list found later in this section to see whether any of them might affect your applications. Note particularly any that are marked **Incompatible change**; these result in in-

compatibilities with earlier versions of MySQL.

- Read the 4.0 changelog to see what significant new features you can use in 4.0. See [Section D.2, “Changes in release 4.0.x \(Recent; still supported\)”](#).
- If you are running MySQL Server on Windows, see [Section 2.3.14, “Upgrading MySQL on Windows”](#).
- After upgrading, update the grant tables to add new privileges and features. This procedure uses the `mysql_fix_privilege_tables` script and is described in [Section 5.5.1, “mysql\\_fix\\_privilege\\_tables — Upgrade MySQL System Tables”](#).
- If you are using replication, see [Section 6.6, “Upgrading a Replication Setup”](#), for information on upgrading your replication setup.
- Edit any MySQL startup scripts or option files so that they do not use any of the options described as deprecated later in this section.
- Convert your old `ISAM` tables to `MyISAM` format. One way to do this is with the `mysql_convert_table_format` script. (This is a Perl script; it requires that `DBI` be installed.) To convert all of the tables in a given database, use this command:

```
shell> mysql_convert_table_format database db_name
```

Note that the above command should be used only if *all* tables in the database are `ISAM` or `MyISAM` tables. To avoid converting tables of other types to `MyISAM`, you can explicitly list the names of the `ISAM` tables following the database name on the command line.

Individual tables can be changed to `MyISAM` by using the following `ALTER TABLE` statement for each table to be converted:

```
mysql> ALTER TABLE tbl_name TYPE=MyISAM;
```

If you are not sure of the storage engine for a given table, use this statement:

```
mysql> SHOW TABLE STATUS LIKE 'tbl_name';
```

- Ensure that you do not have any MySQL clients that use shared libraries (like the Perl `DBD:mysql` module). If you do, you should recompile them, because the data structures used in `libmysqlclient.so` have changed. The same applies to other MySQL interfaces such as the Python `MySQLdb` module.

MySQL 4.0 works even if you do not perform the preceding actions, but you cannot use the new security privileges in MySQL 4.0 and you may run into problems when upgrading later to MySQL 4.1 or newer. The `ISAM` file format still works in MySQL 4.0, but is deprecated and is not compiled in by default as of MySQL 4.1. `MyISAM` tables should be used instead.

Old clients should work with a MySQL 4.0 server without any problems.

Even if you perform the preceding actions, you can still downgrade to MySQL 3.23.52 or newer if you run into problems with the MySQL 4.0 series. In this case, you must use `mysqldump` to dump any tables that use full-text indexes and reload the dump file into the 3.23 server. This is necessary because 4.0 uses an improved format for full-text indexing that is not backward-compatible.

The following lists describe changes that may affect applications and that you should watch out for when upgrading to version 4.0.

**Server Changes:**

- As of MySQL 4.0.24, the server by default no longer loads user-defined functions unless they have at least one auxiliary symbol defined in addition to the main function symbol. This behavior can be overridden with the `--allow-suspicious-udfs` option. See [Section 19.2.4.6, “User-Defined Function Security Precautions”](#).
- MySQL 4.0 has many new privileges in the `mysql.user` table. See [Section 5.7.3, “Privileges Provided by MySQL”](#).

In order for these new privileges to work, you must update the grant tables. The procedure for this is described in [Section 5.5.1, “mysql\\_fix\\_privilege\\_tables — Upgrade MySQL System Tables”](#). Until you do this, all accounts have the `SHOW DATABASES`, `CREATE TEMPORARY TABLES`, and `LOCK TABLES` privileges. `SUPER` and `EXECUTE` privileges take their value from `PROCESS`. `REPLICATION SLAVE` and `REPLICATION CLIENT` take their values from `FILE`.

If you have any scripts that create new MySQL user accounts, you may want to change them to use the new privileges. If you are not using `GRANT` commands in the scripts, this is a good time to change your scripts to use `GRANT` instead of modifying the grant tables directly.

From version 4.0.2 on, the option `--safe-show-database` is deprecated (and no longer does anything). See [Section 5.6.3, “Security-Related mysqld Options”](#).

If you get `Access denied` errors for new users in version 4.0.2 and up, you should check whether you need some of the new grants that you did not need before. In particular, you need `REPLICATION SLAVE` (instead of `FILE`) for new slave servers.

- `safe_mysqld` has been renamed to `mysqld_safe`. For backward compatibility, binary distributions will for some time include `safe_mysqld` as a symlink to `mysqld_safe`.
- `InnoDB` support is included by default in binary distributions. If you build MySQL from source, `InnoDB` is configured in by default. If you do not use `InnoDB` and want to save memory when running a server that has `InnoDB` support enabled, use the `--skip-innodb` server startup option. To compile MySQL without `InnoDB` support, run `configure` with the `--without-innodb` option.
- Values for the startup parameters `myisam_max_extra_sort_file_size` and `myisam_max_extra_sort_file_size` are given in bytes (prior to 4.0.3, they were given in megabytes).
- `mysqld` has the option `--temp-pool` enabled by default because this gives better performance with some operating systems (most notably Linux).
- The `mysqld` startup options `--skip-locking` and `--enable-locking` were renamed to `--skip-external-locking` and `--external-locking`.
- External system locking of `MyISAM/ISAM` files is turned off by default. You can turn this on with `--external-locking`. (However, this is never needed for most users.)
- The following startup variables and options were renamed:

Name in 3.23	Name in 4.0 (and above)
<code>myisam_bulk_insert_tree_size</code>	<code>bulk_insert_buffer_size</code>
<code>query_cache_startup_type</code>	<code>query_cache_type</code>
<code>record_buffer</code>	<code>read_buffer_size</code>
<code>record_rnd_buffer</code>	<code>read_rnd_buffer_size</code>

<code>sort_buffer</code>	<code>sort_buffer_size</code>
<code>warnings</code>	<code>log-warnings</code>
<code>--err-log</code>	<code>--log-error</code> (for <code>mysqld_safe</code> )

The startup options `record_buffer`, `sort_buffer`, and `warnings` still work in MySQL 4.0 but are deprecated.

### SQL Changes:

- Some keywords are reserved in MySQL 4.0 that were not reserved in MySQL 3.23. See [Section 9.5, “Treatment of Reserved Words in MySQL”](#).
- The following SQL variables have been renamed:

Name in 3.23	Name in 4.0 and above
<code>SQL_BIG_TABLES</code>	<code>BIG_TABLES</code>
<code>SQL_LOW_PRIORITY_UPDATES</code>	<code>LOW_PRIORITY_UPDATES</code>
<code>SQL_MAX_JOIN_SIZE</code>	<code>MAX_JOIN_SIZE</code>
<code>SQL_QUERY_CACHE_TYPE</code>	<code>QUERY_CACHE_TYPE</code>

The older names still work in MySQL 4.0 but are deprecated.

- You must use `SET GLOBAL SQL_SLAVE_SKIP_COUNTER=skip_count` instead of `SET SQL_SLAVE_SKIP_COUNTER=skip_count`.
- `SHOW MASTER STATUS` returns an empty set if binary logging is not enabled.
- `SHOW SLAVE STATUS` returns an empty set if the slave is not initialized.
- `SHOW INDEX` has two more columns in 4.0 than in 3.23 (`Null` and `Index_type`).
- The format of `SHOW OPEN TABLES` changed.
- As of MySQL 4.0.11, `ORDER BY col_name DESC` sorts `NULL` values last. In 3.23 and in earlier 4.0 versions, this was not always consistent.
- `CHECK`, `LOCALTIME`, and `LOCALTIMESTAMP` are reserved words.
- `DOUBLE` and `FLOAT` columns honor the `UNSIGNED` flag on storage (previously, `UNSIGNED` was ignored for these columns).
- The result of all bitwise operators (`|`, `&`, `<<`, `>>`, and `~`) is unsigned. This may cause problems if you are using them in a context where you want a signed result. See [Section 12.8, “Cast Functions and Operators”](#).

**Note:** When you use subtraction between integer values where one is of type `UNSIGNED`, the result is unsigned. In other words, before upgrading to MySQL 4.0, you should check your application for cases in which you are subtracting a value from an unsigned entity and want a negative answer or subtracting an unsigned value from an integer column. You can disable this behavior by using the `--sql-mode=NO_UNSIGNED_SUBTRACTION` option when starting `mysqld`. See [Section 5.2.5, “SQL Modes”](#).

- You should use integers to store values in `BIGINT` columns (instead of using strings as in MySQL 3.23). Using strings still works, but using integers is more efficient.
- In MySQL 3.23, `INSERT INTO ... SELECT` always had `IGNORE` enabled. As of 4.0.1, MySQL stops (and possibly rolls back) by default in case of an error unless you specify `IGNORE`.
- You should use `TRUNCATE TABLE` when you want to delete all rows from a table and you do not need to obtain a count of the number of rows that were deleted. (`DELETE FROM tbl_name` returns a row count in 4.0 and does not reset the `AUTO_INCREMENT` counter, and `TRUNCATE TABLE` is faster.)
- You get an error if you have an active transaction or `LOCK TABLES` statement when trying to execute `TRUNCATE TABLE` or `DROP DATABASE`.
- To use `MATCH ... AGAINST (... IN BOOLEAN MODE)` full-text searches, you must rebuild existing table indexes using `REPAIR TABLE tbl_name USE_FRM`. If you attempt a boolean full-text search without rebuilding the indexes in this manner, the search returns incorrect results. See [Section 12.7.5, “Fine-Tuning MySQL Full-Text Search”](#).
- `LOCATE()` and `INSTR()` are case sensitive if one of the arguments is a binary string. Otherwise they are case insensitive.
- `STRCMP()` uses the current character set when performing comparisons. This makes the default comparison behavior not case sensitive unless one or both of the operands are binary strings.
- `HEX(str)` returns the characters in `str` converted to hexadecimal. If you want to convert a number to hexadecimal, you should ensure that you call `HEX()` with a numeric argument.
- `RAND(seed)` returns a different random number series in 4.0 than in 3.23; this was done to further differentiate `RAND(seed)` and `RAND(seed+1)`.
- The default type returned by `IFNULL(A,B)` is set to be the more “general” of the types of `A` and `B`. (The general-to-specific order is string, `REAL`, `INTEGER`).

### C API Changes:

- The old C API functions `mysql_drop_db()`, `mysql_create_db()`, and `mysql_connect()` are no longer supported in MySQL 4.0 unless MySQL is compiled with `CFLAGS=-DUSE_OLD_FUNCTIONS`. It is preferable to change client programs to use the new 4.0 API instead.
- In the `MYSQL_FIELD` structure, `length` and `max_length` have changed from `unsigned int` to `unsigned long`. This should not cause any problems, except that they may generate warning messages when used as arguments in the `printf()` class of functions.
- Multi-threaded clients should use `mysql_thread_init()` and `mysql_thread_end()`. See [Section 17.2.15, “How to Make a Threaded Client”](#).

### Other Changes:

- If you want to recompile the Perl `DBD::mysql` module, use a recent version. Version 2.9003 is recommended. Versions older than 1.2218 should not be used because they use the deprecated `mysql_drop_db()` call.



### 2.11.3. Copying MySQL Databases to Another Machine

If you are using MySQL 3.23 or later, you can copy the `.frm`, `.MYI`, and `.MYD` files for `MyISAM` tables between different architectures that support the same floating-point format. (MySQL takes care of any byte-swapping issues.) See [Section 14.1, “The MyISAM Storage Engine”](#).

The MySQL `ISAM` data and index files (`.ISD` and `*.ISM`, respectively) are dependent upon the architecture and, in some cases, the operating system. If you want to move applications to another machine having a different architecture or operating system than that of the current machine, you should not try to move a database by simply copying the files to the other machine. Use `mysqldump` instead.

By default, `mysqldump` creates a file containing SQL statements. You can then transfer the file to the other machine and use it as input to the `mysql` client.

Try `mysqldump --help` to see what options are available. If you are moving the data to a newer version of MySQL, you should use `mysqldump --opt` to take advantage of any optimizations that result in a dump file that is smaller and can be processed faster.

The easiest (although not the fastest) way to move a database between two machines is to run the following commands on the machine on which the database is located:

```
shell> mysqladmin -h 'other_hostname' create db_name
shell> mysqldump --opt db_name | mysql -h 'other_hostname' db_name
```

If you want to copy a database from a remote machine over a slow network, you can use these commands:

```
shell> mysqladmin create db_name
shell> mysqldump -h 'other_hostname' --opt --compress db_name | mysql db_name
```

You can also store the dump in a file, transfer the file to the target machine, and then load the file into the database there. For example, you can dump a database to a compressed file on the source machine like this:

```
shell> mysqldump --quick db_name | gzip > db_name.gz
```

Transfer the file containing the database contents to the target machine and run these commands there:

```
shell> mysqladmin create db_name
shell> gunzip < db_name.gz | mysql db_name
```

You can also use `mysqldump` and `mysqlimport` to transfer the database. For very large tables, this is much faster than simply using `mysqldump`. In the following commands, `DUMPDIR` represents the full pathname of the directory you use to store the output from `mysqldump`.

First, create the directory for the output files and dump the database:

```
shell> mkdir DUMPDIR
shell> mysqldump --tab=DUMPDIR db_name
```

Then transfer the files in the `DUMPDIR` directory to a directory on the target machine and load the files into MySQL there:

```
shell> mysqladmin create db_name           # create database
shell> cat DUMPDIR/*.sql | mysql db_name  # create tables in database
shell> mysqlimport db_name DUMPDIR/*.txt  # load data into tables
```

Do not forget to copy the `mysql` database because that is where the `user`, `db`, and `host` grant tables

are stored. You might have to run commands as the MySQL `root` user on the new machine until you have the `mysql` database in place.

After you import the `mysql` database on the new machine, execute `mysqladmin flush-privileges` so that the server reloads the grant table information.

## 2.12. Downgrading MySQL

This section describes what you should do to downgrade to an older MySQL version in the unlikely case that the previous version worked better than the new one.

If you are downgrading within the same release series (for example, from 4.0.20 to 4.0.19) the general rule is that you merely need to install the new binaries on top of the old ones. There is no need to do anything with the databases. As always, however, it is always a good idea to make a backup.

The following items form a checklist of things you should do whenever you perform a downgrade:

- Read the upgrading section for the release series from which you are downgrading to be sure that it does not have any features you really need. [Section 2.11, “Upgrading MySQL”](#).
- If there is a downgrading section for that version, please read it, too!

You can always move the MySQL format files and data files between different versions on the same architecture as long as you stay within versions for the same release series of MySQL.

If you downgrade from one release series to another, there may be incompatibilities in table storage formats. In this case, you can use `mysqldump` to dump your tables before downgrading. After downgrading, reload the dump file using `mysql` or `mysqlimport` to re-create your tables. For examples, see [Section 2.11.3, “Copying MySQL Databases to Another Machine”](#).

The normal symptom of a downward-incompatible table format change when you downgrade is that you cannot open tables. In that case, use the following procedure:

1. Stop the older MySQL server that you are downgrading to.
2. Restart the newer MySQL server you are downgrading from.
3. Dump any tables that were inaccessible to the older server by using `mysqldump` to create a dump file.
4. Stop the newer MySQL server and restart the older one.
5. Reload the dump file into the older server. Your tables should be accessible.

### 2.12.1. Downgrading to MySQL 4.0

The table format in 4.1 changed to include more and new character set information. Because of this, you must use `mysqldump` to dump any tables you have created with the newer MySQL server. For example, if all the tables in a particular database need to be dumped to be reverted back to MySQL 4.0 format, use this command:

```
shell> mysqldump --create-options --compatible=mysql40 db_name > dump_file
```

Then stop the newer server, restart the older server, and read in the dump file:

```
shell> mysql db_name < dump_file
```

In the special case that you are downgrading [MyISAM](#) tables, no special treatment is necessary if all columns in the tables contain only numeric columns or string columns ([CHAR](#), [VARCHAR](#), [TEXT](#), and so forth) that contain only [latin1](#) data. Your 4.1 tables should be directly usable with a 4.0 server.

If you used the [mysql\\_fix\\_privilege\\_tables](#) script to upgrade the grant tables, you can either use the preceding method to convert them to back to MySQL 4.0 or do the following in MySQL 4.1 (or above):

```
ALTER TABLE mysql.user
  CONVERT TO CHARACTER SET latin1 COLLATE latin1_swedish_ci;
ALTER TABLE mysql.db
  CONVERT TO CHARACTER SET latin1 COLLATE latin1_swedish_ci;
ALTER TABLE mysql.host
  CONVERT TO CHARACTER SET latin1 COLLATE latin1_swedish_ci;
ALTER TABLE mysql.tables_priv
  CONVERT TO CHARACTER SET latin1 COLLATE latin1_swedish_ci;
ALTER TABLE mysql.columns_priv
  CONVERT TO CHARACTER SET latin1 COLLATE latin1_swedish_ci;
ALTER TABLE mysql.func
  CONVERT TO CHARACTER SET latin1 COLLATE latin1_swedish_ci;
```

## 2.13. Operating System-Specific Notes

### 2.13.1. Linux Notes

This section discusses issues that have been found to occur on Linux. The first few subsections describe general operating system-related issues, problems that can occur when using binary or source distributions, and post-installation issues. The remaining subsections discuss problems that occur with Linux on specific platforms.

Note that most of these problems occur on older versions of Linux. If you are running a recent version, you may see none of them.

#### 2.13.1.1. Linux Operating System Notes

MySQL needs at least Linux version 2.0.

**Warning:** We have seen some strange problems with Linux 2.2.14 and MySQL on SMP systems. Some MySQL users have also reported that they have encountered serious stability problems using MySQL with kernel 2.2.14. If you are using this kernel, you should upgrade to 2.2.19 (or newer) or to a 2.4 or 2.6 kernel. If you have a multiple-CPU machine, you should seriously consider using 2.4 or 2.6 because it gives you a significant speed boost. Your system should also be more stable.

When using LinuxThreads, you should see a minimum of three [mysqld](#) processes running. These are in fact threads. There is one thread for the LinuxThreads manager, one thread to handle connections, and one thread to handle alarms and signals.

#### 2.13.1.2. Linux Binary Distribution Notes

The Linux-Intel binary and RPM releases of MySQL are configured for the highest possible speed. We are always trying to use the fastest stable compiler available.

The binary release is linked with `-static`, which means you do not normally need to worry about which version of the system libraries you have. You need not install LinuxThreads, either. A program linked with `-static` is slightly larger than a dynamically linked program, but also slightly faster (3-5%). However, one problem with a statically linked program is that you cannot use user-defined functions (UDFs). If you are going to write or use UDFs (this is something for C or C++ programmers

only), you must compile MySQL yourself using dynamic linking.

A known issue with binary distributions is that on older Linux systems that use `libc` (such as Red Hat 4.x or Slackware), you get some non-fatal problems with hostname resolution. If your system uses `libc` rather than `glibc2`, you probably will encounter some difficulties with hostname resolution and `getpwnam()`. This happens because `glibc` unfortunately depends on some external libraries to implement hostname resolution and `getpwent()`, even when compiled with `-static`. These problems manifest themselves in two ways:

- You may see the following error message when you run `mysql_install_db`:

```
Sorry, the host 'xxxx' could not be looked up
```

You can deal with this by executing `mysql_install_db --force`, which does not execute the `resolveip` test in `mysql_install_db`. The downside is that you cannot use hostnames in the grant tables: Except for `localhost`, you must use IP numbers instead. If you are using an old version of MySQL that does not support `--force`, you must manually remove the `resolveip` test in `mysql_install_db` using a text editor.

- You also may see the following error when you try to run `mysqld` with the `--user` option:

```
getpwnam: No such file or directory
```

To work around this problem, start `mysqld` by using the `su` command rather than by specifying the `--user` option. This causes the system itself to change the user ID of the `mysqld` process so that `mysqld` need not do so.

Another solution, which solves both problems, is to not use a binary distribution. Get a MySQL source distribution (in RPM or `tar.gz` format) and install that instead.

On some Linux 2.2 versions, you may get the error `Resource temporarily unavailable` when clients make a lot of new connections to a `mysqld` server over TCP/IP. The problem is that Linux has a delay between the time that you close a TCP/IP socket and the time that the system actually frees it. There is room for only a finite number of TCP/IP slots, so you encounter the resource-unavailable error if clients attempt too many new TCP/IP connections during a short time. For example, you may see the error when you run the MySQL `test-connect` benchmark over TCP/IP.

We have inquired about this problem a few times on different Linux mailing lists but have never been able to find a suitable resolution. The only known “fix” is for the clients to use persistent connections, or, if you are running the database server and clients on the same machine, to use Unix socket file connections rather than TCP/IP connections.

### 2.13.1.3. Linux Source Distribution Notes

The following notes regarding `glibc` apply only to the situation when you build MySQL yourself. If you are running Linux on an x86 machine, in most cases it is much better for you to just use our binary. We link our binaries against the best patched version of `glibc` we can find and with the best compiler options, in an attempt to make it suitable for a high-load server. For a typical user, even for setups with a lot of concurrent connections or tables exceeding the 2GB limit, our binary is the best choice in most cases. After reading the following text, if you are in doubt about what to do, try our binary first to determine whether it meets your needs. If you discover that it is not good enough, you may want to try your own build. In that case, we would appreciate a note about it so that we can build a better binary next time.

MySQL uses LinuxThreads on Linux. If you are using an old Linux version that does not have `glibc2`, you must install LinuxThreads before trying to compile MySQL. You can obtain LinuxThreads at [ht-](#)

[tp://dev.mysql.com/downloads/os-linux.html](http://dev.mysql.com/downloads/os-linux.html).

Note that `glibc` versions before and including version 2.1.1 have a fatal bug in `pthread_mutex_timedwait()` handling, which is used when you issue `INSERT DELAYED` statements. We recommend that you not use `INSERT DELAYED` before upgrading `glibc`.

Note that Linux kernel and the LinuxThread library can by default have only 1,024 threads. If you plan to have more than 1,000 concurrent connections, you need to make some changes to LinuxThreads:

- Increase `PTHREAD_THREADS_MAX` in `sysdeps/unix/sysv/linux/bits/local_lim.h` to 4096 and decrease `STACK_SIZE` in `linuxthreads/internals.h` to 256KB. The paths are relative to the root of `glibc`. (Note that MySQL is not stable with around 600-1000 connections if `STACK_SIZE` is the default of 2MB.)
- Recompile LinuxThreads to produce a new `libpthread.a` library, and relink MySQL against it.

The page <http://www.volano.com/linuxnotes.html> contains additional information about circumventing thread limits in LinuxThreads.

There is another issue that greatly hurts MySQL performance, especially on SMP systems. The mutex implementation in LinuxThreads in `glibc` 2.1 is very bad for programs with many threads that hold the mutex only for a short time. This produces a paradoxical result: If you link MySQL against an unmodified LinuxThreads, removing processors from an SMP actually improves MySQL performance in many cases. We have made a patch available for `glibc` 2.1.3 to correct this behavior (<http://dev.mysql.com/Downloads/Linux/linuxthreads-2.1-patch>).

With `glibc` 2.2.2, MySQL 3.23.36 uses the adaptive mutex, which is much better than even the patched one in `glibc` 2.1.3. Be warned, however, that under some conditions, the current mutex code in `glibc` 2.2.2 overspins, which hurts MySQL performance. The likelihood that this condition occurs can be reduced by renicing the `mysqld` process to the highest priority. We have also been able to correct the overspin behavior with a patch, available at <http://dev.mysql.com/Downloads/Linux/linuxthreads-2.2.2.patch>. It combines the correction of overspin, maximum number of threads, and stack spacing all in one. You need to apply it in the `linuxthreads` directory with `patch -p0 </tmp/linuxthreads-2.2.2.patch`. We hope it is included in some form in future releases of `glibc` 2.2. In any case, if you link against `glibc` 2.2.2, you still need to correct `STACK_SIZE` and `PTHREAD_THREADS_MAX`. We hope that the defaults is corrected to some more acceptable values for high-load MySQL setup in the future, so that the commands needed to produce your own build can be reduced to `./configure; make; make install`.

We recommend that you use these patches to build a special static version of `libpthread.a` and use it only for statically linking against MySQL. We know that the patches are safe for MySQL and significantly improve its performance, but we cannot say anything about other applications. If you link other applications that require LinuxThreads against the patched static version of the library, or build a patched shared version and install it on your system, you do so at your own risk.

If you experience any strange problems during the installation of MySQL, or with some common utilities hanging, it is very likely that they are either library or compiler related. If this is the case, using our binary resolves them.

If you link your own MySQL client programs, you may see the following error at runtime:

```
ld.so.1: fatal: libmysqlclient.so.#:
open failed: No such file or directory
```

This problem can be avoided by one of the following methods:

- Link clients with the `-Wl,r/full/path/to/libmysqlclient.so` flag rather than with `-`

`Lpath`).

- Copy `libmysqlclient.so` to `/usr/lib`.
- Add the pathname of the directory where `libmysqlclient.so` is located to the `LD_RUN_PATH` environment variable before running your client.

If you are using the Fujitsu compiler (`fcc/FCC`), you may have some problems compiling MySQL because the Linux header files are very `gcc` oriented. The following `configure` line should work with `fcc/FCC`:

```
CC=fcc CFLAGS="-O -K fast -K lib -K omitfp -Kpreex -D_GNU_SOURCE \
-DCONST=const -DNO_STRTOLL_PROTO" \
CXX=FCC CXXFLAGS="-O -K fast -K lib \
-K omitfp -K preex --no_exceptions --no_rtti -D_GNU_SOURCE \
-DCONST=const -Dalloca=__builtin_alloca -DNO_STRTOLL_PROTO \
'-D_EXTERN_INLINE=static __inline'" \
./configure \
--prefix=/usr/local/mysql --enable-assembly \
--with-mysqld-ldflags=-all-static --disable-shared \
--with-low-memory
```

### 2.13.1.4. Linux Post-Installation Notes

`mysql.server` can be found in the `support-files` directory under the MySQL installation directory or in a MySQL source tree. You can install it as `/etc/init.d/mysql` for automatic MySQL startup and shutdown. See [Section 2.10.2.2, “Starting and Stopping MySQL Automatically”](#).

If MySQL cannot open enough files or connections, it may be that you have not configured Linux to handle enough files.

In Linux 2.2 and onward, you can check the number of allocated file handles as follows:

```
shell> cat /proc/sys/fs/file-max
shell> cat /proc/sys/fs/dquot-max
shell> cat /proc/sys/fs/super-max
```

If you have more than 16MB of memory, you should add something like the following to your init scripts (for example, `/etc/init.d/boot.local` on SuSE Linux):

```
echo 65536 > /proc/sys/fs/file-max
echo 8192 > /proc/sys/fs/dquot-max
echo 1024 > /proc/sys/fs/super-max
```

You can also run the `echo` commands from the command line as `root`, but these settings are lost the next time your computer restarts.

Alternatively, you can set these parameters on startup by using the `sysctl` tool, which is used by many Linux distributions (SuSE has added it as well, beginning with SuSE Linux 8.0). Just put the following values into a file named `/etc/sysctl.conf`:

```
# Increase some values for MySQL
fs.file-max = 65536
fs.dquot-max = 8192
fs.super-max = 1024
```

You should also add the following to `/etc/my.cnf`:

```
[mysqld_safe]
open-files-limit=8192
```

This should allow the server a limit of 8,192 for the combined number of connections and open files.

The `STACK_SIZE` constant in `LinuxThreads` controls the spacing of thread stacks in the address space. It needs to be large enough so that there is plenty of room for each individual thread stack, but small enough to keep the stack of some threads from running into the global `mysqld` data. Unfortunately, as we have discovered, the Linux implementation of `mmap()` successfully unmaps a mapped region if you ask it to map out an address currently in use, zeroing out the data on the entire page instead of returning an error. So, the safety of `mysqld` or any other threaded application depends on “gentlemanly” behavior of the code that creates threads. The user must take measures to make sure that the number of running threads at any time is sufficiently low for thread stacks to stay away from the global heap. With `mysqld`, you should enforce this behavior by setting a reasonable value for the `max_connections` variable.

If you build MySQL yourself, you can patch `LinuxThreads` for better stack use. See [Section 2.13.1.3, “Linux Source Distribution Notes”](#). If you do not want to patch `LinuxThreads`, you should set `max_connections` to a value no higher than 500. It should be even less if you have a large key buffer, large heap tables, or some other things that make `mysqld` allocate a lot of memory, or if you are running a 2.2 kernel with a 2GB patch. If you are using our binary or RPM version 3.23.25 or later, you can safely set `max_connections` at 1500, assuming no large key buffer or heap tables with lots of data. The more you reduce `STACK_SIZE` in `LinuxThreads` the more threads you can safely create. We recommend values between 128KB and 256KB.

If you use a lot of concurrent connections, you may suffer from a “feature” in the 2.2 kernel that attempts to prevent fork bomb attacks by penalizing a process for forking or cloning a child. This causes MySQL not to scale well as you increase the number of concurrent clients. On single-CPU systems, we have seen this manifested as very slow thread creation: It may take a long time to connect to MySQL (as long as one minute), and it may take just as long to shut it down. On multiple-CPU systems, we have observed a gradual drop in query speed as the number of clients increases. In the process of trying to find a solution, we have received a kernel patch from one of our users who claimed it made a lot of difference for his site. The patch is available at <http://dev.mysql.com/Downloads/Patches/linux-fork.patch>. We have done rather extensive testing of this patch on both development and production systems. It has significantly improved MySQL performance without causing any problems and we recommend it to our users who still run high-load servers on 2.2 kernels.

This issue has been fixed in the 2.4 kernel, so if you are not satisfied with the current performance of your system, rather than patching your 2.2 kernel, it might be easier to upgrade to 2.4. On SMP systems, upgrading also gives you a nice SMP boost in addition to fixing the fairness bug.

We have tested MySQL on the 2.4 kernel on a two-CPU machine and found MySQL scales *much* better. There was virtually no slowdown on query throughput all the way up to 1,000 clients, and the MySQL scaling factor (computed as the ratio of maximum throughput to the throughput for one client) was 180%. We have observed similar results on a four-CPU system: Virtually no slowdown as the number of clients was increased up to 1,000, and a 300% scaling factor. Based on these results, for a high-load SMP server using a 2.2 kernel, we definitely recommend upgrading to the 2.4 kernel at this point.

We have discovered that it is essential to run the `mysqld` process with the highest possible priority on the 2.4 kernel to achieve maximum performance. This can be done by adding a `renice -20 $$` command to `mysqld_safe`. In our testing on a four-CPU machine, increasing the priority resulted in a 60% throughput increase with 400 clients.

We are currently also trying to collect more information on how well MySQL performs with a 2.4 kernel on four-way and eight-way systems. If you have access such a system and have done some benchmarks, please send an email message to [<benchmarks@mysql.com>](mailto:benchmarks@mysql.com) with the results. We will review them for inclusion in the manual.

If you see a dead `mysqld` server process with `ps`, this usually means that you have found a bug in MySQL or that you have a corrupted table. See [Section A.4.2, “What to Do If MySQL Keeps Crashing”](#).

To get a core dump on Linux if `mysqld` dies with a `SIGSEGV` signal, you can start `mysqld` with the `--core-file` option. Note that you also probably need to raise the core file size by adding `ulimit -c 1000000` to `mysqld_safe` or starting `mysqld_safe` with `--core-file-size=1000000`. See [Section 5.4.1, “mysqld\\_safe — MySQL Server Startup Script”](#).

### 2.13.1.5. Linux x86 Notes

MySQL requires `libc` 5.4.12 or newer. It is known to work with `libc` 5.4.46. `glibc` 2.0.6 and later should also work. There have been some problems with the `glibc` RPMs from Red Hat, so if you have problems, check whether there are any updates. The `glibc` 2.0.7-19 and 2.0.7-29 RPMs are known to work.

If you are using Red Hat 8.0 or a new `glibc` 2.2.x library, you may see `mysqld` die in `gethostbyaddr()`. This happens because the new `glibc` library requires a stack size greater than 128KB for this call. To fix the problem, start `mysqld` with the `--thread-stack=192K` option. (Use `-O thread_stack=192K` before MySQL 4.) This stack size is the default on MySQL 4.0.10 and above, so you should not see the problem.

If you are using `gcc` 3.0 and above to compile MySQL, you must install the `libstdc++v3` library before compiling MySQL; if you do not do this, you get an error about a missing `__cxa_pure_virtual` symbol during linking.

On some older Linux distributions, `configure` may produce an error like this:

```
Syntax error in sched.h. Change _P to __P in the
/usr/include/sched.h file.
See the Installation chapter in the Reference Manual.
```

Just do what the error message says. Add an extra underscore to the `_P` macro name that has only one underscore, and then try again.

You may get some warnings when compiling. Those shown here can be ignored:

```
mysqld.cc -o objs-thread/mysqld.o
mysqld.cc: In function `void init_signals()':
mysqld.cc:315: warning: assignment of negative value `-1' to
`long unsigned int'
mysqld.cc: In function `void * signal_hand(void *)':
mysqld.cc:346: warning: assignment of negative value `-1' to
`long unsigned int'
```

If `mysqld` always dumps core when it starts, the problem may be that you have an old `/lib/libc.a`. Try renaming it, and then remove `sql/mysqld` and do a new `make install` and try again. This problem has been reported on some Slackware installations.

If you get the following error when linking `mysqld`, it means that your `libg++.a` is not installed correctly:

```
/usr/lib/libc.a(putc.o): In function `_IO_putc':
putc.o(.text+0x0): multiple definition of `_IO_putc'
```

You can avoid using `libg++.a` by running `configure` like this:

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

If `mysqld` crashes immediately and you are running Red Hat 5.0 with a version of `glibc` older than 2.0.7-5, you should make sure that you have installed all `glibc` patches. There is a lot of information about this in the MySQL mail archives, available online at <http://lists.mysql.com/>.



### 2.13.1.6. Linux SPARC Notes

In some implementations, `readdir_r()` is broken. The symptom is that the `SHOW DATABASES` statement always returns an empty set. This can be fixed by removing `HAVE_READDIR_R` from `config.h` after configuring and before compiling.

### 2.13.1.7. Linux Alpha Notes

MySQL 3.23.12 is the first MySQL version that is tested on Linux-Alpha. If you plan to use MySQL on Linux-Alpha, you should ensure that you have this version or newer.

We have tested MySQL on Alpha with our benchmarks and test suite, and it appears to work nicely.

We currently build the MySQL binary packages on SuSE Linux 7.0 for AXP, kernel 2.4.4-SMP, Compaq C compiler (V6.2-505) and Compaq C++ compiler (V6.3-006) on a Compaq DS20 machine with an Alpha EV6 processor.

You can find the preceding compilers at <http://www.support.compaq.com/alpha-tools/>. By using these compilers rather than `gcc`, we get about 9-14% better MySQL performance.

Note that until MySQL version 3.23.52 and 4.0.2, we optimized the binary for the current CPU only (by using the `-fast` compile option). This means that for older versions, you can use our Alpha binaries only if you have an Alpha EV6 processor.

For all subsequent releases, we added the `-arch generic` flag to our compile options, which ensures that the binary runs on all Alpha processors. We also compile statically to avoid library problems. The `configure` command looks like this:

```
CC=ccc CFLAGS="-fast -arch generic" CXX=cxx \
CXXFLAGS="-fast -arch generic -noexceptions -nortti" \
./configure --prefix=/usr/local/mysql --disable-shared \
--with-extra-charsets=complex --enable-thread-safe-client \
--with-mysqld-ldflags=-non_shared --with-client-ldflags=-non_shared
```

If you want to use `egcs`, the following `configure` line worked for us:

```
CFLAGS="-O3 -fomit-frame-pointer" CXX=gcc \
CXXFLAGS="-O3 -fomit-frame-pointer -felide-constructors \
-fno-exceptions -fno-rtti" \
./configure --prefix=/usr/local/mysql --disable-shared
```

Some known problems when running MySQL on Linux-Alpha:

- Debugging threaded applications like MySQL does not work with `gdb` 4.18. You should use `gdb` 5.1 instead.
- If you try linking `mysqld` statically when using `gcc`, the resulting image dumps core at startup time. In other words, *do not* use `--with-mysqld-ldflags=-all-static` with `gcc`.

### 2.13.1.8. Linux PowerPC Notes

MySQL should work on MkLinux with the newest `glibc` package (tested with `glibc` 2.0.7).

### 2.13.1.9. Linux MIPS Notes

To get MySQL to work on Qube2 (Linux Mips), you need the newest `glibc` libraries. `glibc-2.0.7-29C2` is known to work. You must also use the `egcs` C++ compiler (`egcs` 1.0.2-9, `gcc`

2.95.2 or newer).

### 2.13.1.10. Linux IA-64 Notes

To get MySQL to compile on Linux IA-64, we use the following `configure` command for building with `gcc 2.96`:

```
CC=gcc \  
CFLAGS="-O3 -fno-omit-frame-pointer" \  
CXX=gcc \  
CXXFLAGS="-O3 -fno-omit-frame-pointer -felide-constructors \  
-fno-exceptions -fno-rtti" \  
./configure --prefix=/usr/local/mysql \  
"--with-comment=Official MySQL binary" \  
"--with-extra-charsets=complex
```

On IA-64, the MySQL client binaries use shared libraries. This means that if you install our binary distribution at a location other than `/usr/local/mysql`, you need to add the path of the directory where you have `libmysqlclient.so` installed either to the `/etc/ld.so.conf` file or to the value of your `LD_LIBRARY_PATH` environment variable.

See [Section A.3.1, “Problems Linking to the MySQL Client Library”](#).

### 2.13.1.11. SELinux Notes

RHEL4 comes with SELinux, which supports tighter access control for processes. If SELinux is enabled (`SELINUX` in `/etc/selinux/config` is set to `enforcing`, `SELINUXTYPE` is set to either `targeted` or `strict`), you might encounter problems installing MySQL AB RPM packages.

Red Hat has an update that solves this. It involves an update of the “security policy” specification to handle the install structure of the RPMs provided by MySQL AB. For further information, see [https://bugzilla.redhat.com/bugzilla/show\\_bug.cgi?id=167551](https://bugzilla.redhat.com/bugzilla/show_bug.cgi?id=167551) and <http://rhn.redhat.com/errata/RHBA-2006-0049.html>.

## 2.13.2. Mac OS X Notes

On Mac OS X, `tar` cannot handle long filenames. If you need to unpack a `.tar.gz` distribution, use `gnutar` instead.

### 2.13.2.1. Mac OS X 10.x (Darwin)

MySQL should work without major problems on Mac OS X 10.x (Darwin).

Known issues:

- If you have problems with performance under heavy load, try using the `--skip-thread-priority` option to `mysqld`. This runs all threads with the same priority. On Mac OS X, this gives better performance, at least until Apple fixes its thread scheduler.
- The connection times (`wait_timeout`, `interactive_timeout` and `net_read_timeout`) values are not honored. The symptom is that persistent connections can hang for a very long time without getting closed down and that a 'kill' for a thread will not take affect until the thread does it a new command

This is probably a signal handling problem in the thread library where the signal does not break a pending read and we hope that a future update to the thread libraries will fix this.

Our binary for Mac OS X is compiled on Darwin 6.3 with the following `configure` line:

```
CC=gcc CFLAGS="-O3 -fno-omit-frame-pointer" CXX=gcc \
CXXFLAGS="-O3 -fno-omit-frame-pointer -felide-constructors \
-fno-exceptions -fno-rtti" \
./configure --prefix=/usr/local/mysql \
--with-extra-charsets=complex --enable-thread-safe-client \
--enable-local-infile --disable-shared
```

See [Section 2.5, “Installing MySQL on Mac OS X”](#).

### 2.13.2.2. Mac OS X Server 1.2 (Rhapsody)

For current versions of Mac OS X Server, no operating system changes are necessary before compiling MySQL. Compiling for the Server platform is the same as for the client version of Mac OS X. (However, note that MySQL comes preinstalled on Mac OS X Server, so you need not build it yourself.)

For older versions (Mac OS X Server 1.2, a.k.a. Rhapsody), you must first install a pthread package before trying to configure MySQL.

See [Section 2.5, “Installing MySQL on Mac OS X”](#).

### 2.13.3. Solaris Notes

For information about installing MySQL on Solaris using PKG distributions, see [Section 2.6, “Installing MySQL on Solaris”](#).

On Solaris, you may run into trouble even before you get the MySQL distribution unpacked. Solaris `tar` cannot handle long filenames, so you may see an error like this when you unpack MySQL:

```
x mysql-3.22.12-beta/bench/Results/ATIS-mysql_odbc-NT_4.0-cmp-db2,
informix,ms-sql,mysql,oracle,solid,sybase, 0 bytes, 0 tape blocks
tar: directory checksum error
```

In this case, you must use GNU `tar` (`gtar`) to unpack the distribution. You can find a precompiled copy for Solaris at <http://dev.mysql.com/downloads/os-solaris.html>.

Sun native threads work only on Solaris 2.5 and higher. For Solaris 2.4 and earlier, MySQL automatically uses MIT-pthreads. See [Section 2.9.5, “MIT-pthreads Notes”](#).

If you get the following error from `configure`, it means that you have something wrong with your compiler installation:

```
checking for restartable system calls... configure: error cannot
run test programs while cross compiling
```

In this case, you should upgrade your compiler to a newer version. You may also be able to solve this problem by inserting the following row into the `config.cache` file:

```
ac_cv_sys_restartable_syscalls=${ac_cv_sys_restartable_syscalls='no'}
```

If you are using Solaris on a SPARC, the recommended compiler is `gcc` 2.95.2 or 3.2. You can find this at <http://gcc.gnu.org/>. Note that `egcs` 1.1.1 and `gcc` 2.8.1 do not work reliably on SPARC!

The recommended `configure` line when using `gcc` 2.95.2 is:

```
CC=gcc CFLAGS="-O3" \
CXX=gcc CXXFLAGS="-O3 -felide-constructors -fno-exceptions -fno-rtti" \
./configure --prefix=/usr/local/mysql --with-low-memory \
--enable-assembler
```

If you have an UltraSPARC system, you can get 4% better performance by adding `-mcpu=v8 -Wa, -xarch=v8plusa` to the `CFLAGS` and `CXXFLAGS` environment variables.

If you have Sun's Forte 5.0 (or newer) compiler, you can run `configure` like this:

```
CC=cc CFLAGS="-Xa -fast -native -xstrconst -mt" \
CXX=CC CXXFLAGS="-noex -mt" \
./configure --prefix=/usr/local/mysql --enable-asmbler
```

To create a 64-bit binary with Sun's Forte compiler, use the following configuration options:

```
CC=cc CFLAGS="-Xa -fast -native -xstrconst -mt -xarch=v9" \
CXX=CC CXXFLAGS="-noex -mt -xarch=v9" ASFLAGS="-xarch=v9" \
./configure --prefix=/usr/local/mysql --enable-asmbler
```

To create a 64-bit Solaris binary using `gcc`, add `-m64` to `CFLAGS` and `CXXFLAGS` and remove `--enable-asmbler` from the `configure` line. This works only with MySQL 4.0 and up; MySQL 3.23 does not include the required modifications to support this.

In the MySQL benchmarks, we got a 4% speedup on an UltraSPARC when using Forte 5.0 in 32-bit mode compared to using `gcc` 3.2 with the `-mcpu` flag.

If you create a 64-bit `mysqld` binary, it is 4% slower than the 32-bit binary, but can handle more threads and memory.

When using Solaris 10 for x86\_64, you should mount any filesystems on which you intend to store `InnoDB` files with the `forcedirectio` option. (By default mounting is done without this option.) Failing to do so will cause a significant drop in performance when using the `InnoDB` storage engine on this platform.

If you get a problem with `fdatasync` or `sched_yield`, you can fix this by adding `LIBS=-lrt` to the `configure` line

For compilers older than WorkShop 5.3, you might have to edit the `configure` script. Change this line:

```
#if !defined(__STDC__) || __STDC__ != 1
```

To this:

```
#if !defined(__STDC__)
```

If you turn on `__STDC__` with the `-Xc` option, the Sun compiler cannot compile with the Solaris `pthread.h` header file. This is a Sun bug (broken compiler or broken include file).

If `mysqld` issues the following error message when you run it, you have tried to compile MySQL with the Sun compiler without enabling the `-mt` multi-thread option:

```
libc internal error: _rmutex_unlock: rmutex not held
```

Add `-mt` to `CFLAGS` and `CXXFLAGS` and recompile.

If you are using the SFW version of `gcc` (which comes with Solaris 8), you must add `/opt/sfw/lib` to the environment variable `LD_LIBRARY_PATH` before running `configure`.

If you are using the `gcc` available from `sunfreeware.com`, you may have many problems. To avoid this, you should recompile `gcc` and GNU `binutils` on the machine where you are running them.

If you get the following error when compiling MySQL with `gcc`, it means that your `gcc` is not configured for your version of Solaris:

```
shell> gcc -O3 -g -O2 -DDEBUG_OFF -o thr_alarm ...
./thr_alarm.c: In function `signal_hand':
./thr_alarm.c:556: too many arguments to function `sigwait'
```

The proper thing to do in this case is to get the newest version of `gcc` and compile it with your current `gcc` compiler. At least for Solaris 2.5, almost all binary versions of `gcc` have old, unusable include files that break all programs that use threads, and possibly other programs!

Solaris does not provide static versions of all system libraries (`libpthreads` and `libdl`), so you cannot compile MySQL with `--static`. If you try to do so, you get one of the following errors:

```
ld: fatal: library -ldl: not found
undefined reference to `dlopen'
cannot find -lrt
```

If you link your own MySQL client programs, you may see the following error at runtime:

```
ld.so.1: fatal: libmysqlclient.so.#:
open failed: No such file or directory
```

This problem can be avoided by one of the following methods:

- Link clients with the `-Wl,r/full/path/to/libmysqlclient.so` flag rather than with `-Lpath`).
- Copy `libmysqlclient.so` to `/usr/lib`.
- Add the pathname of the directory where `libmysqlclient.so` is located to the `LD_RUN_PATH` environment variable before running your client.

If you have problems with `configure` trying to link with `-lz` when you do not have `zlib` installed, you have two options:

- If you want to be able to use the compressed communication protocol, you need to get and install `zlib` from [ftp.gnu.org](http://ftp.gnu.org).
- Run `configure` with the `--with-named-z-libs=no` option when building MySQL.

If you are using `gcc` and have problems with loading user-defined functions (UDFs) into MySQL, try adding `-lgcc` to the link line for the UDF.

If you would like MySQL to start automatically, you can copy `support-files/mysql.server` to `/etc/init.d` and create a symbolic link to it named `/etc/rc3.d/S99mysql.server`.

If too many processes try to connect very rapidly to `mysqld`, you should see this error in the MySQL log:

```
Error in accept: Protocol error
```

You might try starting the server with the `--back_log=50` option as a workaround for this. (Use `-O back_log=50` before MySQL 4.)

Solaris does not support core files for `setuid()` applications, so you cannot get a core file from `mysqld` if you are using the `--user` option.

### 2.13.3.1. Solaris 2.7/2.8 Notes

Normally, you can use a Solaris 2.6 binary on Solaris 2.7 and 2.8. Most of the Solaris 2.6 issues also apply for Solaris 2.7 and 2.8.

MySQL 3.23.4 and above should be able to detect new versions of Solaris automatically and enable workarounds for the following problems.

Solaris 2.7 and 2.8 have some bugs in the include files. You may see the following error when you use `gcc`:

```
/usr/include/widec.h:42: warning: `getwc' redefined
/usr/include/wchar.h:326: warning: this is the location of the previous
definition
```

If this occurs, you can fix the problem by copying `/usr/include/widec.h` to `../lib/gcc-lib/os/gcc-version/include` and changing line 41 from this:

```
#if !defined(lint) && !defined(__lint)
```

To this:

```
#if !defined(lint) && !defined(__lint) && !defined(getwc)
```

Alternatively, you can edit `/usr/include/widec.h` directly. Either way, after you make the fix, you should remove `config.cache` and run `configure` again.

If you get the following errors when you run `make`, it is because `configure` did not detect the `curses.h` file (probably because of the error in `/usr/include/widec.h`):

```
In file included from mysql.cc:50:
/usr/include/term.h:1060: syntax error before `;'
/usr/include/term.h:1081: syntax error before `;'
```

The solution to this problem is to do one of the following:

- Configure with `CFLAGS=-DHAVE_CURSES_H CXXFLAGS=-DHAVE_CURSES_H` and run `./configure`.
- Edit `/usr/include/widec.h` as indicated in the preceding discussion and re-run `configure`.
- Remove the `#define HAVE_TERM` line from the `config.h` file and run `make` again.

If your linker cannot find `-lz` when linking client programs, the problem is probably that your `libz.so` file is installed in `/usr/local/lib`. You can fix this problem by one of the following methods:

- Add `/usr/local/lib` to `LD_LIBRARY_PATH`.
- Add a link to `libz.so` from `/lib`.
- If you are using Solaris 8, you can install the optional `zlib` from your Solaris 8 CD distribution.

- Run `configure` with the `--with-named-z-libs=no` option when building MySQL.

### 2.13.3.2. Solaris x86 Notes

On Solaris 8 on x86, `mysqld` dumps core if you remove the debug symbols using `strip`.

If you are using `gcc` or `egcs` on Solaris x86 and you experience problems with core dumps under load, you should use the following `configure` command:

```
CC=gcc CFLAGS="-O3 -fomit-frame-pointer -DHAVE_CURSES_H" \  
CXX=gcc \  
CXXFLAGS="-O3 -fomit-frame-pointer -felide-constructors \  
-fno-exceptions -fno-rtti -DHAVE_CURSES_H" \  
./configure --prefix=/usr/local/mysql
```

This avoids problems with the `libstdc++` library and with C++ exceptions.

If this does not help, you should compile a debug version and run it with a trace file or under `gdb`. See [Section E.1.3, “Debugging `mysqld` under `gdb`”](#).

### 2.13.4. BSD Notes

This section provides information about using MySQL on variants of BSD Unix.

#### 2.13.4.1. FreeBSD Notes

FreeBSD 4.x or newer is recommended for running MySQL, because the thread package is much more integrated. To get a secure and stable system, you should use only FreeBSD kernels that are marked `-RELEASE`.

The easiest (and preferred) way to install MySQL is to use the `mysql-server` and `mysql-client` ports available at <http://www.freebsd.org/>. Using these ports gives you the following benefits:

- A working MySQL with all optimizations enabled that are known to work on your version of FreeBSD.
- Automatic configuration and build.
- Startup scripts installed in `/usr/local/etc/rc.d`.
- The ability to use `pkg_info -L` to see which files are installed.
- The ability to use `pkg_delete` to remove MySQL if you no longer want it on your machine.

It is recommended you use MIT-pthreads on FreeBSD 2.x, and native threads on FreeBSD 3 and up. It is possible to run with native threads on some late 2.2.x versions, but you may encounter problems shutting down `mysqld`.

Unfortunately, certain function calls on FreeBSD are not yet fully thread-safe. Most notably, this includes the `gethostbyname()` function, which is used by MySQL to convert hostnames into IP addresses. Under certain circumstances, the `mysqld` process suddenly causes 100% CPU load and is unresponsive. If you encounter this problem, try to start MySQL using the `--skip-name-resolve` option.

Alternatively, you can link MySQL on FreeBSD 4.x against the LinuxThreads library, which avoids a few of the problems that the native FreeBSD thread implementation has. For a very good comparison of

LinuxThreads versus native threads, see Jeremy Zawodny's article *FreeBSD or Linux for your MySQL Server?* at <http://jeremy.zawodny.com/blog/archives/000697.html>.

Known problem when using LinuxThreads on FreeBSD is:

- The connection times (`wait_timeout`, `interactive_timeout` and `net_read_timeout`) values are not honored. The symptom is that persistent connections can hang for a very long time without getting closed down and that a 'kill' for a thread will not take affect until the thread does it a new command

This is probably a signal handling problem in the thread library where the signal does not break a pending read. This is supposed to be fixed in FreeBSD 5.0

The MySQL build process requires GNU make (`gmake`) to work. If GNU `make` is not available, you must install it first before compiling MySQL.

The recommended way to compile and install MySQL on FreeBSD with `gcc` (2.95.2 and up) is:

```
CC=gcc CFLAGS="-O2 -fno-strength-reduce" \
CXX=gcc CXXFLAGS="-O2 -fno-rtti -fno-exceptions \
-felide-constructors -fno-strength-reduce" \
./configure --prefix=/usr/local/mysql --enable-assembler
gmake
gmake install
cd /usr/local/mysql
bin/mysql_install_db --user=mysql
bin/mysqld_safe &
```

If you notice that `configure` uses MIT-pthreads, you should read the MIT-pthreads notes. See [Section 2.9.5, "MIT-pthreads Notes"](#).

If you get an error from `make install` that it cannot find `/usr/include/pthreads`, `configure` did not detect that you need MIT-pthreads. To fix this problem, remove `config.cache`, and then re-run `configure` with the `--with-mit-threads` option.

Be sure that your name resolver setup is correct. Otherwise, you may experience resolver delays or failures when connecting to `mysqld`. Also make sure that the `localhost` entry in the `/etc/hosts` file is correct. The file should start with a line similar to this:

```
127.0.0.1      localhost localhost.your.domain
```

FreeBSD is known to have a very low default file handle limit. See [Section A.2.17, "File Not Found"](#). Start the server by using the `--open-files-limit` option for `mysqld_safe`, or raise the limits for the `mysqld` user in `/etc/login.conf` and rebuild it with `cap_mkdb /etc/login.conf`. Also be sure that you set the appropriate class for this user in the password file if you are not using the default (use `chpass mysqld-user-name`). See [Section 5.4.1, "mysqld\\_safe — MySQL Server Startup Script"](#).

FreeBSD limits the size of a process to 512MB, even if you have much more RAM available on the system. So you may get an error such as this:

```
Out of memory (Needed 16391 bytes)
```

In current versions of FreeBSD (at least 4.x and greater), you may increase this limit by adding the following entries to the `/boot/loader.conf` file and rebooting the machine (these are not settings that can be changed at run time with the `sysctl` command):

```
kern.maxdsiz="1073741824" # 1GB
kern.dflsiz="1073741824" # 1GB
```



```
kern.maxssiz="134217728" # 128MB
```

For older versions of FreeBSD, you must recompile your kernel to change the maximum data segment size for a process. In this case, you should look at the `MAXDSIZ` option in the `LINT` config file for more information.

If you get problems with the current date in MySQL, setting the `TZ` variable should help. See [Appendix F, Environment Variables](#).

### 2.13.4.2. NetBSD Notes

To compile on NetBSD, you need GNU `make`. Otherwise, the build process fails when `make` tries to run `lint` on C++ files.

### 2.13.4.3. OpenBSD 2.5 Notes

On OpenBSD 2.5, you can compile MySQL with native threads with the following options:

```
CFLAGS=-pthread CXXFLAGS=-pthread ./configure --with-mit-threads=no
```

### 2.13.4.4. BSD/OS Version 2.x Notes

If you get the following error when compiling MySQL, your `ulimit` value for virtual memory is too low:

```
item_func.h: In method
`Item_func_ge::Item_func_ge(const Item_func_ge &)`:
item_func.h:28: virtual memory exhausted
make[2]: *** [item_func.o] Error 1
```

Try using `ulimit -v 80000` and run `make` again. If this does not work and you are using `bash`, try switching to `csh` or `sh`; some BSDI users have reported problems with `bash` and `ulimit`.

If you are using `gcc`, you may also use have to use the `--with-low-memory` flag for `configure` to be able to compile `sql_yacc.cc`.

If you get problems with the current date in MySQL, setting the `TZ` variable should help. See [Appendix F, Environment Variables](#).

### 2.13.4.5. BSD/OS Version 3.x Notes

Upgrade to BSD/OS 3.1. If that is not possible, install BSDIpatch M300-038.

Use the following command when configuring MySQL:

```
env CXX=shlcc++ CC=shlcc2 \
./configure \
--prefix=/usr/local/mysql \
--localstatedir=/var/mysql \
--without-perl \
--with-unix-socket-path=/var/mysql/mysql.sock
```

The following is also known to work:

```
env CC=gcc CXX=gcc CXXFLAGS=-O3 \
./configure \
--prefix=/usr/local/mysql \
--with-unix-socket-path=/var/mysql/mysql.sock
```

You can change the directory locations if you wish, or just use the defaults by not specifying any locations.

If you have problems with performance under heavy load, try using the `--skip-thread-priority` option to `mysqld`. This runs all threads with the same priority. On BSDI 3.1, this gives better performance, at least until BSDI fixes its thread scheduler.

If you get the error `virtual memory exhausted` while compiling, you should try using `ulimit -v 80000` and running `make` again. If this does not work and you are using `bash`, try switching to `csh` or `sh`; some BSDI users have reported problems with `bash` and `ulimit`.

### 2.13.4.6. BSD/OS Version 4.x Notes

BSDI 4.x has some thread-related bugs. If you want to use MySQL on this, you should install all thread-related patches. At least M400-023 should be installed.

On some BSDI 4.x systems, you may get problems with shared libraries. The symptom is that you cannot execute any client programs, for example, `mysqladmin`. In this case, you need to reconfigure not to use shared libraries with the `--disable-shared` option to `configure`.

Some customers have had problems on BSDI 4.0.1 that the `mysqld` binary after a while cannot open tables. This occurs because some library/system-related bug causes `mysqld` to change current directory without having asked for that to happen.

The fix is to either upgrade MySQL to at least version 3.23.34 or, after running `configure`, remove the line `#define HAVE_REALPATH` from `config.h` before running `make`.

Note that this means that you cannot symbolically link a database directories to another database directory or symbolic link a table to another database on BSDI. (Making a symbolic link to another disk is okay).

## 2.13.5. Other Unix Notes

### 2.13.5.1. HP-UX Version 10.20 Notes

There are a couple of small problems when compiling MySQL on HP-UX. We recommend that you use `gcc` instead of the HP-UX native compiler, because `gcc` produces better code.

We recommend using `gcc` 2.95 on HP-UX. Do not use high optimization flags (such as `-O6`) because they may not be safe on HP-UX.

The following `configure` line should work with `gcc` 2.95:

```
CFLAGS="-I/opt/dce/include -fpic" \
CXXFLAGS="-I/opt/dce/include -felide-constructors -fno-exceptions \
-fno-rtti" \
CXX=gcc \
./configure --with-pthread \
--with-named-thread-libs='-ldce' \
--prefix=/usr/local/mysql --disable-shared
```

The following `configure` line should work with `gcc` 3.1:

```
CFLAGS="-DHPUX -I/opt/dce/include -O3 -fPIC" CXX=gcc \
CXXFLAGS="-DHPUX -I/opt/dce/include -felide-constructors \
-fno-exceptions -fno-rtti -O3 -fPIC" \
./configure --prefix=/usr/local/mysql \
--with-extra-charsets=complex --enable-thread-safe-client \
--enable-local-infile --with-pthread \
--with-named-thread-libs=-ldce --with-lib-ccflags=-fPIC \
--disable-shared
```

## 2.13.5.2. HP-UX Version 11.x Notes

For HP-UX 11.x, we recommend MySQL 3.23.15 or later.

Because of some critical bugs in the standard HP-UX libraries, you should install the following patches before trying to run MySQL on HP-UX 11.0:

```
PHKL_22840 Streams cumulative
PHNE_22397 ARPA cumulative
```

This solves the problem of getting `EWOULDBLOCK` from `recv()` and `EBADF` from `accept()` in threaded applications.

If you are using `gcc` 2.95.1 on an unpatched HP-UX 11.x system, you may get the following error:

```
In file included from /usr/include/unistd.h:11,
                 from ../include/global.h:125,
                 from mysql_priv.h:15,
                 from item.cc:19:
/usr/include/sys/unistd.h:184: declaration of C function ...
/usr/include/sys/unistd.h:440: previous declaration ...
In file included from item.h:306,
                 from mysql_priv.h:158,
                 from item.cc:19:
```

The problem is that HP-UX does not define `pthread_atfork()` consistently. It has conflicting prototypes in `/usr/include/sys/unistd.h:184` and `/usr/include/sys/unistd.h:440`.

One solution is to copy `/usr/include/sys/unistd.h` into `mysql/include` and edit `unistd.h` and change it to match the definition in `pthread.h`. Look for this line:

```
extern int pthread_atfork(void (*prepare)(), void (*parent)(),
                        void (*child)());
```

Change it to look like this:

```
extern int pthread_atfork(void (*prepare)(void), void (*parent)(void),
                        void (*child)(void));
```

After making the change, the following `configure` line should work:

```
CFLAGS="-fomit-frame-pointer -O3 -fpic" CXX=gcc \
CXXFLAGS="-felide-constructors -fno-exceptions -fno-rtti -O3" \
./configure --prefix=/usr/local/mysql --disable-shared
```

If you are using MySQL 4.0.5 with the HP-UX compiler, you can use the following command (which has been tested with `cc` B.11.11.04):

```
CC=cc CXX=aCC CFLAGS=+DD64 CXXFLAGS=+DD64 ./configure \
--with-extra-character-set=complex
```

You can ignore any errors of the following type:

```
aCC: warning 901: unknown option: `-' : use +help for online
documentation
```

If you get the following error from `configure`, verify that you do not have the path to the K&R compiler before the path to the HP-UX C and C++ compiler:

```
checking for cc option to accept ANSI C... no
configure: error: MySQL requires an ANSI C compiler (and a C++ compiler).
Try gcc. See the Installation chapter in the Reference Manual.
```

Another reason for not being able to compile is that you did not define the `+DD64` flags as just described.

Another possibility for HP-UX 11 is to use MySQL binaries for HP-UX 10.20. We have received reports from some users that these binaries work fine on HP-UX 11.00. If you encounter problems, be sure to check your HP-UX patch level.

### 2.13.5.3. IBM-AIX notes

Automatic detection of `xlc` is missing from Autoconf, so a number of variables need to be set before running `configure`. The following example uses the IBM compiler:

```
export CC="xlc_r -ma -O3 -qstrict -qoptimize=3 -qmaxmem=8192 "
export CXX="xlc_r -ma -O3 -qstrict -qoptimize=3 -qmaxmem=8192 "
export CFLAGS="-I /usr/local/include"
export LDFLAGS="-L /usr/local/lib"
export CPPFLAGS=$CFLAGS
export CXXFLAGS=$CFLAGS

./configure --prefix=/usr/local \
            --localstatedir=/var/mysql \
            --sbindir='/usr/local/bin' \
            --libexecdir='/usr/local/bin' \
            --enable-thread-safe-client \
            --enable-large-files
```

The preceding options are used to compile the MySQL distribution that can be found at <http://www-frec.bull.com/>.

If you change the `-O3` to `-O2` in the preceding `configure` line, you must also remove the `-qstrict` option. This is a limitation in the IBM C compiler.

If you are using `gcc` or `egcs` to compile MySQL, you *must* use the `-fno-exceptions` flag, because the exception handling in `gcc/egcs` is not thread-safe! (This is tested with `egcs` 1.1.) There are also some known problems with IBM's assembler that may cause it to generate bad code when used with `gcc`.

We recommend the following `configure` line with `egcs` and `gcc` 2.95 on AIX:

```
CC="gcc -pipe -mcpu=power -Wa,-many" \
CXX="gcc -pipe -mcpu=power -Wa,-many" \
CXXFLAGS="-felide-constructors -fno-exceptions -fno-rtti" \
./configure --prefix=/usr/local/mysql --with-low-memory
```

The `-Wa,-many` option is necessary for the compile to be successful. IBM is aware of this problem but is in no hurry to fix it because of the workaround that is available. We do not know if the `-fno-exceptions` is required with `gcc` 2.95, but because MySQL does not use exceptions and the option generates faster code, we recommend that you should always use it with `egcs` and `gcc`.

If you get a problem with assembler code, try changing the `-mcpu=xxx` option to match your CPU. Typically `power2`, `power`, or `powerpc` may need to be used. Alternatively, you might need to use `604` or `604e`. We are not positive but suspect that `power` would likely be safe most of the time, even on a `power2` machine.

If you do not know which CPU is present, execute a `uname -m` command. It produces a string that looks like `000514676700` whose format is `xyyyyyyyymmss` where `xx` and `ss` are always `00`, `yyyyyy` is a unique system ID and `mm` is the ID of the CPU Planar. A chart of these values can be found at [http://www16.boulder.ibm.com/pseries/en\\_US/cmds/aixcmds5/uname.htm](http://www16.boulder.ibm.com/pseries/en_US/cmds/aixcmds5/uname.htm).

This gives you a machine type and model which you can use to determine what type of CPU you have.

If you have problems with signals (MySQL dies unexpectedly under high load), you may have found an OS bug with threads and signals. In this case, you can tell MySQL not to use signals by configuring as follows:

```
CFLAGS=-DDONT_USE_THR_ALARM CXX=gcc \  
CXXFLAGS="-felide-constructors -fno-exceptions -fno-rtti \  
-DDONT_USE_THR_ALARM" \  
./configure --prefix=/usr/local/mysql --with-debug \  
--with-low-memory
```

This does not affect the performance of MySQL, but has the side effect that you cannot kill clients that are “sleeping” on a connection with `mysqladmin kill` or `mysqladmin shutdown`. Instead, the client dies when it issues its next command.

On some versions of AIX, linking with `libbind.a` makes `getservbyname()` dump core. This is an AIX bug and should be reported to IBM.

For AIX 4.2.1 and `gcc`, you have to make the following changes.

After configuring, edit `config.h` and `include/my_config.h` and change the line that says this:

```
#define HAVE_SNPRINTF 1
```

to this:

```
#undef HAVE_SNPRINTF
```

And finally, in `mysqld.cc`, you need to add a prototype for `initgroups()`.

```
#ifdef _AIX41  
extern "C" int initgroups(const char *,int);  
#endif
```

If you need to allocate a lot of memory to the `mysqld` process, it is not sufficient merely to use `ulimit -d unlimited`. You may also have to modify `mysqld_safe`, adding a line something like this:

```
export LDR_CNTRL='MAXDATA=0x80000000'
```

You can find more information about using very large amounts of memory at [http://publib16.boulder.ibm.com/pseries/en\\_US/aixprgdd/genprog/lrg\\_prg\\_support.htm](http://publib16.boulder.ibm.com/pseries/en_US/aixprgdd/genprog/lrg_prg_support.htm).

Users of AIX 4.3 should use `gmake` instead of the `make` utility included with AIX.

## 2.13.5.4. SunOS 4 Notes

On SunOS 4, MIT-pthreads is needed to compile MySQL. This in turn means you need GNU `make`.

Some SunOS 4 systems have problems with dynamic libraries and `libtool`. You can use the following `configure` line to avoid this problem:

```
./configure --disable-shared --with-mysqld-ldflags=-all-static
```

When compiling `readline`, you may get warnings about duplicate defines. These can be ignored.

When compiling `mysqld`, there are some `implicit declaration of function` warnings.

---

These can be ignored.

### 2.13.5.5. Alpha-DEC-UNIX Notes (Tru64)

If you are using `egcs` 1.1.2 on Digital Unix, you should upgrade to `gcc` 2.95.2, because `egcs` on DEC has some serious bugs.

When compiling threaded programs under Digital Unix, the documentation recommends using the `-pthread` option for `cc` and `cxx` and the `-lmach -lexc` libraries (in addition to `-lpthread`). You should run `configure` something like this:

```
CC="cc -pthread" CXX="cxx -pthread -O" \
./configure --with-named-thread-libs="-lpthread -lmach -lexc -lc"
```

When compiling `mysqld`, you may see a couple of warnings like this:

```
mysqld.cc: In function void handle_connections():
mysqld.cc:626: passing long unsigned int *' as argument 3 of
accept(int,sockaddr *, int *)'
```

You can safely ignore these warnings. They occur because `configure` can detect only errors, not warnings.

If you start the server directly from the command line, you may have problems with it dying when you log out. (When you log out, your outstanding processes receive a `SIGHUP` signal.) If so, try starting the server like this:

```
nohup mysqld [options] &
```

`nohup` causes the command following it to ignore any `SIGHUP` signal sent from the terminal. Alternatively, start the server by running `mysqld_safe`, which invokes `mysqld` using `nohup` for you. See [Section 5.4.1, “mysqld\\_safe — MySQL Server Startup Script”](#).

If you get a problem when compiling `mysys/get_opt.c`, just remove the `#define _NO_PROTO` line from the start of that file.

If you are using Compaq's CC compiler, the following `configure` line should work:

```
CC="cc -pthread"
CFLAGS="-O4 -ansi_alias -ansi_args -fast -inline speed \
        -speculate all -arch host"
CXX="cxx -pthread"
CXXFLAGS="-O4 -ansi_alias -ansi_args -fast -inline speed \
        -speculate all -arch host -noexceptions -nortti"
export CC CFLAGS CXX CXXFLAGS
./configure \
  --prefix=/usr/local/mysql \
  --with-low-memory \
  --enable-large-files \
  --enable-shared=yes \
  --with-named-thread-libs="-lpthread -lmach -lexc -lc"
gnumake
```

If you get a problem with `libtool` when compiling with shared libraries as just shown, when linking `mysql`, you should be able to get around this by issuing these commands:

```
cd mysql
/bin/sh ../libtool --mode=link cxx -pthread -O3 -DDEBUG_OFF \
  -O4 -ansi_alias -ansi_args -fast -inline speed \
  -speculate all \ -arch host -DUNDEF_HAVE_GETHOSTBYNAME_R \
  -o mysql mysql.o readline.o sql_string.o completion_hash.o \
  ../readline/libreadline.a -lcurses \
  ../libmysql/.libs/libmysqlclient.so -lm
cd ..
gnumake
```

```
gnumake install
scripts/mysql_install_db
```

### 2.13.5.6. Alpha-DEC-OSF/1 Notes

If you have problems compiling and have DEC `CC` and `gcc` installed, try running `configure` like this:

```
CC=cc CFLAGS=-O CXX=gcc CXXFLAGS=-O3 \
./configure --prefix=/usr/local/mysql
```

If you get problems with the `c_asm.h` file, you can create and use a 'dummy' `c_asm.h` file with:

```
touch include/c_asm.h
CC=gcc CFLAGS=-I./include \
CXX=gcc CXXFLAGS=-O3 \
./configure --prefix=/usr/local/mysql
```

Note that the following problems with the `ld` program can be fixed by downloading the latest DEC (Compaq) patch kit from: <http://ftp.support.compaq.com/public/unix/>.

On OSF/1 V4.0D and compiler "DEC C V5.6-071 on Digital Unix V4.0 (Rev. 878)," the compiler had some strange behavior (undefined `asm` symbols). `/bin/ld` also appears to be broken (problems with `_exit` undefined errors occurring while linking `mysqld`). On this system, we have managed to compile MySQL with the following `configure` line, after replacing `/bin/ld` with the version from OSF 4.0C:

```
CC=gcc CXX=gcc CXXFLAGS=-O3 ./configure --prefix=/usr/local/mysql
```

With the Digital compiler "C++ V6.1-029," the following should work:

```
CC=cc -pthread
CFLAGS=-O4 -ansi_alias -ansi_args -fast -inline speed \
    -speculate all -arch host
CXX=cxx -pthread
CXXFLAGS=-O4 -ansi_alias -ansi_args -fast -inline speed \
    -speculate all -arch host -noexceptions -nortti
export CC CFLAGS CXX CXXFLAGS
./configure --prefix=/usr/mysql/mysql \
    --with-mysqld-ldflags=-all-static --disable-shared \
    --with-named-thread-libs="-lmach -lexc -lc"
```

In some versions of OSF/1, the `alloca()` function is broken. Fix this by removing the line in `config.h` that defines `'HAVE_ALLOCA'`.

The `alloca()` function also may have an incorrect prototype in `/usr/include/alloca.h`. This warning resulting from this can be ignored.

`configure` uses the following thread libraries automatically: -  
`--with-named-thread-libs="-lpthread -lmach -lexc -lc"`.

When using `gcc`, you can also try running `configure` like this:

```
CFLAGS=-D_PTHREAD_USE_D4 CXX=gcc CXXFLAGS=-O3 ./configure ...
```

If you have problems with signals (MySQL dies unexpectedly under high load), you may have found an OS bug with threads and signals. In this case, you can tell MySQL not to use signals by configuring with:

```
CFLAGS=-DDONT_USE_THR_ALARM \
CXXFLAGS=-DDONT_USE_THR_ALARM \
```

```
./configure ...
```

This does not affect the performance of MySQL, but has the side effect that you cannot kill clients that are “sleeping” on a connection with `mysqladmin kill` or `mysqladmin shutdown`. Instead, the client dies when it issues its next command.

With `gcc 2.95.2`, you may encounter the following compile error:

```
sql_acl.cc:1456: Internal compiler error in `scan_region',
at except.c:2566
Please submit a full bug report.
```

To fix this, you should change to the `sql` directory and do a cut-and-paste of the last `gcc` line, but change `-O3` to `-O0` (or add `-O0` immediately after `gcc` if you do not have any `-O` option on your compile line). After this is done, you can just change back to the top-level directory and run `make` again.

## 2.13.5.7. SGI Irix Notes

If you are using Irix 6.5.3 or newer, `mysqld` is able to create threads only if you run it as a user that has `CAP_SCHED_MGT` privileges (such as `root`) or if you give the `mysqld` server this privilege with the following shell command:

```
chcap "CAP_SCHED_MGT+epi" /opt/mysql/libexec/mysqld
```

You may have to undefine some symbols in `config.h` after running `configure` and before compiling.

In some Irix implementations, the `alloca()` function is broken. If the `mysqld` server dies on some `SELECT` statements, remove the lines from `config.h` that define `HAVE_ALLOC` and `HAVE_ALLOCA_H`. If `mysqladmin create` does not work, remove the line from `config.h` that defines `HAVE_READDIR_R`. You may have to remove the `HAVE_TERM_H` line as well.

SGI recommends that you install all the patches on this page as a set:

[http://support.sgi.com/surfzone/patches/patchset/6.2\\_indigo.rps.html](http://support.sgi.com/surfzone/patches/patchset/6.2_indigo.rps.html)

At the very minimum, you should install the latest kernel rollup, the latest `rld` rollup, and the latest `libc` rollup.

You definitely need all the POSIX patches on this page, for pthreads support:

[http://support.sgi.com/surfzone/patches/patchset/6.2\\_posix.rps.html](http://support.sgi.com/surfzone/patches/patchset/6.2_posix.rps.html)

If you get the something like the following error when compiling `mysql.cc`:

```
"/usr/include/curses.h", line 82: error(1084):
invalid combination of type
```

Type the following in the top-level directory of your MySQL source tree:

```
extra/replace bool curses_bool < /usr/include/curses.h > include/curses.h
make
```

There have also been reports of scheduling problems. If only one thread is running, performance is slow. Avoid this by starting another client. This may lead to a two-to-tenfold increase in execution speed thereafter for the other thread. This is a poorly understood problem with Irix threads; you may have to improvise to find solutions until this can be fixed.



If you are compiling with `gcc`, you can use the following `configure` command:

```
CC=gcc CXX=gcc CXXFLAGS=-O3 \
./configure --prefix=/usr/local/mysql --enable-thread-safe-client \
--with-named-thread-libs=-lpthread
```

On Irix 6.5.11 with native Irix C and C++ compilers ver. 7.3.1.2, the following is reported to work

```
CC=cc CXX=CC CFLAGS='-O3 -n32 -TARG:platform=IP22 -I/usr/local/include \
-L/usr/local/lib' CXXFLAGS='-O3 -n32 -TARG:platform=IP22 \
-I/usr/local/include -L/usr/local/lib' \
./configure --prefix=/usr/local/mysql --with-innodb --with-berkeley-db \
--with-libwrap=/usr/local \
--with-named-curses-libs=/usr/local/lib/libncurses.a
```

### 2.13.5.8. SCO UNIX and OpenServer 5.0.x Notes

The current port is tested only on `sco3.2v5.0.5`, `sco3.2v5.0.6`, and `sco3.2v5.0.7` systems. There has also been progress on a port to `sco3.2v4.2`. Open Server 5.0.8 (Legend) has native threads and allows files greater than 2GB. The current maximum file size is 2GB.

We have been able to compile MySQL with the following `configure` command on OpenServer with `gcc 2.95.3`.

```
CC=gcc CFLAGS="-D_FILE_OFFSET_BITS=64 -O3" \
CXX=gcc CXXFLAGS="-D_FILE_OFFSET_BITS=64 -O3" \
./configure --prefix=/usr/local/mysql \
--enable-thread-safe-client --with-innodb \
--with-openssl --with-vio --with-extra-charsets=complex
```

`gcc` is available at <ftp://ftp.sco.com/pub/openserver5/opensrc/gnutools-5.0.7Kj>.

This development system requires the OpenServer Execution Environment Supplement oss646B on OpenServer 5.0.6 and oss656B and The OpenSource libraries found in `gwlibs`. All OpenSource tools are in the `opensrc` directory. They are available at <ftp://ftp.sco.com/pub/openserver5/opensrc/>.

We recommend using the latest production release of MySQL.

SCO provides operating system patches at <ftp://ftp.sco.com/pub/openserver5> for OpenServer 5.0.[0-6] and <ftp://ftp.sco.com/pub/openserverv5/507> for OpenServer 5.0.7.

SCO provides information about security fixes at <ftp://ftp.sco.com/pub/security/OpenServer> for OpenServer 5.0.x.

The maximum file size on an OpenServer 5.0.x system is 2GB.

The total memory which can be allocated for streams buffers, clists, and lock records cannot exceed 60MB on OpenServer 5.0.x.

Streams buffers are allocated in units of 4096 byte pages, clists are 70 bytes each, and lock records are 64 bytes each, so:

```
(NSTRPAGES × 4096) + (NCLIST × 70) + (MAX_FLCKREC × 64) <= 62914560
```

Follow this procedure to configure the Database Services option. If you are unsure whether an application requires this, see the documentation provided with the application.

1. Log in as `root`.

2. Enable the SUDS driver by editing the `/etc/conf/sdevice.d/suds` file. Change the `N` in the second field to a `Y`.
3. Use `mkdev aio` or the Hardware/Kernel Manager to enable support for asynchronous I/O and re-link the kernel. To allow users to lock down memory for use with this type of I/O, update the `aiomemlock(F)` file. This file should be updated to include the names of users that can use AIO and the maximum amounts of memory they can lock down.
4. Many applications use `setuid` binaries so that you need to specify only a single user. See the documentation provided with the application to determine whether this is the case for your application.

After you complete this process, reboot the system to create a new kernel incorporating these changes.

By default, the entries in `/etc/conf/cf.d/mtune` are set as follows:

Value	Default	Min	Max
-----	-----	---	---
NBUF	0	24	450000
NHBUF	0	32	524288
NMPBUF	0	12	512
MAX_INODE	0	100	64000
MAX_FILE	0	100	64000
CTBUFSIZE	128	0	256
MAX_PROC	0	50	16000
MAX_REGION	0	500	160000
NCLIST	170	120	16640
MAXUP	100	15	16000
NOFILES	110	60	11000
NHINODE	128	64	8192
NAUTOUP	10	0	60
NGROUPS	8	0	128
BDFLUSHR	30	1	300
MAX_FLCKREC	0	50	16000
PUTBUFSZ	8000	2000	20000
MAXSLICE	100	25	100
ULIMIT	4194303	2048	4194303
* Streams Parameters			
NSTREAM	64	1	32768
NSTRPUSH	9	9	9
NMUXLINK	192	1	4096
STRMSGSZ	16384	4096	524288
STRCTLSZ	1024	1024	1024
STRMAXBLK	524288	4096	524288
NSTRPAGES	500	0	8000
STRSPLITFRAC	80	50	100
NLOG	3	3	3
NUMSP	64	1	256
NUMTIM	16	1	8192
NUMTRW	16	1	8192
* Semaphore Parameters			
SEMMAP	10	10	8192
SEMMNI	10	10	8192
SEMMNS	60	60	8192
SEMMNU	30	10	8192
SEMMSL	25	25	150
SEMOPM	10	10	1024
SEMUME	10	10	25
SEVMVMX	32767	32767	32767
SEMAEM	16384	16384	16384
* Shared Memory Parameters			
SHMMAX	524288	131072	2147483647
SHMMIN	1	1	1
SHMNI	100	100	2000
FILE	0	100	64000
NMOUNT	0	4	256
NPROC	0	50	16000
NREGION	0	500	160000

We recommend setting these values as follows:

- `NOFILES` should be 4096 or 2048.
- `MAXUP` should be 2048.

To make changes to the kernel, use the `idtune name parameter` command. `idtune` modifies the `/etc/conf/cf.d/stune` file for you. For example, to change `SEMMS` to `200`, execute this command as `root`:

```
# /etc/conf/bin/idtune SEMMS 200
```

Then rebuild and reboot the kernel by issuing this command:

```
# /etc/conf/bin/idbuild -B && init 6
```

We recommend tuning the system, but the proper parameter values to use depend on the number of users accessing the application or database and size of the database (that is, the used buffer pool). The following kernel parameters can be set with `idtune`:

- `SHMMAX` (recommended setting: 128MB) and `SHMSEG` (recommended setting: 15). These parameters have an influence on the MySQL database engine to create user buffer pools.
- `NOFILES` and `MAXUP` should be set to at least 2048.
- `MAXPROC` should be set to at least 3000/4000 (depends on number of users) or more.
- We also recommend using the following formulas to calculate values for `SEMMSL`, `SEMMS`, and `SEMNU`:

```
SEMMSL = 13
```

13 is what has been found to be the best for both Progress and MySQL.

```
SEMMS = SEMMSL × number of db servers to be run on the system
```

Set `SEMMS` to the value of `SEMMSL` multiplied by the number of database servers (maximum) that you are running on the system at one time.

```
SEMNU = SEMMS
```

Set the value of `SEMNU` to equal the value of `SEMMS`. You could probably set this to 75% of `SEMMS`, but this is a conservative estimate.

You need to at least install the SCO OpenServer Linker and Application Development Libraries or the OpenServer Development System to use `gcc`. You cannot use the GCC Dev system without installing one of these.

You should get the FSU Pthreads package and install it first. This can be found at <http://moss.csc.ncsu.edu/~mueller/ftp/pub/PART/pthreads.tar.gz>. You can also get a precompiled package from [ftp://ftp.zenez.com/pub/zenez/prgms/FSU-threads-3.14.tar.gz](http://ftp.zenez.com/pub/zenez/prgms/FSU-threads-3.14.tar.gz).

FSU Pthreads can be compiled with SCO Unix 4.2 with `tcpip`, or using OpenServer 3.0 or Open Desktop 3.0 (OS 3.0 ODT 3.0) with the SCO Development System installed using a good port of GCC 2.5.x. For ODT or OS 3.0, you need a good port of GCC 2.5.x. There are a lot of problems without a good port. The port for this product requires the SCO Unix Development system. Without it, you are missing the

libraries and the linker that is needed. You also need [SCO-3.2v4.2-includes.tar.gz](ftp://ftp.zenez.com/pub/zenez/prgms/SCO-3.2v4.2-includes.tar.gz). This file contains the changes to the SCO Development include files that are needed to get MySQL to build. You need to replace the existing system include files with these modified header files. They can be obtained from <ftp://ftp.zenez.com/pub/zenez/prgms/SCO-3.2v4.2-includes.tar.gz>.

To build FSU Pthreads on your system, all you should need to do is run GNU `make`. The `Makefile` in `FSU-threads-3.14.tar.gz` is set up to make FSU-threads.

You can run `./configure` in the `threads/src` directory and select the SCO OpenServer option. This command copies `Makefile.SCO5` to `Makefile`. Then run `make`.

To install in the default `/usr/include` directory, log in as `root`, and then `cd` to the `thread/src` directory and run `make install`.

Remember that you must use GNU `make` to build MySQL.

**Note:** If you don't start `mysqld_safe` as `root`, you should get only the default 110 open files per process. `mysqld` writes a note about this in the log file.

With SCO 3.2V4.2, you should use FSU Pthreads version 3.14 or newer. The following `configure` command should work:

```
CFLAGS="-D_XOPEN_XPG4" CXX=gcc CXXFLAGS="-D_XOPEN_XPG4" \  
./configure \  
  --prefix=/usr/local/mysql \  
  --with-named-thread-libs="-lgthreads -lsocket -lgen -lgthreads" \  
  --with-named-curses-libs="-lcurses"
```

You may have problems with some include files. In this case, you can find new SCO-specific include files at <ftp://ftp.zenez.com/pub/zenez/prgms/SCO-3.2v4.2-includes.tar.gz>.

You should unpack this file in the `include` directory of your MySQL source tree.

SCO development notes:

- MySQL should automatically detect FSU Pthreads and link `mysqld` with `-lgthreads -lsocket -lgthreads`.
- The SCO development libraries are re-entrant in FSU Pthreads. SCO claims that its library functions are re-entrant, so they must be re-entrant with FSU Pthreads. FSU Pthreads on OpenServer tries to use the SCO scheme to make re-entrant libraries.
- FSU Pthreads (at least the version at <ftp://ftp.zenez.com>) comes linked with GNU `malloc`. If you encounter problems with memory usage, make sure that `gmalloc.o` is included in `libgthreads.a` and `libgthreads.so`.
- In FSU Pthreads, the following system calls are pthreads-aware: `read()`, `write()`, `getmsg()`, `connect()`, `accept()`, `select()`, and `wait()`.
- The CSSA-2001-SCO.35.2 (the patch is listed in custom as `erg711905-dscr_remap` security patch (version 2.0.0)) breaks FSU threads and makes `mysqld` unstable. You have to remove this one if you want to run `mysqld` on an OpenServer 5.0.6 machine.
- If you use SCO OpenServer 5, you may need to recompile FSU pthreads with `-DDRAFT7` in `CFLAGS`. Otherwise, `InnoDB` may hang at a `mysqld` startup.
- SCO provides operating system patches at <ftp://ftp.sco.com/pub/openserver5> for OpenServer 5.0.x.
- SCO provides security fixes and `libsocket.so.2` at <ftp://ftp.sco.com/pub/security/OpenServer>

and <ftp://ftp.sco.com/pub/security/sse> for OpenServer 5.0.x.

- Pre-OSR506 security fixes. Also, the `telnetd` fix at <ftp://stage.caldera.com/pub/security/openserver/> or <ftp://stage.caldera.com/pub/security/openserver/CSSA-2001-SCO.10/> as both `libsocket.so.2` and `libresolv.so.1` with instructions for installing on pre-OSR506 systems.

It's probably a good idea to install these patches before trying to compile/use MySQL.

Beginning with Legend/OpenServer 6.0.0, there are native threads and no 2GB file size limit.

### 2.13.5.9. SCO OpenServer 6.0.x Notes

OpenServer 6 includes these key improvements:

- Larger file support up to 1 TB
- Multiprocessor support increased from 4 to 32 processors
- Increased memory support up to 64GB
- Extending the power of UnixWare into OpenServer 6
- Dramatic performance improvement

OpenServer 6.0.0 commands are organized as follows:

- `/bin` is for commands that behave exactly the same as on OpenServer 5.0.x.
- `/u95/bin` is for commands that have better standards conformance, for example Large File System (LFS) support.
- `/udk/bin` is for commands that behave the same as on UnixWare 7.1.4. The default is for the LFS support.

The following is a guide to setting `PATH` on OpenServer 6. If the user wants the traditional OpenServer 5.0.x then `PATH` should be `/bin` first. If the user wants LFS support, the path should be `/u95/bin:/bin`. If the user wants UnixWare 7 support first, the path would be `/udk/bin:/u95/bin:/bin:`.

We recommend using the latest production release of MySQL. Should you choose to use an older release of MySQL on OpenServer 6.0.x, you must use a version of MySQL at least as recent as 3.22.13 to get fixes for some portability and OS problems.

MySQL distribution files with names of the following form are `tar` archives of media are `tar` archives of media images suitable for installation with the SCO Software Manager (`/etc/custom`) on SCO OpenServer 6:

```
mysql-PRODUCT-4.1.22-sco-osr6-i686.VOLS.tar
```

A distribution where `PRODUCT` is `pro-cert` is the Commercially licensed MySQL Pro Certified server. A distribution where `PRODUCT` is `pro-gpl-cert` is the MySQL Pro Certified server licensed under the terms of the General Public License (GPL).

Select whichever distribution you wish to install and, after download, extract the `tar` archive into an empty directory. For example:

```
shell> mkdir /tmp/mysql-pro
shell> cd /tmp/mysql-pro
shell> tar xf /tmp/mysql-pro-cert-4.1.22-sco-osr6-i686.VOLS.tar
```

Prior to installation, back up your data in accordance with the procedures outlined in [Section 2.11](#), “Upgrading MySQL”.

Remove any previously installed `pkgadd` version of MySQL:

```
shell> pkginfo mysql 2>&1 > /dev/null && pkgrm mysql
```

Install MySQL Pro from media images using the SCO Software Manager:

```
shell> /etc/custom -p SCO:MySQL -i -z /tmp/mysql-pro
```

Alternatively, the SCO Software Manager can be displayed graphically by clicking on the [Software Manager](#) icon on the desktop, selecting [Software -> Install New](#), selecting the host, selecting [Media Images](#) for the Media Device, and entering `/tmp/mysql-pro` as the Image Directory.

After installation, run `mkdev mysql` as the `root` user to configure your newly installed MySQL Pro Certified server.

**Note:** The installation procedure for VOLS packages does not create the `mysql` user and group that the package uses by default. You should either create the `mysql` user and group, or else select a different user and group using an option in `mkdev mysql`.

If you wish to configure your MySQL Pro server to interface with the Apache Web server via PHP, download and install the PHP update from SCO at <http://ftp.sco.com/pub/updates/OpenServer/SCOSA-2006.17/>.

We have been able to compile MySQL with the following `configure` command on OpenServer 6.0.x:

```
CC=cc CFLAGS="-D_FILE_OFFSET_BITS=64 -O3" \
CXX=CC CXXFLAGS="-D_FILE_OFFSET_BITS=64 -O3" \
./configure --prefix=/usr/local/mysql \
--enable-thread-safe-client --with-berkeley-db \
--with-extra-charsets=complex \
--build=i686-unknown-sysv5SCO_SV6.0.0
```

If you use `gcc`, you must use `gcc` 2.95.3 or newer.

```
CC=gcc CXX=g++ ... ./configure ...
```

The version of Berkeley DB that comes with either UnixWare 7.1.4 or OpenServer 6.0.0 is not used when building MySQL. MySQL instead uses its own version of Berkeley DB. The `configure` command needs to build both a static and a dynamic library in `src_directory/bdb/build_unix/`, but it does not with MySQL's own [BDB](#) version. The workaround is as follows.

1. Configure as normal for MySQL.
2. `cd bdb/build_unix/`
3. `cp -p Makefile Makefile.sav`
4. Use same options and run `../dist/configure`.

5. Run `gmake`.
6. `cp -p Makefile.sav Makefile`
7. Change location to the top source directory and run `gmake`.

This allows both the shared and dynamic libraries to be made and work.

SCO provides OpenServer 6 operating system patches at <ftp://ftp.sco.com/pub/openserver6>.

SCO provides information about security fixes at <ftp://ftp.sco.com/pub/security/OpenServer>.

By default, the maximum file size on a OpenServer 6.0.0 system is 1TB. Some operating system utilities have a limitation of 2GB. The maximum possible file size on UnixWare 7 is 1TB with VXFS or HTFS.

OpenServer 6 can be configured for large file support (file sizes greater than 2GB) by tuning the UNIX kernel.

By default, the entries in `/etc/conf/cf.d/mtune` are set as follows:

Value	Default	Min	Max
SVMLIM	0x9000000	0x1000000	0x7FFFFFFF
HVMLIM	0x9000000	0x1000000	0x7FFFFFFF

To make changes to the kernel, use the `idtune name parameter` command. `idtune` modifies the `/etc/conf/cf.d/stune` file for you. We recommend setting the kernel values by executing the following commands as `root`:

```
# /etc/conf/bin/idtune SDATLIM 0x7FFFFFFF
# /etc/conf/bin/idtune HDATLIM 0x7FFFFFFF
# /etc/conf/bin/idtune SVMLIM 0x7FFFFFFF
# /etc/conf/bin/idtune HVMLIM 0x7FFFFFFF
# /etc/conf/bin/idtune SFNOLIM 2048
# /etc/conf/bin/idtune HFNOLIM 2048
```

Then rebuild and reboot the kernel by issuing this command:

```
# /etc/conf/bin/idbuild -B && init 6
```

We recommend tuning the system, but the proper parameter values to use depend on the number of users accessing the application or database and size of the database (that is, the used buffer pool). The following kernel parameters can be set with `idtune`:

- `SHMMAX` (recommended setting: 128MB) and `SHMSEG` (recommended setting: 15). These parameters have an influence on the MySQL database engine to create user buffer pools.
- `SFNOLIM` and `HFNOLIM` should be at maximum 2048.
- `NPROC` should be set to at least 3000/4000 (depends on number of users).
- We also recommend using the following formulas to calculate values for `SEMMSL`, `SEMMNS`, and `SEMMLIM`:

```
SEMMSL = 13
```

13 is what has been found to be the best for both Progress and MySQL.

```
SEMMNS = SEMMSL × number of db servers to be run on the system
```

Set `SEMMNS` to the value of `SEMMSL` multiplied by the number of database servers (maximum) that you are running on the system at one time.

```
SEMMNU = SEMMNS
```

Set the value of `SEMMNU` to equal the value of `SEMMNS`. You could probably set this to 75% of `SEMMNS`, but this is a conservative estimate.

## 2.13.5.10. SCO UnixWare 7.1.x and OpenUNIX 8.0.0 Notes

We recommend using the latest production release of MySQL. Should you choose to use an older release of MySQL on UnixWare 7.1.x, you must use a version of MySQL at least as recent as 3.22.13 to get fixes for some portability and OS problems.

We have been able to compile MySQL with the following `configure` command on UnixWare 7.1.x:

```
CC="cc" CFLAGS="-I/usr/local/include" \
CXX="CC" CXXFLAGS="-I/usr/local/include" \
./configure --prefix=/usr/local/mysql \
--enable-thread-safe-client --with-berkeley-db=./bdb \
--with-innodb --with-openssl --with-extra-charsets=complex
```

If you want to use `gcc`, you must use `gcc 2.95.3` or newer.

```
CC=gcc CXX=g++ ... ./configure ...
```

The version of Berkeley DB that comes with either UnixWare 7.1.4 or OpenServer 6.0.0 is not used when building MySQL. MySQL instead uses its own version of Berkeley DB. The `configure` command needs to build both a static and a dynamic library in `src_directory/bdb/build_unix/`, but it does not with MySQL's own `BDB` version. The workaround is as follows.

1. Configure as normal for MySQL.
2. `cd bdb/build_unix/`
3. `cp -p Makefile Makefile.sav`
4. Use same options and run `../dist/configure`.
5. Run `gmake`.
6. `cp -p Makefile.sav Makefile`
7. Change to top source directory and run `gmake`.

This allows both the shared and dynamic libraries to be made and work.

SCO provides operating system patches at <ftp://ftp.sco.com/pub/unixware7> for UnixWare 7.1.1, <ftp://ftp.sco.com/pub/unixware7/713/> for UnixWare 7.1.3, <ftp://ftp.sco.com/pub/unixware7/714/> for UnixWare 7.1.4, and <ftp://ftp.sco.com/pub/openunix8> for OpenUNIX 8.0.0.

SCO provides information about security fixes at <ftp://ftp.sco.com/pub/security/OpenUNIX> for OpenUNIX and <ftp://ftp.sco.com/pub/security/UnixWare> for UnixWare.



The UnixWare 7 file size limit is 1 TB with VXFS. Some OS utilities have a limitation of 2GB.

On UnixWare 7.1.4 you do not need to do anything to get large file support, but to enable large file support on prior versions of UnixWare 7.1.x, run `fsadm`.

```
# fsadm -Fvxfs -o largefiles /
# fsadm / * Note
# ulimit unlimited
# /etc/conf/bin/idtune SFSZLIM 0x7FFFFFFF ** Note
# /etc/conf/bin/idtune HFSZLIM 0x7FFFFFFF ** Note
# /etc/conf/bin/idbuild -B

* This should report "largefiles".
** 0x7FFFFFFF represents infinity for these values.
```

Reboot the system using `shutdown`.

By default, the entries in `/etc/conf/cf.d/mtune` are set as follows:

Value	Default	Min	Max
SVMMLIM	0x9000000	0x1000000	0x7FFFFFFF
HVMMLIM	0x9000000	0x1000000	0x7FFFFFFF

To make changes to the kernel, use the `idtune name parameter` command. `idtune` modifies the `/etc/conf/cf.d/stune` file for you. We recommend setting the kernel values by executing the following commands as `root`:

```
# /etc/conf/bin/idtune SDATLIM 0x7FFFFFFF
# /etc/conf/bin/idtune HDATLIM 0x7FFFFFFF
# /etc/conf/bin/idtune SVMMLIM 0x7FFFFFFF
# /etc/conf/bin/idtune HVMMLIM 0x7FFFFFFF
# /etc/conf/bin/idtune SFNOLIM 2048
# /etc/conf/bin/idtune HFNOLIM 2048
```

Then rebuild and reboot the kernel by issuing this command:

```
# /etc/conf/bin/idbuild -B && init 6
```

We recommend tuning the system, but the proper parameter values to use depend on the number of users accessing the application or database and size of the database (that is, the used buffer pool). The following kernel parameters can be set with `idtune`:

- `SHMMAX` (recommended setting: 128MB) and `SHMSEG` (recommended setting: 15). These parameters have an influence on the MySQL database engine to create user buffer pools.
- `SFNOLIM` and `HFNOLIM` should be at maximum 2048.
- `NPROC` should be set to at least 3000/4000 (depends on number of users).
- We also recommend using the following formulas to calculate values for `SEMMSL`, `SEMMNS`, and `SEMMLU`:

```
SEMMSL = 13
```

13 is what has been found to be the best for both Progress and MySQL.

```
SEMMNS = SEMMSL × number of db servers to be run on the system
```

Set `SEMMNS` to the value of `SEMMSL` multiplied by the number of database servers (maximum) that

you are running on the system at one time.

```
SEMMNU = SEMMNS
```

Set the value of `SEMMNU` to equal the value of `SEMMNS`. You could probably set this to 75% of `SEMMNS`, but this is a conservative estimate.

## 2.13.6. OS/2 Notes

MySQL uses quite a few open files. Because of this, you should add something like the following to your `CONFIG.SYS` file:

```
SET EMXOPT=-c -n -h1024
```

If you do not do this, you may encounter the following error:

```
File 'xxxx' not found (Errcode: 24)
```

When using MySQL with OS/2 Warp 3, FixPack 29 or above is required. With OS/2 Warp 4, FixPack 4 or above is required. This is a requirement of the Pthreads library. MySQL must be installed on a partition with a type that supports long filenames, such as HPFS, FAT32, and so on.

The `INSTALL.COMD` script must be run from OS/2's own `CMD.EXE` and may not work with replacement shells such as `4OS2.EXE`.

The `scripts/mysql-install-db` script has been renamed. It is called `install.cmd` and is a REXX script, which sets up the default MySQL security settings and creates the Workplace Shell icons for MySQL.

Dynamic module support is compiled in but not fully tested. Dynamic modules should be compiled using the Pthreads runtime library.

```
gcc -zdll -zmt -Zcrt.dll=pthrdrt1 -I../include -I../regex -I.. \
-o example udf_example.cc -L../lib -lmysqlclient udf_example.def
mv example.dll example.udf
```

**Note:** Due to limitations in OS/2, UDF module name stems must not exceed eight characters. Modules are stored in the `/mysql2/udf` directory; the `safe-mysqld.cmd` script puts this directory in the `BEGINLIBPATH` environment variable. When using UDF modules, specified extensions are ignored--it is assumed to be `.udf`. For example, in Unix, the shared module might be named `example.so` and you would load a function from it like this:

```
mysql> CREATE FUNCTION metaphon RETURNS STRING SONAME 'example.so';
```

In OS/2, the module would be named `example.udf`, but you would not specify the module extension:

```
mysql> CREATE FUNCTION metaphon RETURNS STRING SONAME 'example';
```

## 2.14. Perl Installation Notes

Perl support for MySQL is provided by means of the `DBI/DBD` client interface. The interface requires Perl 5.6.1 or later. It *does not work* if you have an older version of Perl.

If you want to use transactions with Perl DBI, you need to have `DBD: :mysql` version 1.2216 or newer. `DBD: :mysql` 2.9003 or newer is recommended.

If you are using the MySQL 4.1 or newer client library, you must use `DBD::mysql` 2.9003 or newer.

As of MySQL 3.22.8, Perl support is no longer included with MySQL distributions. You can obtain the necessary modules from <http://search.cpan.org> for Unix, or by using the ActiveState `ppm` program on Windows. The following sections describe how to do this.

Perl support for MySQL must be installed if you want to run the MySQL benchmark scripts. See [Section 7.1.4, “The MySQL Benchmark Suite”](#).

## 2.14.1. Installing Perl on Unix

MySQL Perl support requires that you have installed MySQL client programming support (libraries and header files). Most installation methods install the necessary files. However, if you installed MySQL from RPM files on Linux, be sure that you've installed the developer RPM. The client programs are in the client RPM, but client programming support is in the developer RPM.

If you want to install Perl support, the files you need can be obtained from the CPAN (Comprehensive Perl Archive Network) at <http://search.cpan.org>.

The easiest way to install Perl modules on Unix is to use the `CPAN` module. For example:

```
shell> perl -MCPAN -e shell
cpan> install DBI
cpan> install DBD::mysql
```

The `DBD::mysql` installation runs a number of tests. These tests attempt to connect to the local MySQL server using the default username and password. (The default username is your login name on Unix, and `ODBC` on Windows. The default password is “no password.”) If you cannot connect to the server with those values (for example, if your account has a password), the tests fail. You can use `force install DBD::mysql` to ignore the failed tests.

`DBI` requires the `Data::Dumper` module. It may be installed; if not, you should install it before installing `DBI`.

It is also possible to download the module distributions in the form of compressed `tar` archives and build the modules manually. For example, to unpack and build a `DBI` distribution, use a procedure such as this:

1. Unpack the distribution into the current directory:

```
shell> gunzip < DBI-VERSION.tar.gz | tar xvf -
```

This command creates a directory named `DBI-VERSION`.

2. Change location into the top-level directory of the unpacked distribution:

```
shell> cd DBI-VERSION
```

3. Build the distribution and compile everything:

```
shell> perl Makefile.PL
shell> make
shell> make test
shell> make install
```

The `make test` command is important because it verifies that the module is working. Note that when

you run that command during the `DBD:mysql` installation to exercise the interface code, the MySQL server must be running or the test fails.

It is a good idea to rebuild and reinstall the `DBD:mysql` distribution whenever you install a new release of MySQL, particularly if you notice symptoms such as that all your `DBI` scripts fail after you upgrade MySQL.

If you do not have access rights to install Perl modules in the system directory or if you want to install local Perl modules, the following reference may be useful: <http://servers.digitaldaze.com/extensions/perl/modules.html#modules>

Look under the heading “Installing New Modules that Require Locally Installed Modules.”

## 2.14.2. Installing ActiveState Perl on Windows

On Windows, you should do the following to install the MySQL `DBD` module with ActiveState Perl:

1. Get ActiveState Perl from <http://www.activestate.com/Products/ActivePerl/> and install it.
2. Open a console window (a “DOS window”).
3. If necessary, set the `HTTP_proxy` variable. For example, you might try a setting like this:

```
set HTTP_proxy=my.proxy.com:3128
```

4. Start the PPM program:

```
C:\> C:\perl\bin\ppm.pl
```

5. If you have not previously done so, install `DBI`:

```
ppm> install DBI
```

6. If this succeeds, run the following command:

```
ppm> install \  
ftp://ftp.de.uu.net/pub/CPAN/authors/id/JWIED/DBD-mysql-1.2212.x86.ppd
```

This procedure should work with ActiveState Perl 5.6 or newer.

If you cannot get the procedure to work, you should install the MyODBC driver instead and connect to the MySQL server through ODBC:

```
use DBI;  
$dbh= DBI->connect("DBI:ODBC:$dsn",$user,$password) ||  
die "Got error $DBI::errstr when connecting to $dsn\n";
```

## 2.14.3. Problems Using the Perl `DBI/DBD` Interface

If Perl reports that it cannot find the `./mysql/mysql.so` module, the problem is probably that Perl cannot locate the `libmysqlclient.so` shared library. You should be able to fix this problem by one of the following methods:

- Compile the `DBD:mysql` distribution with `perl Makefile.PL -static -config` rather than `perl Makefile.PL`.

- Copy `libmysqlclient.so` to the directory where your other shared libraries are located (probably `/usr/lib` or `/lib`).
- Modify the `-L` options used to compile `DBD:mysql` to reflect the actual location of `libmysqlclient.so`.
- On Linux, you can add the pathname of the directory where `libmysqlclient.so` is located to the `/etc/ld.so.conf` file.
- Add the pathname of the directory where `libmysqlclient.so` is located to the `LD_RUN_PATH` environment variable. Some systems use `LD_LIBRARY_PATH` instead.

Note that you may also need to modify the `-L` options if there are other libraries that the linker fails to find. For example, if the linker cannot find `libc` because it is in `/lib` and the link command specifies `-L/usr/lib`, change the `-L` option to `-L/lib` or add `-L/lib` to the existing link command.

If you get the following errors from `DBD:mysql`, you are probably using `gcc` (or using an old binary compiled with `gcc`):

```
/usr/bin/perl: can't resolve symbol '__moddi3'
/usr/bin/perl: can't resolve symbol '__divdi3'
```

Add `-L/usr/lib/gcc-lib/... -lgcc` to the link command when the `mysql.so` library gets built (check the output from `make` for `mysql.so` when you compile the Perl client). The `-L` option should specify the pathname of the directory where `libgcc.a` is located on your system.

Another cause of this problem may be that Perl and MySQL are not both compiled with `gcc`. In this case, you can solve the mismatch by compiling both with `gcc`.

You may see the following error from `DBD:mysql` when you run the tests:

```
t/00base.....install_driver(mysql) failed:
Can't load './blib/arch/auto/DBD/mysql/mysql.so' for module DBD:mysql:
./blib/arch/auto/DBD/mysql/mysql.so: undefined symbol:
uncompress at /usr/lib/perl5/5.00503/i586-linux/DynaLoader.pm line 169.
```

This means that you need to include the `-lz` compression library on the link line. That can be done by changing the following line in the file `lib/DBD/mysql/Install.pm`:

```
$sysliblist .= " -lm";
```

Change that line to:

```
$sysliblist .= " -lm -lz";
```

After this, you *must* run `make realclean` and then proceed with the installation from the beginning.

If you want to install DBI on SCO, you have to edit the `Makefile` in `DBI-xxx` and each subdirectory. Note that the following assumes `gcc 2.95.2` or newer:

OLD:	NEW:
CC = cc	CC = gcc
CCDLFLAGS = -KPIC -Wl,-Bexport	CCDLFLAGS = -fpic
CCDLFLAGS = -wl,-Bexport	CCDLFLAGS =
LD = ld	LD = gcc -G -fpic
LDDLFLAGS = -G -L/usr/local/lib	LDDLFLAGS = -L/usr/local/lib
LDFLAGS = -belf -L/usr/local/lib	LDFLAGS = -L/usr/local/lib
LD = ld	LD = gcc -G -fpic
OPTIMISE = -Od	OPTIMISE = -O1

```
OLD:
CCCFLAGS = -belf -dy -w0 -U M_XENIX -DPERL_SCO5 -I/usr/local/include

NEW:
CCFLAGS = -U M_XENIX -DPERL_SCO5 -I/usr/local/include
```

These changes are necessary because the Perl dynaloader does not load the `DBI` modules if they were compiled with `icc` or `cc`.

If you want to use the Perl module on a system that does not support dynamic linking (such as SCO), you can generate a static version of Perl that includes `DBI` and `DBD::mysql`. The way this works is that you generate a version of Perl with the `DBI` code linked in and install it on top of your current Perl. Then you use that to build a version of Perl that additionally has the `DBD` code linked in, and install that.

On SCO, you must have the following environment variables set:

```
LD_LIBRARY_PATH=/lib:/usr/lib:/usr/local/lib:/usr/progressive/lib
```

Or:

```
LD_LIBRARY_PATH=/usr/lib:/lib:/usr/local/lib:/usr/ccs/lib:\
/usr/progressive/lib:/usr/skunk/lib
LIBPATH=/usr/lib:/lib:/usr/local/lib:/usr/ccs/lib:\
/usr/progressive/lib:/usr/skunk/lib
MANPATH=scohelp:/usr/man:/usr/local1/man:/usr/local/man:\
/usr/skunk/man:
```

First, create a Perl that includes a statically linked `DBI` module by running these commands in the directory where your `DBI` distribution is located:

```
shell> perl Makefile.PL -static -config
shell> make
shell> make install
shell> make perl
```

Then you must install the new Perl. The output of `make perl` indicates the exact `make` command you need to execute to perform the installation. On SCO, this is `make -f Makefile.aperl inst_perl MAP_TARGET=perl`.

Next, use the just-created Perl to create another Perl that also includes a statically linked `DBD::mysql` by running these commands in the directory where your `DBD::mysql` distribution is located:

```
shell> perl Makefile.PL -static -config
shell> make
shell> make install
shell> make perl
```

Finally, you should install this new Perl. Again, the output of `make perl` indicates the command to use.

---

# Chapter 3. Tutorial

This chapter provides a tutorial introduction to MySQL by showing how to use the `mysql` client program to create and use a simple database. `mysql` (sometimes referred to as the “terminal monitor” or just “monitor”) is an interactive program that allows you to connect to a MySQL server, run queries, and view the results. `mysql` may also be used in batch mode: you place your queries in a file beforehand, then tell `mysql` to execute the contents of the file. Both ways of using `mysql` are covered here.

To see a list of options provided by `mysql`, invoke it with the `--help` option:

```
shell> mysql --help
```

This chapter assumes that `mysql` is installed on your machine and that a MySQL server is available to which you can connect. If this is not true, contact your MySQL administrator. (If *you* are the administrator, you need to consult the relevant portions of this manual, such as [Chapter 5, Database Administration](#).)

This chapter describes the entire process of setting up and using a database. If you are interested only in accessing an existing database, you may want to skip over the sections that describe how to create the database and the tables it contains.

Because this chapter is tutorial in nature, many details are necessarily omitted. Consult the relevant sections of the manual for more information on the topics covered here.

## 3.1. Connecting to and Disconnecting from the Server

To connect to the server, you will usually need to provide a MySQL user name when you invoke `mysql` and, most likely, a password. If the server runs on a machine other than the one where you log in, you will also need to specify a host name. Contact your administrator to find out what connection parameters you should use to connect (that is, what host, user name, and password to use). Once you know the proper parameters, you should be able to connect like this:

```
shell> mysql -h host -u user -p
Enter password: *****
```

`host` and `user` represent the host name where your MySQL server is running and the user name of your MySQL account. Substitute appropriate values for your setup. The `*****` represents your password; enter it when `mysql` displays the `Enter password:` prompt.

If that works, you should see some introductory information followed by a `mysql>` prompt:

```
shell> mysql -h host -u user -p
Enter password: *****
Welcome to the MySQL monitor.  Commands end with ; or \g.
Your MySQL connection id is 25338 to server version: 4.1.22-standard

Type 'help;' or '\h' for help. Type '\c' to clear the buffer.

mysql>
```

The `mysql>` prompt tells you that `mysql` is ready for you to enter commands.

If you are logging in on the same machine that MySQL is running on, you can omit the host, and simply use the following:

```
shell< mysql -u user -p
```

If, when you attempt to log in, you get an error message such as ERROR 2002 (HY000): Can't connect to local MySQL server through socket '/tmp/mysql.sock' (2), it means that that MySQL server daemon (Unix) or service (Windows) is not running. Consult the administrator or see the section of [Chapter 2, \*Installing and Upgrading MySQL\*](#) that is appropriate to your operating system.

For help with other problems often encountered when trying to log in, see [Section A.2, “Common Errors When Using MySQL Programs”](#).

Some MySQL installations allow users to connect as the anonymous (unnamed) user to the server running on the local host. If this is the case on your machine, you should be able to connect to that server by invoking `mysql` without any options:

```
shell> mysql
```

After you have connected successfully, you can disconnect any time by typing `QUIT` (or `\q`) at the `mysql>` prompt:

```
mysql> QUIT
Bye
```

On Unix, you can also disconnect by pressing Control-D.

Most examples in the following sections assume that you are connected to the server. They indicate this by the `mysql>` prompt.

## 3.2. Entering Queries

Make sure that you are connected to the server, as discussed in the previous section. Doing so does not in itself select any database to work with, but that's okay. At this point, it's more important to find out a little about how to issue queries than to jump right in creating tables, loading data into them, and retrieving data from them. This section describes the basic principles of entering commands, using several queries you can try out to familiarize yourself with how `mysql` works.

Here's a simple command that asks the server to tell you its version number and the current date. Type it in as shown here following the `mysql>` prompt and press Enter:

```
mysql> SELECT VERSION(), CURRENT_DATE;
+-----+-----+
| VERSION() | CURRENT_DATE |
+-----+-----+
| 4.1.14-Max | 2005-09-03   |
+-----+-----+
1 row in set (0.01 sec)
mysql>
```

This query illustrates several things about `mysql`:

- A command normally consists of an SQL statement followed by a semicolon. (There are some exceptions where a semicolon may be omitted. `QUIT`, mentioned earlier, is one of them. We'll get to others later.)
- When you issue a command, `mysql` sends it to the server for execution and displays the results, then prints another `mysql>` prompt to indicate that it is ready for another command.
- `mysql` displays query output in tabular form (rows and columns). The first row contains labels for the columns. The rows following are the query results. Normally, column labels are the names of the columns you fetch from database tables. If you're retrieving the value of an expression rather than a table column (as in the example just shown), `mysql` labels the column using the expression itself.



- `mysql` shows how many rows were returned and how long the query took to execute, which gives you a rough idea of server performance. These values are imprecise because they represent wall clock time (not CPU or machine time), and because they are affected by factors such as server load and network latency. (For brevity, the “rows in set” line is sometimes not shown in the remaining examples in this chapter.)

Keywords may be entered in any lettercase. The following queries are equivalent:

```
mysql> SELECT VERSION(), CURRENT_DATE;
mysql> select version(), current_date;
mysql> SeLeCt vErSiOn(), cUrReNt_DaTe;
```

Here's another query. It demonstrates that you can use `mysql` as a simple calculator:

```
mysql> SELECT SIN(PI()/4), (4+1)*5;
+-----+-----+
| SIN(PI()/4) | (4+1)*5 |
+-----+-----+
| 0.707107    | 25      |
+-----+-----+
```

The queries shown thus far have been relatively short, single-line statements. You can even enter multiple statements on a single line. Just end each one with a semicolon:

```
mysql> SELECT VERSION(); SELECT NOW();
+-----+
| VERSION() |
+-----+
| 4.1.14-Max |
+-----+

+-----+
| NOW() |
+-----+
| 2005-09-03 12:27:16 |
+-----+
```

A command need not be given all on a single line, so lengthy commands that require several lines are not a problem. `mysql` determines where your statement ends by looking for the terminating semicolon, not by looking for the end of the input line. (In other words, `mysql` accepts free-format input: it collects input lines but does not execute them until it sees the semicolon.)

Here's a simple multiple-line statement:

```
mysql> SELECT
-> USER()
-> ,
-> CURRENT_DATE;
+-----+-----+
| USER() | CURRENT_DATE |
+-----+-----+
| jon@localhost | 2005-09-03 |
+-----+-----+
```

In this example, notice how the prompt changes from `mysql>` to `->` after you enter the first line of a multiple-line query. This is how `mysql` indicates that it has not yet seen a complete statement and is waiting for the rest. The prompt is your friend, because it provides valuable feedback. If you use that feedback, you can always be aware of what `mysql` is waiting for.

If you decide you do not want to execute a command that you are in the process of entering, cancel it by typing `\c`:

```
mysql> SELECT
-> USER()
```

```

-> \c
mysql>

```

Here, too, notice the prompt. It switches back to `mysql>` after you type `\c`, providing feedback to indicate that `mysql` is ready for a new command.

The following table shows each of the prompts you may see and summarizes what they mean about the state that `mysql` is in:

Prompt	Meaning
<code>mysql&gt;</code>	Ready for new command.
<code>-&gt;</code>	Waiting for next line of multiple-line command.
<code>'&gt;</code>	Waiting for next line, waiting for completion of a string that began with a single quote ('').
<code>"&gt;</code>	Waiting for next line, waiting for completion of a string that began with a double quote ("").
<code>`&gt;</code>	Waiting for next line, waiting for completion of an identifier that began with a backtick (`).
<code>/*&gt;</code>	Waiting for next line, waiting for completion of a comment that began with /*.

The ``>` prompt was implemented MySQL 4.0.16. The `/*>` prompt was implemented in MySQL 4.1.12.

Multiple-line statements commonly occur by accident when you intend to issue a command on a single line, but forget the terminating semicolon. In this case, `mysql` waits for more input:

```

mysql> SELECT USER()
->

```

If this happens to you (you think you've entered a statement but the only response is a `->` prompt), most likely `mysql` is waiting for the semicolon. If you don't notice what the prompt is telling you, you might sit there for a while before realizing what you need to do. Enter a semicolon to complete the statement, and `mysql` executes it:

```

mysql> SELECT USER()
-> ;
+-----+
| USER() |
+-----+
| jon@localhost |
+-----+

```

The `'>` and `">` prompts occur during string collection (another way of saying that MySQL is waiting for completion of a string). In MySQL, you can write strings surrounded by either `'` or `"` characters (for example, `'hello'` or `"goodbye"`), and `mysql` lets you enter strings that span multiple lines. When you see a `'>` or `">` prompt, it means that you have entered a line containing a string that begins with a `'` or `"` quote character, but have not yet entered the matching quote that terminates the string. This often indicates that you have inadvertently left out a quote character. For example:

```

mysql> SELECT * FROM my_table WHERE name = 'Smith AND age < 30;
'>

```

If you enter this `SELECT` statement, then press **Enter** and wait for the result, nothing happens. Instead of wondering why this query takes so long, notice the clue provided by the `'>` prompt. It tells you that `mysql` expects to see the rest of an unterminated string. (Do you see the error in the statement? The string `'Smith` is missing the second single quote mark.)

At this point, what do you do? The simplest thing is to cancel the command. However, you cannot just type `\c` in this case, because `mysql` interprets it as part of the string that it is collecting. Instead, enter the closing quote character (so `mysql` knows you've finished the string), then type `\c`:

```
mysql> SELECT * FROM my_table WHERE name = 'Smith AND age < 30;
'> '\c
mysql>
```

The prompt changes back to `mysql>`, indicating that `mysql` is ready for a new command.

The ``>` prompt is similar to the `'>` and `">` prompts, but indicates that you have begun but not completed a backtick-quoted identifier.

It is important to know what the `'>`, `">`, and ``>` prompts signify, because if you mistakenly enter an unterminated string, any further lines you type appear to be ignored by `mysql` — including a line containing `QUIT`. This can be quite confusing, especially if you do not know that you need to supply the terminating quote before you can cancel the current command.

### 3.3. Creating and Using a Database

Once you know how to enter commands, you are ready to access a database.

Suppose that you have several pets in your home (your menagerie) and you would like to keep track of various types of information about them. You can do so by creating tables to hold your data and loading them with the desired information. Then you can answer different sorts of questions about your animals by retrieving data from the tables. This section shows you how to:

- Create a database
- Create a table
- Load data into the table
- Retrieve data from the table in various ways
- Use multiple tables

The menagerie database is simple (deliberately), but it is not difficult to think of real-world situations in which a similar type of database might be used. For example, a database like this could be used by a farmer to keep track of livestock, or by a veterinarian to keep track of patient records. A menagerie distribution containing some of the queries and sample data used in the following sections can be obtained from the MySQL Web site. It is available in both compressed `tar` file and Zip formats at <http://dev.mysql.com/doc/>.

Use the `SHOW` statement to find out what databases currently exist on the server:

```
mysql> SHOW DATABASES;
+-----+
| Database |
+-----+
| mysql    |
| test     |
| tmp      |
+-----+
```

The list of databases is probably different on your machine, but the `mysql` and `test` databases are likely to be among them. The `mysql` database is required because it describes user access privileges. The `test` database is often provided as a workspace for users to try things out.

Note that you may not see all databases if you do not have the `SHOW DATABASES` privilege. See [Section 13.5.1.2, “GRANT Syntax”](#).

If the `test` database exists, try to access it:

```
mysql> USE test
Database changed
```

Note that `USE`, like `QUIT`, does not require a semicolon. (You can terminate such statements with a semicolon if you like; it does no harm.) The `USE` statement is special in another way, too: it must be given on a single line.

You can use the `test` database (if you have access to it) for the examples that follow, but anything you create in that database can be removed by anyone else with access to it. For this reason, you should probably ask your MySQL administrator for permission to use a database of your own. Suppose that you want to call yours `menagerie`. The administrator needs to execute a command like this:

```
mysql> GRANT ALL ON menagerie.* TO 'your_mysql_name'@'your_client_host';
```

where `your_mysql_name` is the MySQL user name assigned to you and `your_client_host` is the host from which you connect to the server.

### 3.3.1. Creating and Selecting a Database

If the administrator creates your database for you when setting up your permissions, you can begin using it. Otherwise, you need to create it yourself:

```
mysql> CREATE DATABASE menagerie;
```

Under Unix, database names are case sensitive (unlike SQL keywords), so you must always refer to your database as `menagerie`, not as `Menagerie`, `MENAGERIE`, or some other variant. This is also true for table names. (Under Windows, this restriction does not apply, although you must refer to databases and tables using the same lettercase throughout a given query. However, for a variety of reasons, our recommended best practice is always to use the same lettercase that was used when the database was created.)

**Note:** If you get an error such as `ERROR 1044 (42000): Access denied for user 'monty'@'localhost' to database 'menagerie'` when attempting to create a database, this means that your user account does not have the necessary privileges to do so. Discuss this with the administrator or see [Section 5.7, “The MySQL Access Privilege System”](#).

Creating a database does not select it for use; you must do that explicitly. To make `menagerie` the current database, use this command:

```
mysql> USE menagerie
Database changed
```

Your database needs to be created only once, but you must select it for use each time you begin a `mysql` session. You can do this by issuing a `USE` statement as shown in the example. Alternatively, you can select the database on the command line when you invoke `mysql`. Just specify its name after any connection parameters that you might need to provide. For example:

```
shell> mysql -h host -u user -p menagerie
Enter password: *****
```

Note that `menagerie` in the command just shown is **not** your password. If you want to supply your password on the command line after the `-p` option, you must do so with no intervening space (for example, as `-pmypassword`, not as `-p mypassword`). However, putting your password on the com-

mand line is not recommended, because doing so exposes it to snooping by other users logged in on your machine.

### 3.3.2. Creating a Table

Creating the database is the easy part, but at this point it's empty, as `SHOW TABLES` tells you:

```
mysql> SHOW TABLES;
Empty set (0.00 sec)
```

The harder part is deciding what the structure of your database should be: what tables you need and what columns should be in each of them.

You want a table that contains a record for each of your pets. This can be called the `pet` table, and it should contain, as a bare minimum, each animal's name. Because the name by itself is not very interesting, the table should contain other information. For example, if more than one person in your family keeps pets, you might want to list each animal's owner. You might also want to record some basic descriptive information such as species and sex.

How about age? That might be of interest, but it's not a good thing to store in a database. Age changes as time passes, which means you'd have to update your records often. Instead, it's better to store a fixed value such as date of birth. Then, whenever you need age, you can calculate it as the difference between the current date and the birth date. MySQL provides functions for doing date arithmetic, so this is not difficult. Storing birth date rather than age has other advantages, too:

- You can use the database for tasks such as generating reminders for upcoming pet birthdays. (If you think this type of query is somewhat silly, note that it is the same question you might ask in the context of a business database to identify clients to whom you need to send out birthday greetings in the current week or month, for that computer-assisted personal touch.)
- You can calculate age in relation to dates other than the current date. For example, if you store death date in the database, you can easily calculate how old a pet was when it died.

You can probably think of other types of information that would be useful in the `pet` table, but the ones identified so far are sufficient: name, owner, species, sex, birth, and death.

Use a `CREATE TABLE` statement to specify the layout of your table:

```
mysql> CREATE TABLE pet (name VARCHAR(20), owner VARCHAR(20),
-> species VARCHAR(20), sex CHAR(1), birth DATE, death DATE);
```

`VARCHAR` is a good choice for the `name`, `owner`, and `species` columns because the column values vary in length. The lengths in those column definitions need not all be the same, and need not be `20`. You can pick any length from `1` to `255`, whatever seems most reasonable to you. (If you make a poor choice and it turns out later that you need a longer field, MySQL provides an `ALTER TABLE` statement.)

Several types of values can be chosen to represent sex in animal records, such as `'m'` and `'f'`, or perhaps `'male'` and `'female'`. It is simplest to use the single characters `'m'` and `'f'`.

The use of the `DATE` data type for the `birth` and `death` columns is a fairly obvious choice.

Once you have created a table, `SHOW TABLES` should produce some output:

```
mysql> SHOW TABLES;
+-----+
| Tables in menagerie |
+-----+
| pet                  |
```

To verify that your table was created the way you expected, use a `DESCRIBE` statement:

```
mysql> DESCRIBE pet;
+-----+
| Field      | Type          | Null | Key | Default | Extra |
+-----+-----+-----+-----+-----+-----+
| name       | varchar(20)   | YES  |     | NULL    |       |
| owner      | varchar(20)   | YES  |     | NULL    |       |
| species    | varchar(20)   | YES  |     | NULL    |       |
| sex        | char(1)       | YES  |     | NULL    |       |
| birth      | date          | YES  |     | NULL    |       |
| death      | date          | YES  |     | NULL    |       |
+-----+-----+-----+-----+-----+-----+
```

You can use `DESCRIBE` any time, for example, if you forget the names of the columns in your table or what types they have.

For more information about MySQL data types, see [Chapter 11, Data Types](#).

### 3.3.3. Loading Data into a Table

After creating your table, you need to populate it. The `LOAD DATA` and `INSERT` statements are useful for this.

Suppose that your pet records can be described as shown here. (Observe that MySQL expects dates in 'YYYY-MM-DD' format; this may be different from what you are used to.)

name	owner	species	sex	birth	death
Fluffy	Harold	cat	f	1993-02-04	
Claws	Gwen	cat	m	1994-03-17	
Buffy	Harold	dog	f	1989-05-13	
Fang	Benny	dog	m	1990-08-27	
Bowser	Diane	dog	m	1979-08-31	1995-07-29
Chirpy	Gwen	bird	f	1998-09-11	
Whistler	Gwen	bird		1997-12-09	
Slim	Benny	snake	m	1996-04-29	

Because you are beginning with an empty table, an easy way to populate it is to create a text file containing a row for each of your animals, then load the contents of the file into the table with a single statement.

You could create a text file `pet.txt` containing one record per line, with values separated by tabs, and given in the order in which the columns were listed in the `CREATE TABLE` statement. For missing values (such as unknown sexes or death dates for animals that are still living), you can use `NULL` values. To represent these in your text file, use `\N` (backslash, capital-N). For example, the record for Whistler the bird would look like this (where the whitespace between values is a single tab character):

```
Whistler      Gwen      bird      \N      1997-12-09      \N
```

To load the text file `pet.txt` into the `pet` table, use this command:

```
mysql> LOAD DATA LOCAL INFILE '/path/pet.txt' INTO TABLE pet;
```

Note that if you created the file on Windows with an editor that uses `\r\n` as a line terminator, you should use:

```
mysql> LOAD DATA LOCAL INFILE '/path/pet.txt' INTO TABLE pet
-> LINES TERMINATED BY '\r\n';
```

(On an Apple machine running OS X, you would likely want to use `LINES TERMINATED BY '\r'`.)

You can specify the column value separator and end of line marker explicitly in the `LOAD DATA` statement if you wish, but the defaults are tab and linefeed. These are sufficient for the statement to read the file `pet.txt` properly.

If the statement fails, it is likely that your MySQL installation does not have local file capability enabled by default. See [Section 5.6.4, “Security Issues with `LOAD DATA LOCAL`”](#), for information on how to change this.

When you want to add new records one at a time, the `INSERT` statement is useful. In its simplest form, you supply values for each column, in the order in which the columns were listed in the `CREATE TABLE` statement. Suppose that Diane gets a new hamster named “Puffball.” You could add a new record using an `INSERT` statement like this:

```
mysql> INSERT INTO pet
-> VALUES ('Puffball', 'Diane', 'hamster', 'f', '1999-03-30', NULL);
```

Note that string and date values are specified as quoted strings here. Also, with `INSERT`, you can insert `NULL` directly to represent a missing value. You do not use `\N` like you do with `LOAD DATA`.

From this example, you should be able to see that there would be a lot more typing involved to load your records initially using several `INSERT` statements rather than a single `LOAD DATA` statement.

### 3.3.4. Retrieving Information from a Table

The `SELECT` statement is used to pull information from a table. The general form of the statement is:

```
SELECT what_to_select
FROM which_table
WHERE conditions_to_satisfy;
```

*what\_to\_select* indicates what you want to see. This can be a list of columns, or `*` to indicate “all columns.” *which\_table* indicates the table from which you want to retrieve data. The `WHERE` clause is optional. If it is present, *conditions\_to\_satisfy* specifies one or more conditions that rows must satisfy to qualify for retrieval.

#### 3.3.4.1. Selecting All Data

The simplest form of `SELECT` retrieves everything from a table:

```
mysql> SELECT * FROM pet;
```

name	owner	species	sex	birth	death
Fluffy	Harold	cat	f	1993-02-04	NULL
Claws	Gwen	cat	m	1994-03-17	NULL
Buffy	Harold	dog	f	1989-05-13	NULL
Fang	Benny	dog	m	1990-08-27	NULL
Bowser	Diane	dog	m	1979-08-31	1995-07-29
Chirpy	Gwen	bird	f	1998-09-11	NULL
Whistler	Gwen	bird	NULL	1997-12-09	NULL
Slim	Benny	snake	m	1996-04-29	NULL
Puffball	Diane	hamster	f	1999-03-30	NULL

This form of `SELECT` is useful if you want to review your entire table, for example, after you've just loaded it with your initial dataset. For example, you may happen to think that the birth date for Bowser doesn't seem quite right. Consulting your original pedigree papers, you find that the correct birth year should be 1989, not 1979.

There are at least two ways to fix this:

- Edit the file `pet.txt` to correct the error, then empty the table and reload it using `DELETE` and `LOAD DATA`:

```
mysql> DELETE FROM pet;
mysql> LOAD DATA LOCAL INFILE 'pet.txt' INTO TABLE pet;
```

However, if you do this, you must also re-enter the record for Puffball.

- Fix only the erroneous record with an `UPDATE` statement:

```
mysql> UPDATE pet SET birth = '1989-08-31' WHERE name = 'Bowser';
```

The `UPDATE` changes only the record in question and does not require you to reload the table.

### 3.3.4.2. Selecting Particular Rows

As shown in the preceding section, it is easy to retrieve an entire table. Just omit the `WHERE` clause from the `SELECT` statement. But typically you don't want to see the entire table, particularly when it becomes large. Instead, you're usually more interested in answering a particular question, in which case you specify some constraints on the information you want. Let's look at some selection queries in terms of questions about your pets that they answer.

You can select only particular rows from your table. For example, if you want to verify the change that you made to Bowser's birth date, select Bowser's record like this:

```
mysql> SELECT * FROM pet WHERE name = 'Bowser';
+-----+-----+-----+-----+-----+-----+
| name | owner | species | sex | birth      | death      |
+-----+-----+-----+-----+-----+-----+
| Bowser | Diane | dog      | m   | 1989-08-31 | 1995-07-29 |
+-----+-----+-----+-----+-----+-----+
```

The output confirms that the year is correctly recorded as 1989, not 1979.

String comparisons normally are case-insensitive, so you can specify the name as `'bowser'`, `'BOWSER'`, and so forth. The query result is the same.

You can specify conditions on any column, not just `name`. For example, if you want to know which animals were born during or after 1998, test the `birth` column:

```
mysql> SELECT * FROM pet WHERE birth >= '1998-1-1';
+-----+-----+-----+-----+-----+-----+
| name      | owner | species | sex | birth      | death      |
+-----+-----+-----+-----+-----+-----+
| Chirpy    | Gwen  | bird    | f   | 1998-09-11 | NULL       |
| Puffball  | Diane | hamster | f   | 1999-03-30 | NULL       |
+-----+-----+-----+-----+-----+-----+
```

You can combine conditions, for example, to locate female dogs:

```
mysql> SELECT * FROM pet WHERE species = 'dog' AND sex = 'f';
```



name	owner	species	sex	birth	death
Buffy	Harold	dog	f	1989-05-13	NULL

The preceding query uses the **AND** logical operator. There is also an **OR** operator:

```
mysql> SELECT * FROM pet WHERE species = 'snake' OR species = 'bird';
```

name	owner	species	sex	birth	death
Chirpy	Gwen	bird	f	1998-09-11	NULL
Whistler	Gwen	bird	NULL	1997-12-09	NULL
Slim	Benny	snake	m	1996-04-29	NULL

**AND** and **OR** may be intermixed, although **AND** has higher precedence than **OR**. If you use both operators, it is a good idea to use parentheses to indicate explicitly how conditions should be grouped:

```
mysql> SELECT * FROM pet WHERE (species = 'cat' AND sex = 'm')
-> OR (species = 'dog' AND sex = 'f');
```

name	owner	species	sex	birth	death
Claws	Gwen	cat	m	1994-03-17	NULL
Buffy	Harold	dog	f	1989-05-13	NULL

### 3.3.4.3. Selecting Particular Columns

If you do not want to see entire rows from your table, just name the columns in which you are interested, separated by commas. For example, if you want to know when your animals were born, select the **name** and **birth** columns:

```
mysql> SELECT name, birth FROM pet;
```

name	birth
Fluffy	1993-02-04
Claws	1994-03-17
Buffy	1989-05-13
Fang	1990-08-27
Bowser	1989-08-31
Chirpy	1998-09-11
Whistler	1997-12-09
Slim	1996-04-29
Puffball	1999-03-30

To find out who owns pets, use this query:

```
mysql> SELECT owner FROM pet;
```

owner
Harold
Gwen
Harold
Benny
Diane
Gwen
Gwen
Benny
Diane

Notice that the query simply retrieves the **owner** column from each record, and some of them appear more than once. To minimize the output, retrieve each unique output record just once by adding the

keyword `DISTINCT`:

```
mysql> SELECT DISTINCT owner FROM pet;
+-----+
| owner |
+-----+
| Benny |
| Diane |
| Gwen  |
| Harold|
+-----+
```

You can use a `WHERE` clause to combine row selection with column selection. For example, to get birth dates for dogs and cats only, use this query:

```
mysql> SELECT name, species, birth FROM pet
-> WHERE species = 'dog' OR species = 'cat';
+-----+-----+-----+
| name  | species | birth  |
+-----+-----+-----+
| Fluffy| cat     | 1993-02-04 |
| Claws | cat     | 1994-03-17 |
| Buffy | dog     | 1989-05-13 |
| Fang  | dog     | 1990-08-27 |
| Bowser| dog     | 1989-08-31 |
+-----+-----+-----+
```

### 3.3.4.4. Sorting Rows

You may have noticed in the preceding examples that the result rows are displayed in no particular order. It's often easier to examine query output when the rows are sorted in some meaningful way. To sort a result, use an `ORDER BY` clause.

Here are animal birthdays, sorted by date:

```
mysql> SELECT name, birth FROM pet ORDER BY birth;
+-----+-----+
| name  | birth  |
+-----+-----+
| Buffy | 1989-05-13 |
| Bowser| 1989-08-31 |
| Fang  | 1990-08-27 |
| Fluffy| 1993-02-04 |
| Claws | 1994-03-17 |
| Slim  | 1996-04-29 |
| Whistler| 1997-12-09 |
| Chirpy| 1998-09-11 |
| Puffball| 1999-03-30 |
+-----+-----+
```

On character type columns, sorting — like all other comparison operations — is normally performed in a case-insensitive fashion. This means that the order is undefined for columns that are identical except for their case. You can force a case-sensitive sort for a column by using `BINARY` like so: `ORDER BY BINARY col_name`.

The default sort order is ascending, with smallest values first. To sort in reverse (descending) order, add the `DESC` keyword to the name of the column you are sorting by:

```
mysql> SELECT name, birth FROM pet ORDER BY birth DESC;
+-----+-----+
| name  | birth  |
+-----+-----+
| Puffball| 1999-03-30 |
| Chirpy | 1998-09-11 |
| Whistler| 1997-12-09 |
| Slim   | 1996-04-29 |
| Claws  | 1994-03-17 |
| Fluffy | 1993-02-04 |
| Fang   | 1990-08-27 |
| Bowser | 1989-08-31 |
+-----+-----+
```

```
| Buffy | 1989-05-13 |
+-----+-----+
```

You can sort on multiple columns, and you can sort different columns in different directions. For example, to sort by type of animal in ascending order, then by birth date within animal type in descending order (youngest animals first), use the following query:

```
mysql> SELECT name, species, birth FROM pet
-> ORDER BY species, birth DESC;
+-----+-----+-----+
| name   | species | birth   |
+-----+-----+-----+
| Chirpy | bird    | 1998-09-11 |
| Whistler | bird   | 1997-12-09 |
| Claws  | cat     | 1994-03-17 |
| Fluffy | cat     | 1993-02-04 |
| Fang   | dog     | 1990-08-27 |
| Bowser | dog     | 1989-08-31 |
| Buffy  | dog     | 1989-05-13 |
| Puffball | hamster | 1999-03-30 |
| Slim   | snake   | 1996-04-29 |
+-----+-----+-----+
```

Note that the `DESC` keyword applies only to the column name immediately preceding it (`birth`); it does not affect the `species` column sort order.

### 3.3.4.5. Date Calculations

MySQL provides several functions that you can use to perform calculations on dates, for example, to calculate ages or extract parts of dates.

To determine how many years old each of your pets is, compute the difference in the year part of the current date and the birth date, then subtract one if the current date occurs earlier in the calendar year than the birth date. The following query shows, for each pet, the birth date, the current date, and the age in years.

```
mysql> SELECT name, birth, CURDATE(),
-> (YEAR(CURDATE())-YEAR(birth))
-> - (RIGHT(CURDATE(),5)<RIGHT(birth,5))
-> AS age
-> FROM pet;
+-----+-----+-----+-----+
| name   | birth   | CURDATE() | age |
+-----+-----+-----+-----+
| Fluffy | 1993-02-04 | 2003-08-19 | 10 |
| Claws  | 1994-03-17 | 2003-08-19 | 9   |
| Buffy  | 1989-05-13 | 2003-08-19 | 14  |
| Fang   | 1990-08-27 | 2003-08-19 | 12  |
| Bowser | 1989-08-31 | 2003-08-19 | 13  |
| Chirpy | 1998-09-11 | 2003-08-19 | 4   |
| Whistler | 1997-12-09 | 2003-08-19 | 5   |
| Slim   | 1996-04-29 | 2003-08-19 | 7   |
| Puffball | 1999-03-30 | 2003-08-19 | 4   |
+-----+-----+-----+-----+
```

Here, `YEAR()` pulls out the year part of a date and `RIGHT()` pulls off the rightmost five characters that represent the `MM-DD` (calendar year) part of the date. The part of the expression that compares the `MM-DD` values evaluates to 1 or 0, which adjusts the year difference down a year if `CURDATE()` occurs earlier in the year than `birth`. The full expression is somewhat ungainly, so an *alias* (`age`) is used to make the output column label more meaningful.

The query works, but the result could be scanned more easily if the rows were presented in some order. This can be done by adding an `ORDER BY name` clause to sort the output by name:

```
mysql> SELECT name, birth, CURDATE(),
-> (YEAR(CURDATE())-YEAR(birth))
-> - (RIGHT(CURDATE(),5)<RIGHT(birth,5))
-> AS age
```

```
-> FROM pet ORDER BY name;
```

name	birth	CURDATE()	age
Bowser	1989-08-31	2003-08-19	13
Buffy	1989-05-13	2003-08-19	14
Chirpy	1998-09-11	2003-08-19	4
Claws	1994-03-17	2003-08-19	9
Fang	1990-08-27	2003-08-19	12
Fluffy	1993-02-04	2003-08-19	10
Puffball	1999-03-30	2003-08-19	4
Slim	1996-04-29	2003-08-19	7
Whistler	1997-12-09	2003-08-19	5

To sort the output by `age` rather than `name`, just use a different `ORDER BY` clause:

```
mysql> SELECT name, birth, CURDATE(),
-> (YEAR(CURDATE())-YEAR(birth))
-> - (RIGHT(CURDATE(),5)<RIGHT(birth,5))
-> AS age
-> FROM pet ORDER BY age;
```

name	birth	CURDATE()	age
Chirpy	1998-09-11	2003-08-19	4
Puffball	1999-03-30	2003-08-19	4
Whistler	1997-12-09	2003-08-19	5
Slim	1996-04-29	2003-08-19	7
Claws	1994-03-17	2003-08-19	9
Fluffy	1993-02-04	2003-08-19	10
Fang	1990-08-27	2003-08-19	12
Bowser	1989-08-31	2003-08-19	13
Buffy	1989-05-13	2003-08-19	14

A similar query can be used to determine age at death for animals that have died. You determine which animals these are by checking whether the `death` value is `NULL`. Then, for those with non-`NULL` values, compute the difference between the `death` and `birth` values:

```
mysql> SELECT name, birth, death,
-> (YEAR(death)-YEAR(birth)) - (RIGHT(death,5)<RIGHT(birth,5))
-> AS age
-> FROM pet WHERE death IS NOT NULL ORDER BY age;
```

name	birth	death	age
Bowser	1989-08-31	1995-07-29	5

The query uses `death IS NOT NULL` rather than `death <> NULL` because `NULL` is a special value that cannot be compared using the usual comparison operators. This is discussed later. See [Section 3.3.4.6, “Working with NULL Values”](#).

What if you want to know which animals have birthdays next month? For this type of calculation, year and day are irrelevant; you simply want to extract the month part of the `birth` column. MySQL provides several functions for extracting parts of dates, such as `YEAR()`, `MONTH()`, and `DAYOFMONTH()`. `MONTH()` is the appropriate function here. To see how it works, run a simple query that displays the value of both `birth` and `MONTH(birth)`:

```
mysql> SELECT name, birth, MONTH(birth) FROM pet;
```

name	birth	MONTH(birth)
Fluffy	1993-02-04	2
Claws	1994-03-17	3
Buffy	1989-05-13	5
Fang	1990-08-27	8
Bowser	1989-08-31	8
Chirpy	1998-09-11	9
Whistler	1997-12-09	12
Slim	1996-04-29	4

```
| Puffball | 1999-03-30 |          3 |
+-----+-----+-----+
```

Finding animals with birthdays in the upcoming month is also simple. Suppose that the current month is April. Then the month value is 4 and you can look for animals born in May (month 5) like this:

```
mysql> SELECT name, birth FROM pet WHERE MONTH(birth) = 5;
+-----+-----+
| name | birth |
+-----+-----+
| Buffy | 1989-05-13 |
+-----+-----+
```

There is a small complication if the current month is December. You cannot merely add one to the month number (12) and look for animals born in month 13, because there is no such month. Instead, you look for animals born in January (month 1).

You can write the query so that it works no matter what the current month is, so that you do not have to use the number for a particular month. `DATE_ADD()` allows you to add a time interval to a given date. If you add a month to the value of `CURDATE()`, then extract the month part with `MONTH()`, the result produces the month in which to look for birthdays:

```
mysql> SELECT name, birth FROM pet
-> WHERE MONTH(birth) = MONTH(DATE_ADD(CURDATE(), INTERVAL 1 MONTH));
```

A different way to accomplish the same task is to add 1 to get the next month after the current one after using the modulo function (`MOD`) to wrap the month value to 0 if it is currently 12:

```
mysql> SELECT name, birth FROM pet
-> WHERE MONTH(birth) = MOD(MONTH(CURDATE()), 12) + 1;
```

Note that `MONTH` returns a number between 1 and 12. And `MOD(something, 12)` returns a number between 0 and 11. So the addition has to be after the `MOD()`, otherwise we would go from November (11) to January (1).

### 3.3.4.6. Working with NULL Values

The `NULL` value can be surprising until you get used to it. Conceptually, `NULL` means “a missing unknown value” and it is treated somewhat differently from other values. To test for `NULL`, you cannot use the arithmetic comparison operators such as `=`, `<`, or `<>`. To demonstrate this for yourself, try the following query:

```
mysql> SELECT 1 = NULL, 1 <> NULL, 1 < NULL, 1 > NULL;
+-----+-----+-----+-----+
| 1 = NULL | 1 <> NULL | 1 < NULL | 1 > NULL |
+-----+-----+-----+-----+
| NULL | NULL | NULL | NULL |
+-----+-----+-----+-----+
```

Clearly you get no meaningful results from these comparisons. Use the `IS NULL` and `IS NOT NULL` operators instead:

```
mysql> SELECT 1 IS NULL, 1 IS NOT NULL;
+-----+-----+
| 1 IS NULL | 1 IS NOT NULL |
+-----+-----+
| 0 | 1 |
+-----+-----+
```

Note that in MySQL, 0 or `NULL` means false and anything else means true. The default truth value from a boolean operation is 1.

This special treatment of `NULL` is why, in the previous section, it was necessary to determine which animals are no longer alive using `death IS NOT NULL` instead of `death <> NULL`.

Two `NULL` values are regarded as equal in a `GROUP BY`.

When doing an `ORDER BY`, `NULL` values are presented first if you do `ORDER BY ... ASC` and last if you do `ORDER BY ... DESC`.

Note that MySQL 4.0.2 to 4.0.10 incorrectly always sorts `NULL` values first regardless of the sort direction.

A common error when working with `NULL` is to assume that it is not possible to insert a zero or an empty string into a column defined as `NOT NULL`, but this is not the case. These are in fact values, whereas `NULL` means “not having a value.” You can test this easily enough by using `IS [NOT] NULL` as shown:

```
mysql> SELECT 0 IS NULL, 0 IS NOT NULL, '' IS NULL, '' IS NOT NULL;
+-----+-----+-----+-----+
| 0 IS NULL | 0 IS NOT NULL | '' IS NULL | '' IS NOT NULL |
+-----+-----+-----+-----+
| 0 | 1 | 0 | 1 |
+-----+-----+-----+-----+
```

Thus it is entirely possible to insert a zero or empty string into a `NOT NULL` column, as these are in fact `NOT NULL`. See [Section A.5.3, “Problems with NULL Values”](#).

### 3.3.4.7. Pattern Matching

MySQL provides standard SQL pattern matching as well as a form of pattern matching based on extended regular expressions similar to those used by Unix utilities such as `vi`, `grep`, and `sed`.

SQL pattern matching allows you to use `'_'` to match any single character and `'%'` to match an arbitrary number of characters (including zero characters). In MySQL, SQL patterns are case-insensitive by default. Some examples are shown here. Note that you do not use `=` or `<>` when you use SQL patterns; use the `LIKE` or `NOT LIKE` comparison operators instead.

To find names beginning with `'b'`:

```
mysql> SELECT * FROM pet WHERE name LIKE 'b%';
+-----+-----+-----+-----+-----+-----+
| name | owner | species | sex | birth | death |
+-----+-----+-----+-----+-----+-----+
| Buffy | Harold | dog | f | 1989-05-13 | NULL |
| Bowser | Diane | dog | m | 1989-08-31 | 1995-07-29 |
+-----+-----+-----+-----+-----+-----+
```

To find names ending with `'fy'`:

```
mysql> SELECT * FROM pet WHERE name LIKE '%fy';
+-----+-----+-----+-----+-----+-----+
| name | owner | species | sex | birth | death |
+-----+-----+-----+-----+-----+-----+
| Fluffy | Harold | cat | f | 1993-02-04 | NULL |
| Buffy | Harold | dog | f | 1989-05-13 | NULL |
+-----+-----+-----+-----+-----+-----+
```

To find names containing a `'w'`:

```
mysql> SELECT * FROM pet WHERE name LIKE '%w%';
+-----+-----+-----+-----+-----+-----+
| name | owner | species | sex | birth | death |
+-----+-----+-----+-----+-----+-----+
| Claws | Gwen | cat | m | 1994-03-17 | NULL |
| Bowser | Diane | dog | m | 1989-08-31 | 1995-07-29 |
+-----+-----+-----+-----+-----+-----+
```

Whistler	Gwen	bird	NULL	1997-12-09	NULL
----------	------	------	------	------------	------

To find names containing exactly five characters, use five instances of the ‘\_’ pattern character:

```
mysql> SELECT * FROM pet WHERE name LIKE '_____';
```

name	owner	species	sex	birth	death
Claws	Gwen	cat	m	1994-03-17	NULL
Buffy	Harold	dog	f	1989-05-13	NULL

The other type of pattern matching provided by MySQL uses extended regular expressions. When you test for a match for this type of pattern, use the `REGEXP` and `NOT REGEXP` operators (or `RLIKE` and `NOT RLIKE`, which are synonyms).

Some characteristics of extended regular expressions are:

- ‘.’ matches any single character.
- A character class ‘[...]’ matches any character within the brackets. For example, ‘[abc]’ matches ‘a’, ‘b’, or ‘c’. To name a range of characters, use a dash. ‘[a-z]’ matches any letter, whereas ‘[0-9]’ matches any digit.
- ‘\*’ matches zero or more instances of the thing preceding it. For example, ‘x\*’ matches any number of ‘x’ characters, ‘[0-9]\*’ matches any number of digits, and ‘.\*’ matches any number of anything.
- A `REGEXP` pattern match succeeds if the pattern matches anywhere in the value being tested. (This differs from a `LIKE` pattern match, which succeeds only if the pattern matches the entire value.)
- To anchor a pattern so that it must match the beginning or end of the value being tested, use ‘^’ at the beginning or ‘\$’ at the end of the pattern.

To demonstrate how extended regular expressions work, the `LIKE` queries shown previously are rewritten here to use `REGEXP`.

To find names beginning with ‘b’, use ‘^’ to match the beginning of the name:

```
mysql> SELECT * FROM pet WHERE name REGEXP '^b';
```

name	owner	species	sex	birth	death
Buffy	Harold	dog	f	1989-05-13	NULL
Bowser	Diane	dog	m	1989-08-31	1995-07-29

Prior to MySQL 3.23.4, `REGEXP` is case sensitive, and the previous query will return no rows. In this case, to match either lowercase or uppercase ‘b’, use this query instead:

```
mysql> SELECT * FROM pet WHERE name REGEXP '^[bb]';
```

From MySQL 3.23.4 on, if you really want to force a `REGEXP` comparison to be case sensitive, use the `BINARY` keyword to make one of the strings a binary string. This query matches only lowercase ‘b’ at the beginning of a name:

```
mysql> SELECT * FROM pet WHERE name REGEXP BINARY '^b';
```

To find names ending with 'fy', use '\$' to match the end of the name:

```
mysql> SELECT * FROM pet WHERE name REGEXP 'fy$';
```

name	owner	species	sex	birth	death
Fluffy	Harold	cat	f	1993-02-04	NULL
Buffy	Harold	dog	f	1989-05-13	NULL

To find names containing a 'w', use this query:

```
mysql> SELECT * FROM pet WHERE name REGEXP 'w';
```

name	owner	species	sex	birth	death
Claws	Gwen	cat	m	1994-03-17	NULL
Bowser	Diane	dog	m	1989-08-31	1995-07-29
Whistler	Gwen	bird	NULL	1997-12-09	NULL

Because a regular expression pattern matches if it occurs anywhere in the value, it is not necessary in the previous query to put a wildcard on either side of the pattern to get it to match the entire value like it would be if you used an SQL pattern.

To find names containing exactly five characters, use '^' and '\$' to match the beginning and end of the name, and five instances of '.' in between:

```
mysql> SELECT * FROM pet WHERE name REGEXP '^.....$';
```

name	owner	species	sex	birth	death
Claws	Gwen	cat	m	1994-03-17	NULL
Buffy	Harold	dog	f	1989-05-13	NULL

You could also write the previous query using the {n} ("repeat-n-times") operator:

```
mysql> SELECT * FROM pet WHERE name REGEXP '^.{5}$';
```

name	owner	species	sex	birth	death
Claws	Gwen	cat	m	1994-03-17	NULL
Buffy	Harold	dog	f	1989-05-13	NULL

[Appendix G, Regular Expressions](#), provides more information about the syntax for regular expressions.

### 3.3.4.8. Counting Rows

Databases are often used to answer the question, "How often does a certain type of data occur in a table?" For example, you might want to know how many pets you have, or how many pets each owner has, or you might want to perform various kinds of census operations on your animals.

Counting the total number of animals you have is the same question as "How many rows are in the pet table?" because there is one record per pet. `COUNT(*)` counts the number of rows, so the query to count your animals looks like this:

```
mysql> SELECT COUNT(*) FROM pet;
```

COUNT(*)
9



Earlier, you retrieved the names of the people who owned pets. You can use `COUNT()` if you want to find out how many pets each owner has:

```
mysql> SELECT owner, COUNT(*) FROM pet GROUP BY owner;
```

owner	COUNT(*)
Benny	2
Diane	2
Gwen	3
Harold	2

Note the use of `GROUP BY` to group all records for each `owner`. Without it, all you get is an error message:

```
mysql> SELECT owner, COUNT(*) FROM pet;
ERROR 1140: Mixing of GROUP columns (MIN(),MAX(),COUNT()...)
with no GROUP columns is illegal if there is no GROUP BY clause
```

`COUNT()` and `GROUP BY` are useful for characterizing your data in various ways. The following examples show different ways to perform animal census operations.

Number of animals per species:

```
mysql> SELECT species, COUNT(*) FROM pet GROUP BY species;
```

species	COUNT(*)
bird	2
cat	2
dog	3
hamster	1
snake	1

Number of animals per sex:

```
mysql> SELECT sex, COUNT(*) FROM pet GROUP BY sex;
```

sex	COUNT(*)
NULL	1
f	4
m	4

(In this output, `NULL` indicates that the sex is unknown.)

Number of animals per combination of species and sex:

```
mysql> SELECT species, sex, COUNT(*) FROM pet GROUP BY species, sex;
```

species	sex	COUNT(*)
bird	NULL	1
bird	f	1
cat	f	1
cat	m	1
dog	f	1
dog	m	2
hamster	f	1
snake	m	1

You need not retrieve an entire table when you use `COUNT()`. For example, the previous query, when performed just on dogs and cats, looks like this:

```
mysql> SELECT species, sex, COUNT(*) FROM pet
-> WHERE species = 'dog' OR species = 'cat'
-> GROUP BY species, sex;
```

species	sex	COUNT(*)
cat	f	1
cat	m	1
dog	f	1
dog	m	2

Or, if you wanted the number of animals per sex only for animals whose sex is known:

```
mysql> SELECT species, sex, COUNT(*) FROM pet
-> WHERE sex IS NOT NULL
-> GROUP BY species, sex;
```

species	sex	COUNT(*)
bird	f	1
cat	f	1
cat	m	1
dog	f	1
dog	m	2
hamster	f	1
snake	m	1

### 3.3.4.9. Using More Than one Table

The `pet` table keeps track of which pets you have. If you want to record other information about them, such as events in their lives like visits to the vet or when litters are born, you need another table. What should this table look like? It needs:

- To contain the pet name so that you know which animal each event pertains to.
- A date so that you know when the event occurred.
- A field to describe the event.
- An event type field, if you want to be able to categorize events.

Given these considerations, the `CREATE TABLE` statement for the `event` table might look like this:

```
mysql> CREATE TABLE event (name VARCHAR(20), date DATE,
-> type VARCHAR(15), remark VARCHAR(255));
```

As with the `pet` table, it's easiest to load the initial records by creating a tab-delimited text file containing the information:

name	date	type	remark
Fluffy	1995-05-15	litter	4 kittens, 3 female, 1 male
Buffy	1993-06-23	litter	5 puppies, 2 female, 3 male
Buffy	1994-06-19	litter	3 puppies, 3 female
Chirpy	1999-03-21	vet	needed beak straightened
Slim	1997-08-03	vet	broken rib
Bowser	1991-10-12	kennel	
Fang	1991-10-12	kennel	

Fang	1998-08-28	birthday	Gave him a new chew toy
Claws	1998-03-17	birthday	Gave him a new flea collar
Whistler	1998-12-09	birthday	First birthday

Load the records like this:

```
mysql> LOAD DATA LOCAL INFILE 'event.txt' INTO TABLE event;
```

Based on what you have learned from the queries that you have run on the `pet` table, you should be able to perform retrievals on the records in the `event` table; the principles are the same. But when is the `event` table by itself insufficient to answer questions you might ask?

Suppose that you want to find out the ages at which each pet had its litters. We saw earlier how to calculate ages from two dates. The litter date of the mother is in the `event` table, but to calculate her age on that date you need her birth date, which is stored in the `pet` table. This means the query requires both tables:

```
mysql> SELECT pet.name,
-> (YEAR(date)-YEAR(birth)) - (RIGHT(date,5)<RIGHT(birth,5)) AS age,
-> remark
-> FROM pet INNER JOIN event
-> ON pet.name = event.name
-> WHERE event.type = 'litter';
```

name	age	remark
Fluffy	2	4 kittens, 3 female, 1 male
Buffy	4	5 puppies, 2 female, 3 male
Buffy	5	3 puppies, 3 female

There are several things to note about this query:

- The `FROM` clause joins two tables because the query needs to pull information from both of them.
- When combining (joining) information from multiple tables, you need to specify how records in one table can be matched to records in the other. This is easy because they both have a `name` column. The query uses `WHERE` clause to match up records in the two tables based on the `name` values.

The query uses an `INNER JOIN` to combine the tables. An `INNER JOIN` allows for rows from either table to appear in the result if and only if both tables meet the conditions specified in the `ON` clause. In this example, the `ON` clause specifies that the `name` column in the `pet` table must match the `name` column in the `event` table. If a name appears in one table but not the other, the row will not appear in the result because the condition in the `ON` clause fails.

- Because the `name` column occurs in both tables, you must be specific about which table you mean when referring to the column. This is done by prepending the table name to the column name.

You need not have two different tables to perform a join. Sometimes it is useful to join a table to itself, if you want to compare records in a table to other records in that same table. For example, to find breeding pairs among your pets, you can join the `pet` table with itself to produce candidate pairs of males and females of like species:

```
mysql> SELECT p1.name, p1.sex, p2.name, p2.sex, p1.species
-> FROM pet AS p1 INNER JOIN pet AS p2
-> ON p1.species = p2.species AND p1.sex = 'f' AND p2.sex = 'm';
```

name	sex	name	sex	species
------	-----	------	-----	---------

Fluffy	f	Claws	m	cat
Buffy	f	Fang	m	dog
Buffy	f	Bowser	m	dog

In this query, we specify aliases for the table name to refer to the columns and keep straight which instance of the table each column reference is associated with.

### 3.4. Getting Information About Databases and Tables

What if you forget the name of a database or table, or what the structure of a given table is (for example, what its columns are called)? MySQL addresses this problem through several statements that provide information about the databases and tables it supports.

You have previously seen `SHOW DATABASES`, which lists the databases managed by the server. To find out which database is currently selected, use the `DATABASE()` function:

```
mysql> SELECT DATABASE();
+-----+
| DATABASE() |
+-----+
| menagerie |
+-----+
```

If you have not yet selected any database, the result is `NULL` (or the empty string before MySQL 4.1.1).

To find out what tables the default database contains (for example, when you are not sure about the name of a table), use this command:

```
mysql> SHOW TABLES;
+-----+
| Tables in menagerie |
+-----+
| event               |
| pet                 |
+-----+
```

If you want to find out about the structure of a table, the `DESCRIBE` command is useful; it displays information about each of a table's columns:

```
mysql> DESCRIBE pet;
+-----+-----+-----+-----+-----+-----+
| Field | Type          | Null | Key | Default | Extra |
+-----+-----+-----+-----+-----+-----+
| name  | varchar(20)   | YES  |     | NULL    |       |
| owner | varchar(20)   | YES  |     | NULL    |       |
| species | varchar(20) | YES  |     | NULL    |       |
| sex   | char(1)       | YES  |     | NULL    |       |
| birth | date          | YES  |     | NULL    |       |
| death | date          | YES  |     | NULL    |       |
+-----+-----+-----+-----+-----+-----+
```

`Field` indicates the column name, `Type` is the data type for the column, `NULL` indicates whether the column can contain `NULL` values, `Key` indicates whether the column is indexed, and `Default` specifies the column's default value.

If you have indexes on a table, `SHOW INDEX FROM tbl_name` produces information about them.

### 3.5. Using `mysql` in Batch Mode

In the previous sections, you used `mysql` interactively to enter queries and view the results. You can also run `mysql` in batch mode. To do this, put the commands you want to run in a file, then tell `mysql` to read its input from the file:

```
shell> mysql < batch-file
```

If you are running `mysql` under Windows and have some special characters in the file that cause problems, you can do this:

```
C:\> mysql -e "source batch-file"
```

If you need to specify connection parameters on the command line, the command might look like this:

```
shell> mysql -h host -u user -p < batch-file
Enter password: *****
```

When you use `mysql` this way, you are creating a script file, then executing the script.

If you want the script to continue even if some of the statements in it produce errors, you should use the `--force` command-line option.

Why use a script? Here are a few reasons:

- If you run a query repeatedly (say, every day or every week), making it a script allows you to avoid retyping it each time you execute it.
- You can generate new queries from existing ones that are similar by copying and editing script files.
- Batch mode can also be useful while you're developing a query, particularly for multiple-line commands or multiple-statement sequences of commands. If you make a mistake, you don't have to re-type everything. Just edit your script to correct the error, then tell `mysql` to execute it again.
- If you have a query that produces a lot of output, you can run the output through a pager rather than watching it scroll off the top of your screen:

```
shell> mysql < batch-file | more
```

- You can catch the output in a file for further processing:

```
shell> mysql < batch-file > mysql.out
```

- You can distribute your script to other people so that they can also run the commands.
- Some situations do not allow for interactive use, for example, when you run a query from a `cron` job. In this case, you must use batch mode.

The default output format is different (more concise) when you run `mysql` in batch mode than when you use it interactively. For example, the output of `SELECT DISTINCT species FROM pet` looks like this when `mysql` is run interactively:

```
+-----+
| species |
+-----+
| bird    |
| cat     |
| dog     |
| hamster |
| snake   |
+-----+
```

In batch mode, the output looks like this instead:

```
species
bird
cat
dog
hamster
snake
```

If you want to get the interactive output format in batch mode, use `mysql -t`. To echo to the output the commands that are executed, use `mysql -vvv`.

You can also use scripts from the `mysql` prompt by using the `source` or `\.` command:

```
mysql> source filename;
mysql> \. filename
```

## 3.6. Examples of Common Queries

Here are examples of how to solve some common problems with MySQL.

Some of the examples use the table `shop` to hold the price of each article (item number) for certain traders (dealers). Supposing that each trader has a single fixed price per article, then (`article`, `dealer`) is a primary key for the records.

Start the command-line tool `mysql` and select a database:

```
shell> mysql your-database-name
```

(In most MySQL installations, you can use the database named `test`).

You can create and populate the example table with these statements:

```
mysql> CREATE TABLE shop (
-> article INT(4) UNSIGNED ZEROFILL DEFAULT '0000' NOT NULL,
-> dealer CHAR(20) DEFAULT '' NOT NULL,
-> price DOUBLE(16,2) DEFAULT '0.00' NOT NULL,
-> PRIMARY KEY(article, dealer));
mysql> INSERT INTO shop VALUES
-> (1,'A',3.45),(1,'B',3.99),(2,'A',10.99),(3,'B',1.45),
-> (3,'C',1.69),(3,'D',1.25),(4,'D',19.95);
```

After issuing the statements, the table should have the following contents:

```
mysql> SELECT * FROM shop;
+-----+-----+-----+
| article | dealer | price |
+-----+-----+-----+
| 0001   | A     | 3.45  |
| 0001   | B     | 3.99  |
| 0002   | A     | 10.99 |
| 0003   | B     | 1.45  |
| 0003   | C     | 1.69  |
| 0003   | D     | 1.25  |
| 0004   | D     | 19.95 |
+-----+-----+-----+
```

### 3.6.1. The Maximum Value for a Column

“What’s the highest item number?”

```
SELECT MAX(article) AS article FROM shop;
+-----+
| article |
```

```
+-----+
|         4 |
+-----+
```

### 3.6.2. The Row Holding the Maximum of a Certain Column

*Task: Find the number, dealer, and price of the most expensive article.*

In standard SQL (and as of MySQL 4.1), this is easily done with a subquery:

```
SELECT article, dealer, price
FROM   shop
WHERE  price=(SELECT MAX(price) FROM shop);
```

In MySQL versions prior to 4.1, just do it in two steps:

1. Get the maximum price value from the table with a `SELECT` statement.

```
mysql> SELECT MAX(price) FROM shop;
+-----+
| MAX(price) |
+-----+
|      19.95 |
+-----+
```

2. Using the value 19.95 shown by the previous query to be the maximum article price, write a query to locate and display the corresponding record:

```
mysql> SELECT article, dealer, price
-> FROM   shop
-> WHERE  price=19.95;
+-----+-----+-----+
| article | dealer | price |
+-----+-----+-----+
|    0004 | D      | 19.95 |
+-----+-----+-----+
```

Another solution is to sort all rows descending by price and get only the first row using the MySQL-specific `LIMIT` clause:

```
SELECT article, dealer, price
FROM   shop
ORDER BY price DESC
LIMIT 1;
```

**Note:** If there were several most expensive articles, each with a price of 19.95, the `LIMIT` solution would show only one of them.

### 3.6.3. Maximum of Column per Group

*Task: Find the highest price per article.*

```
SELECT article, MAX(price) AS price
FROM   shop
GROUP BY article
```

```
+-----+-----+
| article | price |
+-----+-----+
|    0001 |   3.99 |
|    0002 |  10.99 |
|    0003 |   1.69 |
+-----+-----+
```

```
| 0004 | 19.95 |
+-----+-----+
```

### 3.6.4. The Rows Holding the Group-wise Maximum of a Certain Field

*Task: For each article, find the dealer or dealers with the most expensive price.*

In standard SQL (and as of MySQL 4.1), the problem can be solved with a subquery like this:

```
SELECT article, dealer, price
FROM   shop s1
WHERE  price=(SELECT MAX(s2.price)
              FROM shop s2
              WHERE s1.article = s2.article);
```

In MySQL versions prior to 4.1, it's best do it in several steps:

1. Get the list of (article,maxprice) pairs.
2. For each article, get the corresponding rows that have the stored maximum price.

This can easily be done with a temporary table and a join:

```
CREATE TEMPORARY TABLE tmp (
  article INT(4) UNSIGNED ZEROFILL DEFAULT '0000' NOT NULL,
  price   DOUBLE(16,2)                DEFAULT '0.00' NOT NULL);

LOCK TABLES shop READ;

INSERT INTO tmp SELECT article, MAX(price) FROM shop GROUP BY article;

SELECT shop.article, dealer, shop.price FROM shop, tmp
WHERE shop.article=tmp.article AND shop.price=tmp.price;

UNLOCK TABLES;

DROP TABLE tmp;
```

If you don't use a [TEMPORARY](#) table, you must also lock the `tmp` table.

“Can it be done with a single query?”

Yes, but only by using a quite inefficient trick called the “MAX-CONCAT trick”:

```
SELECT article,
  SUBSTRING( MAX( CONCAT(LPAD(price,6,'0'),dealer) ), 7) AS dealer,
  0.00+LEFT(  MAX( CONCAT(LPAD(price,6,'0'),dealer) ), 6) AS price
FROM   shop
GROUP BY article;
```

```
+-----+-----+-----+
| article | dealer | price |
+-----+-----+-----+
| 0001   | B     | 3.99  |
| 0002   | A     | 10.99 |
| 0003   | C     | 1.69  |
| 0004   | D     | 19.95 |
+-----+-----+-----+
```

The last example can be made a bit more efficient by doing the splitting of the concatenated column in the client.



### 3.6.5. Using User-Defined Variables

You can employ MySQL user variables to remember results without having to store them in temporary variables in the client. (See [Section 9.3, “User-Defined Variables”](#).)

For example, to find the articles with the highest and lowest price you can do this:

```
mysql> SELECT @min_price:=MIN(price),@max_price:=MAX(price) FROM shop;
mysql> SELECT * FROM shop WHERE price=@min_price OR price=@max_price;
```

article	dealer	price
0003	D	1.25
0004	D	19.95

### 3.6.6. Using Foreign Keys

In MySQL 3.23.44 and up, [InnoDB](#) tables support checking of foreign key constraints. See [Section 14.2, “The InnoDB Storage Engine”](#), and [Section 1.9.5.5, “Foreign Keys”](#).

A foreign key constraint is not required merely to join two tables. For storage engines other than [InnoDB](#), it is possible when defining a column to use a `REFERENCES tbl_name(col_name)` clause, which has no actual effect, and *serves only as a memo or comment to you that the column which you are currently defining is intended to refer to a column in another table*. It is extremely important to realize when using this syntax that:

- MySQL does not perform any sort of `CHECK` to make sure that `col_name` actually exists in `tbl_name` (or even that `tbl_name` itself exists).
- MySQL does not perform any sort of action on `tbl_name` such as deleting rows in response to actions taken on rows in the table which you are defining; in other words, this syntax induces no `ON DELETE` or `ON UPDATE` behavior whatsoever. (Although you can write an `ON DELETE` or `ON UPDATE` clause as part of the `REFERENCES` clause, it is also ignored.)
- This syntax creates a *column*; it does **not** create any sort of index or key.
- This syntax will cause an error if used in trying to define an [InnoDB](#) table.

You can use a column so created as a join column, as shown here:

```
CREATE TABLE person (
  id SMALLINT UNSIGNED NOT NULL AUTO_INCREMENT,
  name CHAR(60) NOT NULL,
  PRIMARY KEY (id)
);

CREATE TABLE shirt (
  id SMALLINT UNSIGNED NOT NULL AUTO_INCREMENT,
  style ENUM('t-shirt', 'polo', 'dress') NOT NULL,
  color ENUM('red', 'blue', 'orange', 'white', 'black') NOT NULL,
  owner SMALLINT UNSIGNED NOT NULL REFERENCES person(id),
  PRIMARY KEY (id)
);

INSERT INTO person VALUES (NULL, 'Antonio Paz');

SELECT @last := LAST_INSERT_ID();

INSERT INTO shirt VALUES
(NULL, 'polo', 'blue', @last),
(NULL, 'dress', 'white', @last),
(NULL, 't-shirt', 'blue', @last);

INSERT INTO person VALUES (NULL, 'Lilliana Angelovska');
```

```

SELECT @last := LAST_INSERT_ID();

INSERT INTO shirt VALUES
(NULL, 'dress', 'orange', @last),
(NULL, 'polo', 'red', @last),
(NULL, 'dress', 'blue', @last),
(NULL, 't-shirt', 'white', @last);

SELECT * FROM person;
+-----+-----+
| id | name                |
+-----+-----+
| 1  | Antonio Paz         |
| 2  | Lilliana Angelovska |
+-----+-----+

SELECT * FROM shirt;
+-----+-----+-----+-----+
| id | style  | color  | owner |
+-----+-----+-----+-----+
| 1  | polo   | blue   | 1     |
| 2  | dress  | white  | 1     |
| 3  | t-shirt| blue   | 1     |
| 4  | dress  | orange | 2     |
| 5  | polo   | red    | 2     |
| 6  | dress  | blue   | 2     |
| 7  | t-shirt| white  | 2     |
+-----+-----+-----+-----+

SELECT s.* FROM person p INNER JOIN shirt s
ON s.owner = p.id
WHERE p.name LIKE 'Lilliana%'
AND s.color <> 'white';

+-----+-----+-----+-----+
| id | style  | color  | owner |
+-----+-----+-----+-----+
| 4  | dress  | orange | 2     |
| 5  | polo   | red    | 2     |
| 6  | dress  | blue   | 2     |
+-----+-----+-----+-----+

```

When used in this fashion, the [REFERENCES](#) clause is not displayed in the output of [SHOW CREATE TABLE](#) or [DESCRIBE](#):

```

SHOW CREATE TABLE shirt\G
***** 1. row *****
Table: shirt
Create Table: CREATE TABLE `shirt` (
  `id` smallint(5) unsigned NOT NULL auto_increment,
  `style` enum('t-shirt','polo','dress') NOT NULL,
  `color` enum('red','blue','orange','white','black') NOT NULL,
  `owner` smallint(5) unsigned NOT NULL,
  PRIMARY KEY (`id`)
) ENGINE=MyISAM DEFAULT CHARSET=latin1

```

The use of [REFERENCES](#) in this way as a comment or “reminder” in a column definition works with both [MyISAM](#) and [BerkeleyDB](#) tables.

### 3.6.7. Searching on Two Keys

An [OR](#) using a single key is well optimized, as is the handling of [AND](#).

The one tricky case is that of searching on two different keys combined with [OR](#):

```

SELECT field1_index, field2_index FROM test_table
WHERE field1_index = '1' OR field2_index = '1'

```

In MySQL 4.0 and up, you can solve the problem efficiently by using a [UNION](#) that combines the output

of two separate `SELECT` statements. See [Section 13.2.7.2, “UNION Syntax”](#).

Each `SELECT` searches only one key and can be optimized:

```
SELECT field1_index, field2_index
   FROM test_table WHERE field1_index = '1'
UNION
SELECT field1_index, field2_index
   FROM test_table WHERE field2_index = '1';
```

Prior to MySQL 4.0, you can achieve the same effect by using a `TEMPORARY` table and separate `SELECT` statements. This type of optimization is also very good if you are using very complicated queries where the SQL server does the optimizations in the wrong order.

```
CREATE TEMPORARY TABLE tmp
SELECT field1_index, field2_index
   FROM test_table WHERE field1_index = '1';
INSERT INTO tmp
SELECT field1_index, field2_index
   FROM test_table WHERE field2_index = '1';
SELECT * from tmp;
DROP TABLE tmp;
```

This method of solving the problem is in effect a `UNION` of two queries.

### 3.6.8. Calculating Visits Per Day

The following example shows how you can use the bit group functions to calculate the number of days per month a user has visited a Web page.

```
CREATE TABLE t1 (year YEAR(4), month INT(2) UNSIGNED ZEROFILL,
                 day INT(2) UNSIGNED ZEROFILL);
INSERT INTO t1 VALUES(2000,1,1),(2000,1,20),(2000,1,30),(2000,2,2),
                    (2000,2,23),(2000,2,23);
```

The example table contains year-month-day values representing visits by users to the page. To determine how many different days in each month these visits occur, use this query:

```
SELECT year,month,BIT_COUNT(BIT_OR(1<<day)) AS days FROM t1
   GROUP BY year,month;
```

Which returns:

```
+-----+-----+-----+
| year | month | days |
+-----+-----+-----+
| 2000 |    01 |    3 |
| 2000 |    02 |    2 |
+-----+-----+-----+
```

The query calculates how many different days appear in the table for each year/month combination, with automatic removal of duplicate entries.

### 3.6.9. Using `AUTO_INCREMENT`

The `AUTO_INCREMENT` attribute can be used to generate a unique identity for new rows:

```
CREATE TABLE animals (
  id MEDIUMINT NOT NULL AUTO_INCREMENT,
  name CHAR(30) NOT NULL,
  PRIMARY KEY (id)
);
INSERT INTO animals (name) VALUES
```

```

('dog'),('cat'),('penguin'),
('lax'),('whale'),('ostrich');
SELECT * FROM animals;

```

Which returns:

```

+-----+-----+
| id | name |
+-----+-----+
| 1 | dog |
| 2 | cat |
| 3 | penguin |
| 4 | lax |
| 5 | whale |
| 6 | ostrich |
+-----+-----+

```

You can retrieve the most recent `AUTO_INCREMENT` value with the `LAST_INSERT_ID()` SQL function or the `mysql_insert_id()` C API function. These functions are connection-specific, so their return values are not affected by another connection which is also performing inserts.

**Note:** For a multiple-row insert, `LAST_INSERT_ID()` and `mysql_insert_id()` actually return the `AUTO_INCREMENT` key from the *first* of the inserted rows. This allows multiple-row inserts to be reproduced correctly on other servers in a replication setup.

For `MyISAM` and `BDB` tables you can specify `AUTO_INCREMENT` on a secondary column in a multiple-column index. In this case, the generated value for the `AUTO_INCREMENT` column is calculated as `MAX(auto_increment_column) + 1 WHERE prefix=given-prefix`. This is useful when you want to put data into ordered groups.

```

CREATE TABLE animals (
  grp ENUM('fish','mammal','bird') NOT NULL,
  id MEDIUMINT NOT NULL AUTO_INCREMENT,
  name CHAR(30) NOT NULL,
  PRIMARY KEY (grp,id)
);

INSERT INTO animals (grp,name) VALUES
  ('mammal','dog'),('mammal','cat'),
  ('bird','penguin'),('fish','lax'),('mammal','whale'),
  ('bird','ostrich');

SELECT * FROM animals ORDER BY grp,id;

```

Which returns:

```

+-----+-----+-----+
| grp | id | name |
+-----+-----+-----+
| fish | 1 | lax |
| mammal | 1 | dog |
| mammal | 2 | cat |
| mammal | 3 | whale |
| bird | 1 | penguin |
| bird | 2 | ostrich |
+-----+-----+-----+

```

Note that in this case (when the `AUTO_INCREMENT` column is part of a multiple-column index), `AUTO_INCREMENT` values are reused if you delete the row with the biggest `AUTO_INCREMENT` value in any group. This happens even for `MyISAM` tables, for which `AUTO_INCREMENT` values normally are not reused.

If the `AUTO_INCREMENT` column is part of multiple indexes, MySQL will generate sequence values using the index that begins with the `AUTO_INCREMENT` column, if there is one. For example, if the `animals` table contained indexes `PRIMARY KEY (grp, id)` and `INDEX (id)`, MySQL would

ignore the `PRIMARY KEY` for generating sequence values. As a result, the table would contain a single sequence, not a sequence per `grp` value.

To start with an `AUTO_INCREMENT` value other than 1, you can set that value with `CREATE TABLE` or `ALTER TABLE`, like this:

```
mysql> ALTER TABLE tbl AUTO_INCREMENT = 100;
```

Note that this feature is available for `InnoDB` tables only as of MySQL 4.1.12.

More information about `AUTO_INCREMENT` is available here:

- How to assign the `AUTO_INCREMENT` attribute to a column: [Section 13.1.5, “CREATE TABLE Syntax”](#), and [Section 13.1.2, “ALTER TABLE Syntax”](#).
- How `AUTO_INCREMENT` behaves depending on the SQL mode: [Section 5.2.5, “SQL Modes”](#).
- Find the row that contains the most recent `AUTO_INCREMENT` value: [Section 12.1.3, “Comparison Functions and Operators”](#).
- Set the `AUTO_INCREMENT` value to be used: [Section 13.5.3, “SET Syntax”](#).
- `AUTO_INCREMENT` and replication: [Section 6.7, “Replication Features and Known Problems”](#).

## 3.7. Queries from the Twin Project

At Analytikerna and Lentus, we have been doing the systems and field work for a big research project. This project is a collaboration between the Institute of Environmental Medicine at Karolinska Institutet Stockholm and the Section on Clinical Research in Aging and Psychology at the University of Southern California.

The project involves a screening part where all twins in Sweden older than 65 years are interviewed by telephone. Twins who meet certain criteria are passed on to the next stage. In this latter stage, twins who want to participate are visited by a doctor/nurse team. Some of the examinations include physical and neuropsychological examination, laboratory testing, neuroimaging, psychological status assessment, and family history collection. In addition, data are collected on medical and environmental risk factors.

More information about Twin studies can be found at: [http://www.mep.ki.se/twinreg/index\\_en.html](http://www.mep.ki.se/twinreg/index_en.html)

The latter part of the project is administered with a Web interface written using Perl and MySQL.

Each night all data from the interviews is moved into a MySQL database.

### 3.7.1. Find All Non-distributed Twins

The following query is used to determine who goes into the second part of the project:

```
SELECT
  CONCAT(pl.id, pl.tvab) + 0 AS tvid,
  CONCAT(pl.christian_name, ' ', pl.surname) AS Name,
  pl.postal_code AS Code,
  pl.city AS City,
  pg.abrev AS Area,
  IF(td.participation = 'Aborted', 'A', ' ') AS A,
  pl.dead AS dead1,
  l.event AS event1,
  td.suspect AS tsuspect1,
  id.suspect AS isuspect1,
  td.severe AS tsevere1,
  id.severe AS isevere1,
  p2.dead AS dead2,
```

```

12.event AS event2,
h2.nurse AS nurse2,
h2.doctor AS doctor2,
td2.suspect AS tsuspect2,
id2.suspect AS isuspect2,
td2.severe AS tsevere2,
id2.severe AS isevere2,
1.finish_date
FROM
  twin_project AS tp
  /* For Twin 1 */
  LEFT JOIN twin_data AS td ON tp.id = td.id
    AND tp.tvab = td.tvab
  LEFT JOIN informant_data AS id ON tp.id = id.id
    AND tp.tvab = id.tvab
  LEFT JOIN harmony AS h ON tp.id = h.id
    AND tp.tvab = h.tvab
  LEFT JOIN lentus AS l ON tp.id = l.id
    AND tp.tvab = l.tvab
  /* For Twin 2 */
  LEFT JOIN twin_data AS td2 ON p2.id = td2.id
    AND p2.tvab = td2.tvab
  LEFT JOIN informant_data AS id2 ON p2.id = id2.id
    AND p2.tvab = id2.tvab
  LEFT JOIN harmony AS h2 ON p2.id = h2.id
    AND p2.tvab = h2.tvab
  LEFT JOIN lentus AS l2 ON p2.id = l2.id
    AND p2.tvab = l2.tvab,
  person_data AS p1,
  person_data AS p2,
  postal_groups AS pg
WHERE
  /* p1 gets main twin and p2 gets his/her twin. */
  /* ptvab is a field inverted from tvab */
  p1.id = tp.id AND p1.tvab = tp.tvab AND
  p2.id = p1.id AND p2.ptvab = p1.tvab AND
  /* Just the screening survey */
  tp.survey_no = 5 AND
  /* Skip if partner died before 65 but allow emigration (dead=9) */
  (p2.dead = 0 OR p2.dead = 9 OR
   (p2.dead = 1 AND
    (p2.death_date = 0 OR
     ((TO_DAYS(p2.death_date) - TO_DAYS(p2.birthday)) / 365)
     >= 65)))
  AND
  (
  /* Twin is suspect */
  (td.future_contact = 'Yes' AND td.suspect = 2) OR
  /* Twin is suspect - Informant is Blessed */
  (td.future_contact = 'Yes' AND td.suspect = 1
   AND id.suspect = 1) OR
  /* No twin - Informant is Blessed */
  (ISNULL(td.suspect) AND id.suspect = 1
   AND id.future_contact = 'Yes') OR
  /* Twin broken off - Informant is Blessed */
  (td.participation = 'Aborted'
   AND id.suspect = 1 AND id.future_contact = 'Yes') OR
  /* Twin broken off - No inform - Have partner */
  (td.participation = 'Aborted' AND ISNULL(id.suspect)
   AND p2.dead = 0))
  AND
  1.event = 'Finished'
  /* Get at area code */
  AND SUBSTRING(p1.postal_code, 1, 2) = pg.code
  /* Not already distributed */
  AND (h.nurse IS NULL OR h.nurse=00 OR h.doctor=00)
  /* Has not refused or been aborted */
  AND NOT (h.status = 'Refused' OR h.status = 'Aborted'
  OR h.status = 'Died' OR h.status = 'Other')
ORDER BY
  tvid;

```

Some explanations:

- `CONCAT(p1.id, p1.tvab) + 0 AS tvid`

We want to sort on the concatenated `id` and `tvab` in numerical order. Adding `0` to the result causes MySQL to treat the result as a number.

- column `id`

This identifies a pair of twins. It is a key in all tables.

- column `tvab`

This identifies a twin in a pair. It has a value of `1` or `2`.

- column `ptvab`

This is an inverse of `tvab`. When `tvab` is `1` this is `2`, and vice versa. It exists to save typing and to make it easier for MySQL to optimize the query.

This query demonstrates, among other things, how to do lookups on a table from the same table with a join (`p1` and `p2`). In the example, this is used to check whether a twin's partner died before the age of 65. If so, the row is not returned.

All of the above exist in all tables with twin-related information. We have a key on both `id`, `tvab` (all tables), and `id`, `ptvab` (`person_data`) to make queries faster.

On our production machine (A 200MHz UltraSPARC), this query returns about 150-200 rows and takes less than one second.

The current number of records in the tables used in the query:

Table	Rows
<code>person_data</code>	71074
<code>lentus</code>	5291
<code>twin_project</code>	5286
<code>twin_data</code>	2012
<code>informant_data</code>	663
<code>harmony</code>	381
<code>postal_groups</code>	100

### 3.7.2. Show a Table of Twin Pair Status

Each interview ends with a status code called `event`. The query shown here is used to display a table over all twin pairs combined by event. This indicates in how many pairs both twins are finished, in how many pairs one twin is finished and the other refused, and so on.

```
SELECT
    t1.event,
    t2.event,
    COUNT(*)
FROM
    lentus AS t1,
    lentus AS t2,
    twin_project AS tp
WHERE
    /* We are looking at one pair at a time */
    t1.id = tp.id
    AND t1.tvab=tp.tvab
    AND t1.id = t2.id
```

```
/* Just the screening survey */
AND tp.survey_no = 5
/* This makes each pair only appear once */
AND t1.tvab='1' AND t2.tvab='2'
GROUP BY
t1.event, t2.event;
```

## 3.8. Using MySQL with Apache

There are programs that let you authenticate your users from a MySQL database and also let you write your log files into a MySQL table.

You can change the Apache logging format to be easily readable by MySQL by putting the following into the Apache configuration file:

```
LogFormat \
"%h",%{Y%m%d%H%M%S}t,%>s,"%b",\%{Content-Type}o\, \
"%U",\%{Referer}i\,\%{User-Agent}i\"
```

To load a log file in that format into MySQL, you can use a statement something like this:

```
LOAD DATA INFILE '/local/access_log' INTO TABLE tbl_name
FIELDS TERMINATED BY ',' OPTIONALLY ENCLOSED BY '"' ESCAPED BY '\\'
```

The named table should be created to have columns that correspond to those that the `LogFormat` line writes to the log file.



---

# Chapter 4. Using MySQL Programs

This chapter provides a brief overview of the command-line programs provided by MySQL AB and discusses the general syntax for specifying options when you run these programs. Most programs have options that are specific to their own operation, but the option syntax is similar for all of them. Later chapters provide more detailed descriptions of individual programs, including which options they recognize.

MySQL AB also provides three GUI client programs for use with MySQL Server:

- **MySQL Administrator:** This tool is used for administering MySQL servers, databases, tables, and user accounts.
- **MySQL Query Browser:** This graphical tool is provided by MySQL AB for creating, executing, and optimizing queries on MySQL databases.
- **MySQL Migration Toolkit:** This tool helps you migrate schemas and data from other relational database management systems for use with MySQL.

These GUI programs each have their own manuals that you can access at <http://dev.mysql.com/doc/>.

## 4.1. Overview of MySQL Programs

MySQL AB provides several types of programs:

- The MySQL server and server startup scripts:
  - `mysqld` is the MySQL server.
  - `mysqld_safe`, `mysql.server`, and `mysqld_multi` are server startup scripts.
  - `mysql_install_db` initializes the data directory and the initial databases.
- Client programs that access the server:
  - `mysql` is a command-line client for executing SQL statements interactively or in batch mode.
  - `mysqladmin` is an administrative client.
  - `mysqlcheck` performs table maintenance operations.
  - `mysqldump` and `mysqlhotcopy` make database backups.
  - `mysqlimport` imports data files.
  - `mysqlshow` displays information about databases and tables.

[Chapter 5, \*Database Administration\*](#), discusses these programs further

[Chapter 8, \*Client and Utility Programs\*](#), discusses these programs further

- Utility programs that operate independently of the server:
  - `myisamchk` performs table maintenance operations.

- `mysampack` produces compressed, read-only tables.
- `mysqlbinlog` is a tool for processing binary log files.
- `perror` displays the meaning of error codes.

Chapter 8, *Client and Utility Programs*, discusses these programs further

Most MySQL distributions include all of these programs, except for those programs that are platform-specific. (For example, the server startup scripts are not used on Windows.) The exception is that RPM distributions are more specialized. There is one RPM for the server, another for client programs, and so forth. If you appear to be missing one or more programs, see [Chapter 2, \*Installing and Upgrading MySQL\*](#), for information on types of distributions and what they contain. It may be that you have a distribution that does not include all programs and you need to install something else.

## 4.2. Invoking MySQL Programs

To invoke a MySQL program from the command line (that is, from your shell or command prompt), enter the program name followed by any options or other arguments needed to instruct the program what you want it to do. The following commands show some sample program invocations. “`shell>`” represents the prompt for your command interpreter; it is not part of what you type. The particular prompt you see depends on your command interpreter. Typical prompts are `$` for `sh` or `bash`, `%` for `csch` or `tcsh`, and `C:\>` for the Windows `command.com` or `cmd.exe` command interpreters.

```
shell> mysql -u root test
shell> mysqladmin extended-status variables
shell> mysqlshow --help
shell> mysqldump --user=root personnel
```

Arguments that begin with a single or double dash (`'-`', `'--'`) are option arguments. Options typically specify the type of connection a program should make to the server or affect its operational mode. Option syntax is described in [Section 4.3, “Specifying Program Options”](#).

Non-option arguments (arguments with no leading dash) provide additional information to the program. For example, the `mysql` program interprets the first non-option argument as a database name, so the command `mysql -u root test` indicates that you want to use the `test` database.

Later sections that describe individual programs indicate which options a program understands and describe the meaning of any additional non-option arguments.

Some options are common to a number of programs. The most common of these are the `--host` (or `-h`), `--user` (or `-u`), and `--password` (or `-p`) options that specify connection parameters. They indicate the host where the MySQL server is running, and the username and password of your MySQL account. All MySQL client programs understand these options; they allow you to specify which server to connect to and the account to use on that server.

Other connection options are `--port` (or `-P`) to specify a TCP/IP port number and `--socket` (or `-S`) to specify a Unix socket file on Unix (or named pipe name on Windows).

The default hostname is `localhost`. For client programs on Unix, the hostname `localhost` is special. It causes the client to connect to the MySQL server through a Unix socket file. This occurs even if a `--port` or `-P` option is given to specify a port number. To ensure that the client makes a TCP/IP connection to the local server, use `--host` or `-h` to specify a hostname value of `127.0.0.1`, or the IP address or name of the local server. You can also specify the connection protocol explicitly, even for `localhost`, by using the `--protocol=tcp` option.

You may find it necessary to invoke MySQL programs using the pathname to the `bin` directory in which they are installed. This is likely to be the case if you get a “program not found” error whenever you attempt to run a MySQL program from any directory other than the `bin` directory. To make it more convenient to use MySQL, you can add the pathname of the `bin` directory to your `PATH` environment variable setting. That enables you to run a program by typing only its name, not its entire pathname. For example, if `mysql` is installed in `/usr/local/mysql/bin`, you'll be able to run it by invoking it as `mysql`; it will not be necessary to invoke it as `/usr/local/mysql/bin/mysql`.

Consult the documentation for your command interpreter for instructions on setting your `PATH` variable. The syntax for setting environment variables is interpreter-specific. (Some information is given in [Section 4.3.3, “Using Environment Variables to Specify Options”](#).)

## 4.3. Specifying Program Options

There are several ways to specify options for MySQL programs:

- List the options on the command line following the program name. This is most common for options that apply to a specific invocation of the program.
- List the options in an option file that the program reads when it starts. This is common for options that you want the program to use each time it runs.
- List the options in environment variables. This method is useful for options that you want to apply each time the program runs. In practice, option files are used more commonly for this purpose. However, [Section 5.12.2, “Running Multiple Servers on Unix”](#), discusses one situation in which environment variables can be very helpful. It describes a handy technique that uses such variables to specify the TCP/IP port number and Unix socket file for both the server and client programs.

MySQL programs determine which options are given first by examining environment variables, then by reading option files, and then by checking the command line. This means that environment variables have the lowest precedence and command-line options the highest.

Because options are processed in order, if an option is specified multiple times, the last occurrence takes precedence. The following command causes `mysql` to connect to the server running on `localhost`:

```
shell> mysql -h example.com -h localhost
```

If conflicting or related options are given, later options take precedence over earlier options. The following command runs `mysql` in “no column names” mode:

```
shell> mysql --column-names --skip-column-names
```

An option can be specified by writing it in full or as any unambiguous prefix. For example, the `-compress` option can be given to `mysqldump` as `--compr`, but not as `--comp` because that is ambiguous:

```
shell> mysqldump --comp
mysqldump: ambiguous option '--comp' (compatible, compress)
```

Be aware that the use of option prefixes can cause problems in the event that new options are implemented for a program. A prefix that is unambiguous now might become ambiguous in the future.

You can take advantage of the way that MySQL programs process options by specifying default values for a program's options in an option file. That enables you to avoid typing them each time you run the program, but also allows you to override the defaults if necessary by using command-line options.

### 4.3.1. Using Options on the Command Line

Program options specified on the command line follow these rules:

- Options are given after the command name.
- An option argument begins with one dash or two dashes, depending on whether it has a short name or a long name. Many options have both forms. For example, `-?` and `--help` are the short and long forms of the option that instructs a MySQL program to display its help message.
- Option names are case sensitive. `-v` and `-V` are both legal and have different meanings. (They are the corresponding short forms of the `--verbose` and `--version` options.)
- Some options take a value following the option name. For example, `-h localhost` or `-host=localhost` indicate the MySQL server host to a client program. The option value tells the program the name of the host where the MySQL server is running.
- For a long option that takes a value, separate the option name and the value by an '=' sign. For a short option that takes a value, the option value can immediately follow the option letter, or there can be a space between: `-hlocalhost` and `-h localhost` are equivalent. An exception to this rule is the option for specifying your MySQL password. This option can be given in long form as `-password=pass_val` or as `--password`. In the latter case (with no password value given), the program prompts you for the password. The password option also may be given in short form as `-ppass_val` or as `-p`. However, for the short form, if the password value is given, it must follow the option letter with *no intervening space*. The reason for this is that if a space follows the option letter, the program has no way to tell whether a following argument is supposed to be the password value or some other kind of argument. Consequently, the following two commands have two completely different meanings:

```
shell> mysql -ptest
shell> mysql -p test
```

The first command instructs `mysql` to use a password value of `test`, but specifies no default database. The second instructs `mysql` to prompt for the password value and to use `test` as the default database.

MySQL 4.0 introduced some additional flexibility in the way you specify options. These changes were made in MySQL 4.0.2. Some of them relate to the way you specify options that have “enabled” and “disabled” states, and to the use of options that might be present in one version of MySQL but not another. Those capabilities are discussed in this section. Another change pertains to the way you use options to set program variables. [Section 4.3.4, “Using Options to Set Program Variables”](#), discusses that topic further.

Some options control behavior that can be turned on or off. For example, the `mysql` client supports a `-column-names` option that determines whether or not to display a row of column names at the beginning of query results. By default, this option is enabled. However, you may want to disable it in some instances, such as when sending the output of `mysql` into another program that expects to see only data and not an initial header line.

To disable column names, you can specify the option using any of these forms:

```
--disable-column-names
--skip-column-names
--column-names=0
```

The `--disable` and `--skip` prefixes and the `=0` suffix all have the same effect: They turn the option off.

The “enabled” form of the option may be specified in any of these ways:

```
--column-names
--enable-column-names
--column-names=1
```

Another change to option processing introduced in MySQL 4.0 is that you can use the `--loose` prefix for command-line options. If an option is prefixed by `--loose`, a program does not exit with an error if it does not recognize the option, but instead issues only a warning:

```
shell> mysql --loose-no-such-option
mysql: WARNING: unknown option '--no-such-option'
```

The `--loose` prefix can be useful when you run programs from multiple installations of MySQL on the same machine and list options in an option file. An option that may not be recognized by all versions of a program can be given using the `--loose` prefix (or `loose` in an option file). Versions of the program that recognize the option process it normally, and versions that do not recognize it issue a warning and ignore it. This strategy requires that versions involved be 4.0.2 or later, because earlier versions know nothing of the `--loose` convention.

Another option that may occasionally be useful with `mysql` is the `--execute` or `-e` option, which can be used to pass SQL statements to the server. The statements must be enclosed by single or double quotation marks. If you wish to use quoted values within a statement, you should use double quotes for the statement, and single quotes for any quoted values within the statement. When this option is used, `mysql` executes the statements and exits.

For example, you can use the following command to obtain a list of user accounts:

```
shell> mysql -u root -p --execute="SELECT User, Host FROM user" mysql
Enter password: *****
+-----+-----+
| User | Host |
+-----+-----+
| root | gigan |
|      | gigan |
|      | localhost |
| jon  | localhost |
| root | localhost |
+-----+-----+
shell>
```

Note that the long form (`--execute`) is followed by an equals sign (=).

In the preceding example, the name of the `mysql` database was passed as a separate argument. However, the same statement could have been executed using this command, which specifies no default database:

```
mysql> mysql -u root -p --execute="SELECT User, Host FROM mysql.user"
```

Multiple SQL statements may be passed on the command line, separated by semicolons:

```
shell> mysql -u root -p -e "SELECT VERSION();SELECT NOW()"
Enter password: *****
+-----+
| VERSION() |
+-----+
| 4.1.17-log |
+-----+
+-----+
| NOW() |
+-----+
| 2006-01-05 21:19:04 |
+-----+
```

The `--execute` or `-e` option may also be used to pass commands in an analogous fashion to the `ndb_mgm` management client for MySQL Cluster. See [Section 15.3.6, “Safe Shutdown and Restart”](#), for an example.

## 4.3.2. Using Option Files

Most MySQL programs can read startup options from option files (also sometimes called configuration files). Option files provide a convenient way to specify commonly used options so that they need not be entered on the command line each time you run a program. Option file capability is available from MySQL 3.22 on. For the MySQL server, MySQL provides a number of [preconfigured option files](#).

To determine whether a program reads option files, invoke it with the `--help` option (`--verbose` and `--help` for `mysqld` as of MySQL 4.1.1). If the program reads option files, the help message indicates which files it looks for and which option groups it recognizes.

**Note:** Option files used with MySQL Cluster programs are covered in [Section 15.4, “MySQL Cluster Configuration”](#).

On Windows, MySQL programs read startup options from the following files:

Filename	Purpose
<code>WINDIR\my.ini</code>	Global options
<code>C:\my.cnf</code>	Global options
<code>INSTALLDIR\my.ini</code>	Global Options
<code>defaults-extra-file</code>	The file specified with <code>--defaults-extra-file=path</code> , if any

`WINDIR` represents the location of your Windows directory. This is commonly `C:\WINDOWS` or `C:\WINNT`. You can determine its exact location from the value of the `WINDIR` environment variable using the following command:

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

`INSTALLDIR` represents the installation directory of MySQL. This is typically the case with MySQL 4.1.5 and higher, when installed using the installation and configuration wizards. See [Section 2.3.4.14, “The Location of the my.ini File”](#).

On Unix, MySQL programs read startup options from the following files:

Filename	Purpose
<code>/etc/my.cnf</code>	Global options
<code>DATADIR/my.cnf</code>	Server-specific options
<code>defaults-extra-file</code>	The file specified with <code>--defaults-extra-file=path</code> , if any
<code>~/.my.cnf</code>	User-specific options

`DATADIR` represents the path to the directory in which the server-specific `my.cnf` file resides.

If `MYSQL_HOME` is not set and you start the server using the `mysqld_safe` program, `mysqld_safe` attempts to set `MYSQL_HOME` as follows:

- Let `BASEDIR` and `DATADIR` represent the pathnames of the MySQL base directory and data directory, respectively.

- If there is a `my.cnf` file in `DATADIR` but not in `BASEDIR`, `mysqld_safe` sets `MYSQL_HOME` to `DATADIR`.
- Otherwise, if `MYSQL_HOME` is not set and there is no `my.cnf` file in `DATADIR`, `mysqld_safe` sets `MYSQL_HOME` to `BASEDIR`.

Typically, `DATADIR` is `/usr/local/mysql/data` for a binary installation or `/usr/local/var` for a source installation. Note that this is the data directory location that was specified at configuration time, not the one specified with the `--datadir` option when `mysqld` starts. Use of `--datadir` at runtime has no effect on where the server looks for option files, because it looks for them before processing any options.

MySQL looks for option files in the order just described and reads any that exist. If an option file that you want to use does not exist, create it with a plain text editor.

If multiple instances of a given option are found, the last instance takes precedence. There is one exception: For `mysqld`, the *first* instance of the `--user` option is used as a security precaution, to prevent a user specified in an option file from being overridden on the command line.

**Note:** On Unix platforms, MySQL ignores configuration files that are world-writable. This is intentional, and acts as a security measure.

Any long option that may be given on the command line when running a MySQL program can be given in an option file as well. To get the list of available options for a program, run it with the `--help` option.

The syntax for specifying options in an option file is similar to command-line syntax, except that you omit the leading two dashes. For example, `--quick` or `--host=localhost` on the command line should be specified as `quick` or `host=localhost` in an option file. To specify an option of the form `--loose-opt_name` in an option file, write it as `loose-opt_name`.

Empty lines in option files are ignored. Non-empty lines can take any of the following forms:

- `#comment, ;comment`

Comment lines start with `#` or `;`. As of MySQL 4.0.14, a `#` comment can start in the middle of a line as well.

- `[group]`

`group` is the name of the program or group for which you want to set options. After a group line, any option-setting lines apply to the named group until the end of the option file or another group line is given.

- `opt_name`

This is equivalent to `--opt_name` on the command line.

- `opt_name=value`

This is equivalent to `--opt_name=value` on the command line. In an option file, you can have spaces around the `=` character, something that is not true on the command line. As of MySQL 4.0.16, you can enclose the value within double quotes or single quotes. This is useful if the value contains a `#` comment character or whitespace.

- `set-variable = var_name=value`

Set the program variable `var_name` to the given value. This is equivalent to `--set-variable=var_name=value` on the command line. Spaces are allowed around the first '=' character but not around the second. This syntax is deprecated as of MySQL 4.0. See [Section 4.3.4, “Using Options to Set Program Variables”](#), for more information on setting program variables.

For options that take a numeric value, the value can be given with a suffix of **K**, **M**, or **G** (either uppercase or lowercase) to indicate a multiplier of 1024, 1024<sup>2</sup> or 1024<sup>3</sup>. For example, the following command tells `mysqladmin` to ping the server 1024 times, sleeping 10 seconds between each ping:

```
mysql> mysqladmin --count=1K --sleep=10 ping
```

Leading and trailing blanks are automatically deleted from option names and values. You may use the escape sequences `\b`, `\t`, `\n`, `\r`, `\\`, and `\s` in option values to represent the backspace, tab, newline, carriage return, backslash, and space characters.

Because the `\\` escape sequence represents a single backslash, you must write each `\` as `\\`. Alternatively, you can specify the value using `/` rather than `\` as the pathname separator.

If an option group name is the same as a program name, options in the group apply specifically to that program. For example, the `[mysqld]` and `[mysql]` groups apply to the `mysqld` server and the `mysql` client program, respectively.

The `[client]` option group is read by all client programs (but *not* by `mysqld`). This allows you to specify options that apply to all clients. For example, `[client]` is the perfect group to use to specify the password that you use to connect to the server. (But make sure that the option file is readable and writable only by yourself, so that other people cannot find out your password.) Be sure not to put an option in the `[client]` group unless it is recognized by *all* client programs that you use. Programs that do not understand the option quit after displaying an error message if you try to run them.

Here is a typical global option file:

```
[client]
port=3306
socket=/tmp/mysql.sock

[mysqld]
port=3306
socket=/tmp/mysql.sock
key_buffer_size=16M
max_allowed_packet=8M

[mysqldump]
quick
```

The preceding option file uses `var_name=value` syntax for the lines that set the `key_buffer_size` and `max_allowed_packet` variables. Prior to MySQL 4.0.2, you would need to use `set-variable` syntax instead (described earlier in this section).

Here is a typical user option file:

```
[client]
# The following password will be sent to all standard MySQL clients
password="my_password"

[mysql]
no-auto-rehash
set-variable = connect_timeout=2

[mysqlhotcopy]
interactive-timeout
```



This option file uses `set-variable` syntax to set the `connect_timeout` variable. For MySQL 4.0.2 and up, you can also set the variable using just `connect_timeout=2`.

As of MySQL 4.0.14, if you want to create option groups that should be read only by `mysqld` servers from a specific MySQL release series only, you can do this by using groups with names of `[mysqld-4.0]`, `[mysqld-4.1]`, and so forth. The following group indicates that the `--new` option should be used only by MySQL servers with 4.0.x version numbers:

```
[mysqld-4.0]
new
```

Beginning with MySQL 4.1.11 in the 4.1 series and MySQL 5.0.4 in the 5.0 series, it is possible to use `!include` directives in option files to include other option files and `!includedir` to search specific directories for option files. For example, to include the `/home/mydir/myopt.cnf` file, you can use the following directive:

```
!include /home/me/myopt.cnf
```

To search the `/home/mydir` directory and read option files found there, you would use this directive:

```
!includedir /home/mydir
```

**Note:** Currently, any files to be found and included using the `!includedir` directive on Unix operating systems *must* have filenames ending in `.cnf`. On Windows, this directive checks for files with the `.ini` or `.cnf` extension.

Note that options read from included files are applied in the context of the current option group. Suppose that you were to write the following lines in `my.cnf`:

```
[mysqld]
!include /home/mydir/myopt.cnf
```

In this case, the `myopt.cnf` file is processed only for the server, and the `!include` directive is ignored by any client applications. However, if you were to use the following lines, the directory `/home/mydir/my-dump-options` is checked for option files by `mysqldump` only, and not by the server or by any other client applications:

```
[mysqldump]
!includedir /home/mydir/my-dump-options
```

If you have a source distribution, you can find sample option files named `my-xxxx.cnf` in the `support-files` directory. If you have a binary distribution, look in the `support-files` directory under your MySQL installation directory. On Windows, the sample option files may be located in the MySQL installation directory (see earlier in this section or [Chapter 2, Installing and Upgrading MySQL](#), if you do not know where this is). Currently, there are sample option files for small, medium, large, and very large systems. To experiment with one of these files, copy it to `C:\my.cnf` on Windows or to `.my.cnf` in your home directory on Unix.

**Note:** On Windows, the `.cnf` option file extension might not be displayed.

All MySQL programs that support option files handle the following options. They affect option-file handling, so they must be given on the command line and not in an option file. To work properly, each of these options must immediately follow the command name, with the exception that `-print-defaults` may be used immediately after `--defaults-file` or `--defaults-extra-file`.

- `--no-defaults`

Don't read any option files.

- `--print-defaults`

Print the program name and all options that it gets from option files.

- `--defaults-file=file_name`

Use only the given option file. *file\_name* is the full pathname to the file. If the file does not exist, the program exits with an error.

- `--defaults-extra-file=file_name`

Read this option file after the global option file but (on Unix) before the user option file.

*file\_name* is the full pathname to the file.

In shell scripts, you can use the `my_print_defaults` program to parse option files and see what options would be used by a given program. The following example shows the output that `my_print_defaults` might produce when asked to show the options found in the `[client]` and `[mysql]` groups:

```
shell> my_print_defaults client mysql
--port=3306
--socket=/tmp/mysql.sock
--no-auto-rehash
```

**Note for developers:** Option file handling is implemented in the C client library simply by processing all options in the appropriate group or groups before any command-line arguments. This works well for programs that use the last instance of an option that is specified multiple times. If you have a C or C++ program that handles multiply specified options this way but that doesn't read option files, you need add only two lines to give it that capability. Check the source code of any of the standard MySQL clients to see how to do this.

Several other language interfaces to MySQL are based on the C client library, and some of them provide a way to access option file contents. These include Perl and Python. For details, see the documentation for your preferred interface.

### 4.3.2.1. Preconfigured Option Files

MySQL provides a number of preconfigured option files that can be used as a basis for tuning the MySQL server. Look in your installation directory for files such as `my-small.cnf`, `my-medium.cnf`, `my-large.cnf`, and `my-huge.cnf`, which you can rename and copy to the appropriate location for use as a base configuration file. Regarding names and appropriate location, see the general information provided in [Section 4.3.2, "Using Option Files"](#). On Windows, those files have a `.ini` rather than a `.cnf` extension.

### 4.3.3. Using Environment Variables to Specify Options

To specify an option using an environment variable, set the variable using the syntax appropriate for your command processor. For example, on Windows or NetWare, you can set the `USER` variable to specify your MySQL account name. To do so, use this syntax:

```
SET USER=your_name
```

The syntax on Unix depends on your shell. Suppose that you want to specify the TCP/IP port number

using the `MYSQL_TCP_PORT` variable. Typical syntax (such as for `sh`, `bash`, `zsh`, and so on) is as follows:

```
MYSQL_TCP_PORT=3306
export MYSQL_TCP_PORT
```

The first command sets the variable, and the `export` command exports the variable to the shell environment so that its value becomes accessible to MySQL and other processes.

For `csh` and `tcsh`, use `setenv` to make the shell variable available to the environment:

```
setenv MYSQL_TCP_PORT 3306
```

The commands to set environment variables can be executed at your command prompt to take effect immediately, but the settings persist only until you log out. To have the settings take effect each time you log in, place the appropriate command or commands in a startup file that your command interpreter reads each time it starts. Typical startup files are `AUTOEXEC.BAT` for Windows, `.bash_profile` for `bash`, or `.tcshrc` for `tcsh`. Consult the documentation for your command interpreter for specific details.

[Appendix F, \*Environment Variables\*](#), lists all environment variables that affect MySQL program operation.

### 4.3.4. Using Options to Set Program Variables

Many MySQL programs have internal variables that can be set at runtime. As of MySQL 4.0.2, program variables are set the same way as any other long option that takes a value. For example, `mysql` has a `max_allowed_packet` variable that controls the maximum size of its communication buffer. To set the `max_allowed_packet` variable for `mysql` to a value of 16MB, use either of the following commands:

```
shell> mysql --max_allowed_packet=16777216
shell> mysql --max_allowed_packet=16M
```

The first command specifies the value in bytes. The second specifies the value in megabytes. For variables that take a numeric value, the value can be given with a suffix of `K`, `M`, or `G` (either uppercase or lowercase) to indicate a multiplier of  $1024$ ,  $1024^2$  or  $1024^3$ . (For example, when used to set `max_allowed_packet`, the suffixes indicate units of kilobytes, megabytes, or gigabytes.)

In an option file, variable settings are given without the leading dashes:

```
[mysql]
max_allowed_packet=16777216
```

Or:

```
[mysql]
max_allowed_packet=16M
```

If you like, underscores in a variable name can be specified as dashes. The following option groups are equivalent. Both set the size of the server's key buffer to 512MB:

```
[mysqld]
key_buffer_size=512M

[mysqld]
key-buffer-size=512M
```

Prior to MySQL 4.0.2, program variable names are not recognized as option names. Instead, use the `--set-variable` option to assign a value to a variable:

```
shell> mysql --set-variable=max_allowed_packet=16777216
shell> mysql --set-variable=max_allowed_packet=16M
```

In an option file, omit the leading dashes:

```
[mysql]
set-variable = max_allowed_packet=16777216
```

Or:

```
[mysql]
set-variable = max_allowed_packet=16M
```

With `--set-variable`, underscores in variable names cannot be given as dashes for versions of MySQL older than 4.0.2.

The `--set-variable` option is still recognized in MySQL 4.0.2 and up, but is deprecated.

Many server system variables can also be set at runtime. For details, see [Section 5.2.3.2, “Dynamic System Variables”](#).

---

# Chapter 5. Database Administration

This chapter covers topics that deal with administering a MySQL installation:

- Configuring the server
- Managing user accounts
- Performing backups
- The server log files
- The query cache

## 5.1. Overview of Server-Side Programs

The MySQL server, `mysqld`, is the main program that does most of the work in a MySQL installation. The server is accompanied by several related scripts that perform setup operations when you install MySQL or that assist you in starting and stopping the server. This section provides an overview of the server and related programs. The following sections provide more detailed information about each of these programs.

Each MySQL program takes many different options. Most programs provide a `--help` option that you can use to get a description of the program's different options. For example, try `mysqld --help`.

You can override default option values for MySQL programs by specifying options on the command line or in an option file. [Section 4.3, “Specifying Program Options”](#).

The following list briefly describes the MySQL server and server-related programs:

- `mysqld`

The SQL daemon (that is, the MySQL server). To use client programs, `mysqld` must be running, because clients gain access to databases by connecting to the server. See [Section 5.2, “mysqld — The MySQL Server”](#).

- `mysqld-max`

A version of the server that includes additional features. See [Section 5.3, “The mysqld-max Extended MySQL Server”](#).

- `mysqld_safe`

A server startup script. `mysqld_safe` attempts to start `mysqld-max` if it exists, and `mysqld` otherwise. See [Section 5.4.1, “mysqld\\_safe — MySQL Server Startup Script”](#).

- `mysql.server`

A server startup script. This script is used on systems that use System V-style run directories containing scripts that start system services for particular run levels. It invokes `mysqld_safe` to start the MySQL server. See [Section 5.4.2, “mysql.server — MySQL Server Startup Script”](#).

- `mysqld_multi`

A server startup script that can start or stop multiple servers installed on the system. See [Sec-](#)

tion 5.4.3, “[mysqld\\_multi — Manage Multiple MySQL Servers](#)”.

- [mysql\\_install\\_db](#)

This script creates the MySQL database and initializes the grant tables with default privileges. It is usually executed only once, when first installing MySQL on a system. See [Section 2.10.2, “Unix Post-Installation Procedures”](#).

- [mysql\\_fix\\_privilege\\_tables](#)

This script is used after a MySQL upgrade operation. It updates the grant tables with any changes that have been made in newer versions of MySQL. See [Section 5.5.1, “mysql\\_fix\\_privilege\\_tables — Upgrade MySQL System Tables”](#).

There are several other programs that are run on the server host:

- [make\\_binary\\_distribution](#)

This program makes a binary release of a compiled MySQL. This could be sent by FTP to `/pub/mysql/upload/` on `ftp.mysql.com` for the convenience of other MySQL users.

## 5.2. [mysqld](#) — The MySQL Server

[mysqld](#) is the MySQL server. The following discussion covers these MySQL server configuration topics:

- Startup options that the server supports
- Server system variables
- Server status variables
- How to set the server SQL mode
- The server shutdown process

**Note:** Not all storage engines (also known in older versions of MySQL as “table types”) are supported by all MySQL server binaries and configurations. To find out how to determine which storage engines are supported by your MySQL server installation, see [Section 13.5.4.8, “SHOW ENGINES Syntax”](#).

### 5.2.1. Command Options

When you start the [mysqld](#) server, you can specify program options using any of the methods described in [Section 4.3, “Specifying Program Options”](#). The most common methods are to provide options in an option file or on the command line. However, in most cases it is desirable to make sure that the server uses the same options each time it runs. The best way to ensure this is to list them in an option file. See [Section 4.3.2, “Using Option Files”](#).

[mysqld](#) reads options from the `[mysqld]` and `[server]` groups. [mysqld\\_safe](#) reads options from the `[mysqld]`, `[server]`, `[mysqld_safe]`, and `[safe_mysqld]` groups. `mysql.server` reads options from the `[mysqld]` and `[mysql.server]` groups.

An embedded MySQL server usually reads options from the `[server]`, `[embedded]`, and

`[_SERVER]` groups, where `xxxxxx` is the name of the application into which the server is embedded.

`mysqld` accepts many command options. For a list, execute `mysqld --help`. Before MySQL 4.1.1, `--help` prints the full help message. As of 4.1.1, it prints a brief message; to see the full list, use `mysqld --verbose --help`.

The following list shows some of the most common server options. Additional options are described in other sections:

- Options that affect security: See [Section 5.6.3, “Security-Related `mysqld` Options”](#).
- SSL-related options: See [Section 5.8.7.3, “SSL Command Options”](#).
- Binary log control options: See [Section 5.11.4, “The Binary Log”](#).
- Replication-related options: See [Section 6.8, “Replication Startup Options”](#).
- Options specific to particular storage engines: See [Section 14.1.1, “MyISAM Startup Options”](#), [Section 14.5.3, “BDB Startup Options”](#), [Section 14.2.5, “InnoDB Startup Options and System Variables”](#), and [Section 15.6.5.1, “MySQL Cluster-Related Command Options for `mysqld`”](#).

You can also set the values of server system variables by using variable names as options, as described later in this section.

- `--help, -?`

Display a short help message and exit. Before MySQL 4.1.1, `--help` displays the full help message. As of 4.1.1, it displays an abbreviated message only. Use both the `--verbose` and `--help` options to see the full message.

- `--allow-suspicious-udfs`

This option controls whether user-defined functions that have only an `xxx` symbol for the main function can be loaded. By default, the option is off and only UDFs that have at least one auxiliary symbol can be loaded; this prevents attempts at loading functions from shared object files other than those containing legitimate UDFs. This option was added in MySQL 4.0.24, and 4.1.10a. See [Section 19.2.4.6, “User-Defined Function Security Precautions”](#).

- `--ansi`

Use standard (ANSI) SQL syntax instead of MySQL syntax. For more precise control over the server SQL mode, use the `--sql-mode` option instead. See [Section 1.9.3, “Running MySQL in ANSI Mode”](#), and [Section 5.2.5, “SQL Modes”](#).

- `--basedir=path, -b path`

The path to the MySQL installation directory. All paths are usually resolved relative to this directory.

- `--big-tables`

Allow large result sets by saving all temporary sets in files. This option prevents most “table full” errors, but also slows down queries for which in-memory tables would suffice. Since MySQL 3.23.2, the server is able to handle large result sets automatically by using memory for small temporary tables and switching to disk tables where necessary.

- `--bind-address=IP`

The IP address to bind to.

- `--bootstrap`

This option is used by the `mysql_install_db` script to create the MySQL privilege tables without having to start a full MySQL server.

- `--character-sets-dir=path`

The directory where character sets are installed. See [Section 5.10.1, “The Character Set Used for Data and Sorting”](#).

- `--character-set-client-handshake`

Don't ignore character set information sent by the client. To ignore client information and use the default server character set, use `--skip-character-set-client-handshake`; this makes MySQL 4.1 and higher behave like MySQL 4.0. This option was added in MySQL 4.1.15.

- `--character-set-server=charset_name, -C charset_name`

Use `charset_name` as the default server character set. See [Section 5.10.1, “The Character Set Used for Data and Sorting”](#). If you use this option to specify a non-default character set, you should also use `--collation-server` to specify the collation. This option is available as of MySQL 4.1.3.

- `--chroot=path`

Put the `mysqld` server in a closed environment during startup by using the `chroot()` system call. This is a recommended security measure as of MySQL 4.0. (MySQL 3.23 is not able to provide a `chroot()` jail that is 100% closed.) Note that use of this option somewhat limits `LOAD DATA INFILE` and `SELECT ... INTO OUTFILE`.

- `--collation-server=collation_name`

Use `collation_name` as the default server collation. This option is available as of MySQL 4.1.3. See [Section 5.10.1, “The Character Set Used for Data and Sorting”](#).

- `--console`

(Windows only.) Write error log messages to `stderr` and `stdout` even if `--log-error` is specified. `mysqld` does not close the console window if this option is used.

- `--core-file`

Write a core file if `mysqld` dies. For some systems, you must also specify the `--core-file-size` option to `mysqld_safe`. See [Section 5.4.1, “mysqld\\_safe — MySQL Server Startup Script”](#). Note that on some systems, such as Solaris, you do not get a core file if you are also using the `--user` option.

- `--datadir=path, -h path`

The path to the data directory.

- `--debug[=debug_options], -# [debug_options]`

If MySQL is configured with `--with-debug`, you can use this option to get a trace file of what `mysqld` is doing. The `debug_options` string often is `'d:t:o,file_name'`. The default is `'d:t:i:o,mysqld.trace'`. See [Section E.1.2, “Creating Trace Files”](#).



- `--default-character-set=charset_name, -C charset_name`

Use `charset_name` as the default character set. This option is deprecated in favor of `-character-set-server` as of MySQL 4.1.3. See [Section 5.10.1, “The Character Set Used for Data and Sorting”](#).
- `--default-collation=collation_name`

Use `collation_name` as the default collation. This option is deprecated in favor of `-collation-server` as of MySQL 4.1.3. See [Section 5.10.1, “The Character Set Used for Data and Sorting”](#).
- `--default-storage-engine=type`

This option is a synonym for `--default-table-type`. It is available as of MySQL 4.1.2.
- `--default-table-type=type`

Set the default table type (storage engine) for tables. See [Chapter 14, Storage Engines and Table Types](#).
- `--default-time-zone=timezone`

Set the default server time zone. This option sets the global `time_zone` system variable. If this option is not given, the default time zone is the same as the system time zone (given by the value of the `system_time_zone` system variable). This option is available as of MySQL 4.1.3.
- `--delay-key-write[={OFF|ON|ALL}]`

Specify how to use delayed key writes. Delayed key writing causes key buffers not to be flushed between writes for `MyISAM` tables. `OFF` disables delayed key writes. `ON` enables delayed key writes for those tables that were created with the `DELAY_KEY_WRITE` option. `ALL` delays key writes for all `MyISAM` tables. Available as of MySQL 4.0.3. See [Section 7.5.2, “Tuning Server Parameters”](#), and [Section 14.1.1, “MyISAM Startup Options”](#).

**Note:** If you set this variable to `ALL`, you should not use `MyISAM` tables from within another program (such as another MySQL server or `myisamchk`) when the tables are in use. Doing so leads to index corruption.
- `--delay-key-write-for-all-tables`

Old form of `--delay-key-write=ALL` for use prior to MySQL 4.0.3. As of 4.0.3, use `--delay-key-write` instead.
- `--des-key-file=file_name`

Read the default DES keys from this file. These keys are used by the `DES_ENCRYPT()` and `DES_DECRYPT()` functions.
- `--enable-named-pipe`

Enable support for named pipes. This option applies only on Windows NT, 2000, XP, and 2003 systems, and can be used only with the `mysqld-nt` and `mysqld-max-nt` servers that support named-pipe connections.
- `--exit-info[=flags], -T [flags]`

This is a bit mask of different flags that you can use for debugging the `mysqld` server. Do not use this option unless you know *exactly* what it does!

- `--external-locking`

Enable external locking (system locking), which is disabled by default as of MySQL 4.0. Note that if you use this option on a system on which `lockd` does not fully work (such as Linux), it is easy for `mysqld` to deadlock. This option was named `--enable-locking` before MySQL 4.0.3.

**Note:** If you use this option to enable updates to `MyISAM` tables from many MySQL processes, you must ensure that the following conditions are satisfied:

- You should not use the query cache for queries that use tables that are updated by another process.
- You should not use `--delay-key-write=ALL` or `DELAY_KEY_WRITE=1` on any shared tables.

The easiest way to ensure this is to always use `--external-locking` together with `--delay-key-write=OFF` and `--query-cache-size=0`. (This is not done by default because in many setups it is useful to have a mixture of the preceding options.)

- `--flush`

Flush (synchronize) all changes to disk after each SQL statement. Normally, MySQL does a write of all changes to disk only after each SQL statement and lets the operating system handle the synchronizing to disk. See [Section A.4.2, “What to Do If MySQL Keeps Crashing”](#).

- `--init-file=file_name`

Read SQL statements from this file at startup. Each statement must be on a single line and should not include comments.

- `--innodb-safe-binlog`

Adds consistency guarantees between the content of `InnoDB` tables and the binary log. See [Section 5.11.4, “The Binary Log”](#).

- `--innodb-xxx`

The `InnoDB` options are listed in [Section 14.2.5, “InnoDB Startup Options and System Variables”](#).

- `--language=lang_name, -L lang_name`

Return client error messages in the given language. `lang_name` can be given as the language name or as the full pathname to the directory where the language files are installed. See [Section 5.10.2, “Setting the Error Message Language”](#).

- `--log[=file_name], -l [file_name]`

Log connections and SQL statements received from clients to this file. See [Section 5.11.2, “The General Query Log”](#). If you omit the filename, MySQL uses `host_name.log` as the filename.

- `--log-bin[=base_name]`

Enable binary logging. The server logs all statements that change data to the binary log, which is used for backup and replication. See [Section 5.11.4, “The Binary Log”](#).

The option value, if given, is the basename for the log sequence. The server creates binary log files in sequence by adding a numeric suffix to the basename. It is recommended that you specify a basename (see [Section A.8.4, “Open Issues in MySQL”](#), for the reason). Otherwise, MySQL uses `host_name-bin` as the basename.

- `--log-bin-index[=file_name]`

The index file for binary log filenames. See [Section 5.11.4, “The Binary Log”](#). If you omit the filename, and if you didn't specify one with `--log-bin`, MySQL uses `host_name-bin.index` as the filename.

- `--log-error[=file_name]`

Log errors and startup messages to this file. See [Section 5.11.1, “The Error Log”](#). If you omit the filename, MySQL uses `host_name.err`. If the filename has no extension, the server adds an extension of `.err`.

- `--log-isam[=file_name]`

Log all **ISAM/MyISAM** changes to this file (used only when debugging **ISAM/MyISAM**).

- `--log-long-format`

Log extra information to the update log, binary update log, and slow query log, if they have been activated. For example, the username and timestamp are logged for queries. Before MySQL 4.1, if you are using `--log-slow-queries` and `--log-long-format`, queries that are not using indexes also are logged to the slow query log. `--log-long-format` is deprecated as of MySQL version 4.1, when `--log-short-format` was introduced. (Long log format is the default setting since version 4.1.) Also note that starting with MySQL 4.1, the `--log-queries-not-using-indexes` option is available for the purpose of logging queries that do not use indexes to the slow query log.

- `--log-queries-not-using-indexes`

If you are using this option with `--log-slow-queries`, queries that do not use indexes also are logged to the slow query log. This option is available as of MySQL 4.1. See [Section 5.11.5, “The Slow Query Log”](#).

- `--log-short-format`

Log less information to the update log, binary update log, and slow query log, if they have been activated. For example, the username and timestamp are not logged for queries. This option was introduced in MySQL 4.1.

- `--log-slow-admin-statements`

Log slow administrative statements such as `OPTIMIZE TABLE`, `ANALYZE TABLE`, and `ALTER TABLE` to the slow query log.

This option was added in MySQL 4.1.13. (It is unnecessary in MySQL 4.0 because slow administrative statements are logged by default.)

- `--log-slow-queries[=file_name]`

Log all queries that have taken more than `long_query_time` seconds to execute to this file. See [Section 5.11.5, “The Slow Query Log”](#). Note that the default for the amount of information logged has changed in MySQL 4.1. See the `--log-long-format` and `--log-short-format` options for details.

- `--log-update[=file_name]`

Log updates to `fileN` where `N` is a unique number if not given. See [Section 5.11.3, “The Update Log”](#). The update log is now deprecated; you should use the binary log instead (`--log-bin`). See [Section 5.11.4, “The Binary Log”](#).

- `--log-warnings[=level], -W [level]`

Print out warnings such as `Aborted connection...` to the error log. Enabling this option is recommended, for example, if you use replication (you get more information about what is happening, such as messages about network failures and reconnections). This option is enabled by default as of MySQL 4.0.19 and 4.1.2; to disable it, use `--log-warnings=0`. As of MySQL 4.0.21 and 4.1.3, a *level* argument can be given. If omitted, the default *level* is 1. Aborted connections are not logged to the error log unless the value is greater than 1. See [Section A.2.10, “Communication Errors and Aborted Connections”](#).

Before MySQL 4.0.21 and 4.1.3, this is a boolean option, not an integer-valued option. Before 4.0, this option was named `--warnings`.

- `--low-priority-updates`

Give table-modifying operations (`INSERT`, `REPLACE`, `DELETE`, `UPDATE`) lower priority than selects. This can also be done via `{INSERT | REPLACE | DELETE | UPDATE} LOW_PRIORITY ...` to lower the priority of only one query, or by `SET LOW_PRIORITY_UPDATES=1` to change the priority in one thread. See [Section 7.3.2, “Table Locking Issues”](#).

- `--memlock`

Lock the `mysqld` process in memory. This works on systems such as Solaris that support the `mlockall()` system call. This might help if you have a problem where the operating system is causing `mysqld` to swap on disk. Note that use of this option requires that you run the server as `root`, which is normally not a good idea for security reasons. See [Section 5.6.5, “How to Run MySQL as a Normal User”](#).

- `--myisam-recover[=option[,option]...]`

Set the MyISAM storage engine recovery mode. The option value is any combination of the values of `DEFAULT`, `BACKUP`, `FORCE`, or `QUICK`. If you specify multiple values, separate them by commas. You can also use a value of `" "` to disable this option. If this option is used, each time `mysqld` opens a MyISAM table, it checks whether the table is marked as crashed or wasn't closed properly. (The last option works only if you are running with external locking disabled.) If this is the case, `mysqld` runs a check on the table. If the table was corrupted, `mysqld` attempts to repair it.

The following options affect how the repair works:

Option	Description
<code>DEFAULT</code>	The same as not giving any option to <code>--myisam-recover</code> .
<code>BACKUP</code>	If the data file was changed during recovery, save a backup of the <code>tbl_name.MYD</code> file as <code>tbl_name-datetime.BAK</code> .
<code>FORCE</code>	Run recovery even if we would lose more than one row from the <code>.MYD</code> file.
<code>QUICK</code>	do not check the rows in the table if there are not any delete blocks.

Before the server automatically repairs a table, it writes a note about the repair to the error log. If you want to be able to recover from most problems without user intervention, you should use the options `BACKUP`, `FORCE`. This forces a repair of a table even if some rows would be deleted, but it keeps the old data file as a backup so that you can later examine what happened.

This option is available as of MySQL 3.23.25.

- `--ndb-connectstring=connect_string`

When using the `NDB` storage engine, it is possible to point out the management server that distributes the cluster configuration by setting the connect string option. See [Section 15.4.4.2, “The Cluster connectstring”](#), for syntax.

- `--ndbcluster`

If the binary includes support for the `NDB Cluster` storage engine (from version 4.1.3, the MySQL-Max binaries are built with `NDB Cluster` enabled), this option enables the engine, which is disabled by default. Using the `NDB Cluster` storage engine is necessary for using MySQL Cluster. See [Chapter 15, \*MySQL Cluster\*](#).

- `--new`

The `--new` option can be used to make the server behave as 4.1 in certain respects, easing a 4.0 to 4.1 upgrade:

- Hexadecimal strings such as `0xFF` are treated as strings by default rather than as numbers. (Works in 4.0.12 and up.)
- `TIMESTAMP` is returned as a string with the format `'YYYY-MM-DD HH:MM:SS'`. (Works in 4.0.13 and up.) See [Chapter 11, \*Data Types\*](#).

This option can be used to help you see how your applications behave in MySQL 4.1, without actually upgrading to 4.1.

- `--old-passwords`

Force the server to generate short (pre-4.1) password hashes for new passwords. This is useful for compatibility when the server must support older client programs. See [Section 5.7.9, “Password Hashing as of MySQL 4.1”](#).

- `--old-protocol, -o`

Use the 3.20 protocol for compatibility with some very old clients.

- `--one-thread`

Only use one thread (for debugging under Linux). This option is available only if the server is built with debugging enabled. See [Section E.1, “Debugging a MySQL Server”](#).

- `--open-files-limit=count`

Change the number of file descriptors available to `mysqld`. If this option is not set or is set to 0, `mysqld` uses the value to reserve file descriptors with `setrlimit()`. If the value is 0, `mysqld` reserves `max_connections×5` or `max_connections + table_open_cache×2` files (whichever is larger). You should try increasing this value if `mysqld` gives you the error `Too many open files`.

- `--pid-file=path`

The pathname of the process ID file. This file is used by other programs such as `mysqld_safe` to determine the server's process ID.

- `--port=port_num, -P port_num`

The port number to use when listening for TCP/IP connections. The port number must be 1024 or higher unless the server is started by the `root` system user.

- `--safe-mode`

Skip some optimization stages.

- `--safe-show-database`

With this option, the `SHOW DATABASES` statement displays only the names of those databases for which the user has some kind of privilege. As of MySQL 4.0.2, this option is deprecated and does not do anything (it is enabled by default), because there is a `SHOW DATABASES` privilege that can be used to control access to database names on a per-account basis. See [Section 5.7.3, “Privileges Provided by MySQL”](#).

- `--safe-user-create`

If this option is enabled, a user cannot create new MySQL users by using the `GRANT` statement, if the user doesn't have the `INSERT` privilege for the `mysql.user` table or any column in the table.

- `--secure-auth`

Disallow authentication by clients that attempt to use accounts that have old (pre-4.1) passwords. This option is available as of MySQL 4.1.1.

- `--shared-memory`

Enable shared-memory connections by local clients. This option is available only on Windows. It was added in MySQL 4.1.0.

- `--shared-memory-base-name=name`

The name of shared memory to use for shared-memory connections. This option is available only on Windows. The default name is `MYSQL`. The name is case sensitive. This option was added in MySQL 4.1.0.

- `--skip-bdb`

Disable the `BDB` storage engine. This saves memory and might speed up some operations. Do not use this option if you require `BDB` tables.

- `--skip-concurrent-insert`

Turn off the ability to select and insert at the same time on `MyISAM` tables. (This is to be used only if you think you have found a bug in this feature.) See [Section 7.3.3, “Concurrent Inserts”](#).

- `--skip-delay-key-write`

Ignore the `DELAY_KEY_WRITE` option for all tables. As of MySQL 4.0.3, you should use `--delay-key-write=OFF` instead. See [Section 7.5.2, “Tuning Server Parameters”](#).

- `--skip-external-locking`

Do not use external locking (system locking). With external locking disabled, you must shut down the server to use `myisamchk` or `isamchk`. See [Section 1.4.3, “MySQL Stability”](#). As of MySQL 3.23, you can use the `CHECK TABLE` and `REPAIR TABLE` statements to check and repair `MyISAM` tables. This option previously was named `--skip-locking`.

External locking has been disabled by default since MySQL 4.0.

- `--skip-grant-tables`

This option causes the server not to use the privilege system at all, which gives anyone with access to the server *unrestricted access to all databases*. You can cause a running server to start using the grant tables again by executing `mysqladmin flush-privileges` or `mysqladmin reload` command from a system shell, or by issuing a MySQL `FLUSH PRIVILEGES` statement after connecting to the server. This option also suppresses loading of user-defined functions (UDFs).

- `--skip-host-cache`

Do not use the internal hostname cache for faster name-to-IP resolution. Instead, query the DNS server every time a client connects. See [Section 7.5.5, “How MySQL Uses DNS”](#).

- `--skip-innodb`

Disable the `InnoDB` storage engine. This saves memory and disk space and might speed up some operations. Do not use this option if you require `InnoDB` tables.

- `--skip-isam`

Disable the `ISAM` storage engine. As of MySQL 4.1, `ISAM` is disabled by default, so this option applies only if the server was configured with support for `ISAM`. This option was added in MySQL 4.1.1.

- `--skip-merge`

Disable the `MERGE` storage engine. This option was added in MySQL 4.1.21. It can be used if the following behavior is undesirable: If a user has access to `MyISAM` table `t`, that user can create a `MERGE` table `m` that accesses `t`. However, if the user's privileges on `t` are subsequently revoked, the user can continue to access `t` by doing so through `m`.

- `--skip-name-resolve`

Do not resolve hostnames when checking client connections. Use only IP numbers. If you use this option, all `Host` column values in the grant tables must be IP numbers or `localhost`. See [Section 7.5.5, “How MySQL Uses DNS”](#).

- `--skip-ndbcluster`

Disable the `NDB Cluster` storage engine. This is the default for binaries that were built with `NDB Cluster` storage engine support; the server allocates memory and other resources for this storage engine only if the `--ndbcluster` option is given explicitly. See [Section 15.4.3, “Quick Test Setup of MySQL Cluster”](#), for an example of usage.

- `--skip-networking`

Do not listen for TCP/IP connections at all. All interaction with `mysqld` must be made via named pipes or shared memory (on Windows) or Unix socket files (on Unix). This option is highly recommended for systems where only local clients are allowed. See [Section 7.5.5, “How MySQL Uses DNS”](#).

- `--skip-new`

do not use new, possibly wrong routines.

- `--skip-symlink`

This is the old form of `--skip-symbolic-links`, for use before MySQL 4.0.13.

- `--ssl*`

Options that begin with `--ssl` specify whether to allow clients to connect via SSL and indicate where to find SSL keys and certificates. See [Section 5.8.7.3, “SSL Command Options”](#).

- `--standalone`

Available on Windows NT-based systems only; instructs the MySQL server not to run as a service.

- `--symbolic-links, --skip-symbolic-links`

Enable or disable symbolic link support. This option has different effects on Windows and Unix:

- On Windows, enabling symbolic links allows you to establish a symbolic link to a database directory by creating a `db_name.sym` file that contains the path to the real directory. See [Section 7.6.1.3, “Using Symbolic Links for Databases on Windows”](#).
- On Unix, enabling symbolic links means that you can link a [MyISAM](#) index file or data file to another directory with the `INDEX DIRECTORY` or `DATA DIRECTORY` options of the `CREATE TABLE` statement. If you delete or rename the table, the files that its symbolic links point to also are deleted or renamed. See [Section 7.6.1.2, “Using Symbolic Links for Tables on Unix”](#).

This option was added in MySQL 4.0.13.

- `--skip-safemalloc`

If MySQL is configured with `--with-debug=full`, all MySQL programs check for memory overruns during each memory allocation and memory freeing operation. This checking is very slow, so for the server you can avoid it when you do not need it by using the `--skip-safemalloc` option.

- `--skip-show-database`

With this option, the `SHOW DATABASES` statement is allowed only to users who have the `SHOW DATABASES` privilege, and the statement displays all database names. Without this option, `SHOW DATABASES` is allowed to all users, but displays each database name only if the user has the `SHOW DATABASES` privilege or some privilege for the database. Note that *any* global privilege is considered a privilege for the database.

- `--skip-stack-trace`

do not write stack traces. This option is useful when you are running `mysqld` under a debugger. On some systems, you also must use this option to get a core file. See [Section E.1, “Debugging a MySQL Server”](#).

- `--skip-thread-priority`

Disable using thread priorities for faster response time.

- `--socket=path`

On Unix, this option specifies the Unix socket file to use when listening for local connections. The default value is `/tmp/mysql.sock`. On Windows, the option specifies the pipe name to use when listening for local connections that use a named pipe. The default value is `MySQL` (not case sensitive).

- `--sql-mode=value[,value[,value...]]`

Set the SQL mode. See [Section 5.2.5, “SQL Modes”](#). This option was added in 3.23.41.



- `--temp-pool`

This option causes most temporary files created by the server to use a small set of names, rather than a unique name for each new file. This works around a problem in the Linux kernel dealing with creating many new files with different names. With the old behavior, Linux seems to “leak” memory, because it is being allocated to the directory entry cache rather than to the disk cache.

- `--transaction-isolation=level`

Sets the default transaction isolation level. The `level` value can be `READ-UNCOMMITTED`, `READ-COMMITTED`, `REPEATABLE-READ`, or `SERIALIZABLE`. See [Section 13.4.6, “SET TRANSACTION Syntax”](#).

- `--tmpdir=path, -t path`

The path of the directory to use for creating temporary files. It might be useful if your default `/tmp` directory resides on a partition that is too small to hold temporary tables. Starting from MySQL 4.1.0, this option accepts several paths that are used in round-robin fashion. Paths should be separated by colon characters (‘:’) on Unix and semicolon characters (‘;’) on Windows, NetWare, and OS/2. If the MySQL server is acting as a replication slave, you should not set `--tmpdir` to point to a directory on a memory-based filesystem or to a directory that is cleared when the server host restarts. For more information about the storage location of temporary files, see [Section A.4.4, “Where MySQL Stores Temporary Files”](#). A replication slave needs some of its temporary files to survive a machine restart so that it can replicate temporary tables or `LOAD DATA INFILE` operations. If files in the temporary file directory are lost when the server restarts, replication fails.

- `--user={user_name|user_id}, -u {user_name|user_id}`

Run the `mysqld` server as the user having the name `user_name` or the numeric user ID `user_id`. (“User” in this context refers to a system login account, not a MySQL user listed in the grant tables.)

This option is *mandatory* when starting `mysqld` as `root`. The server changes its user ID during its startup sequence, causing it to run as that particular user rather than as `root`. See [Section 5.6.1, “General Security Guidelines”](#).

Starting from MySQL 3.23.56 and 4.0.12: To avoid a possible security hole where a user adds a `--user=root` option to a `my.cnf` file (thus causing the server to run as `root`), `mysqld` uses only the first `--user` option specified and produces a warning if there are multiple `--user` options. Options in `/etc/my.cnf` and `$MYSQL_HOME/my.cnf` are processed before command-line options, so it is recommended that you put a `--user` option in `/etc/my.cnf` and specify a value other than `root`. The option in `/etc/my.cnf` is found before any other `--user` options, which ensures that the server runs as a user other than `root`, and that a warning results if any other `--user` option is found.

- `--version, -V`

Display version information and exit.

As of MySQL 4.0, you can assign a value to a server system variable by using an option of the form `--var_name=value`. For example, `--key_buffer_size=32M` sets the `key_buffer_size` variable to a value of 32MB.

Note that when you assign a value to a variable, MySQL might automatically correct the value to stay within a given range, or adjust the value to the closest allowable value if only certain values are allowed.

If you want to restrict the maximum value to which a variable can be set at runtime with `SET`, you can define this by using the `--maximum-var_name=value` command-line option.

It is also possible to set variables by using `--set-variable=var_name=value` or `--var_name=value` syntax. *This syntax is deprecated as of MySQL 4.0.*

You can change the values of most system variables for a running server with the `SET` statement. See [Section 13.5.3, “SET Syntax”](#).

[Section 5.2.2, “System Variables”](#), provides a full description for all variables, and additional information for setting them at server startup and runtime. [Section 7.5.2, “Tuning Server Parameters”](#), includes information on optimizing the server by tuning system variables.

## 5.2.2. System Variables

The `mysql` server maintains many system variables that indicate how it is configured. Each system variable has a default value. System variables can be set at server startup using options on the command line or in an option file. As of MySQL 4.0.3, most of them can be changed dynamically while the server is running by means of the `SET` statement, which enables you to modify operation of the server without having to stop and restart it. You can refer to system variable values in expressions.

There are several ways to see the names and values of system variables:

- To see the values that a server will use based on its compiled-in defaults and any option files that it reads, use this command (omit `--verbose` before MySQL 4.1.1):

```
mysqld --verbose --help
```

- To see the values that a server will use based on its compiled-in defaults, ignoring the settings in any option files, use this command (omit `--verbose` before MySQL 4.1.1):

```
mysqld --no-defaults --verbose --help
```

- To see the current values used by a running server, use the `SHOW VARIABLES` statement.

This section provides a description of each system variable. Variables with no version indicated have been present since at least MySQL 3.22.

For additional system variable information, see these sections:

- [Section 5.2.3, “Using System Variables”](#), discusses the syntax for setting and displaying system variable values.
- [Section 5.2.3.2, “Dynamic System Variables”](#), lists the variables that can be set at runtime.
- Information on tuning system variables can be found in [Section 7.5.2, “Tuning Server Parameters”](#).
- [Section 14.2.5, “InnoDB Startup Options and System Variables”](#), lists `InnoDB` system variables.

*Note:* Some of the following variable descriptions refer to “enabling” or “disabling” a variable. These variables can be enabled with the `SET` statement by setting them to `ON` or `1`, or disabled by setting them to `OFF` or `0`. However, to set such a variable on the command line or in an option file, you must set it to `1` or `0`; setting it to `ON` or `OFF` will not work. For example, on the command line, `-delay_key_write=1` works but `--delay_key_write=ON` does not.

Values for buffer sizes, lengths, and stack sizes are given in bytes unless otherwise specified.

- `ansi_mode`

This is `ON` if `mysqld` was started with `--ansi`. See [Section 1.9.3, “Running MySQL in ANSI Mode”](#). This variable was added in MySQL 3.23.6 and removed in 3.23.41. See the description for `sql_mode`.

- `back_log`

The number of outstanding connection requests MySQL can have. This comes into play when the main MySQL thread gets very many connection requests in a very short time. It then takes some time (although very little) for the main thread to check the connection and start a new thread. The `back_log` value indicates how many requests can be stacked during this short time before MySQL momentarily stops answering new requests. You need to increase this only if you expect a large number of connections in a short period of time.

In other words, this value is the size of the listen queue for incoming TCP/IP connections. Your operating system has its own limit on the size of this queue. The manual page for the Unix `listen()` system call should have more details. Check your OS documentation for the maximum value for this variable. `back_log` cannot be set higher than your operating system limit.

- `basedir`

The MySQL installation base directory. This variable can be set with the `--basedir` option.

- `bdb_cache_size`

The size of the buffer that is allocated for caching indexes and rows for `BDB` tables. If you do not use `BDB` tables, you should start `mysqld` with `--skip-bdb` to not allocate memory for this cache. This variable was added in MySQL 3.23.14.

- `bdb_home`

The base directory for `BDB` tables. This should be assigned the same value as the `datadir` variable. This variable was added in MySQL 3.23.14.

- `bdb_log_buffer_size`

The size of the buffer that is allocated for caching indexes and rows for `BDB` tables. If you do not use `BDB` tables, you should set this to 0 or start `mysqld` with `--skip-bdb` in order not to allocate memory for this cache. This variable was added in MySQL 3.23.31.

- `bdb_logdir`

The directory where the `BDB` storage engine writes its log files. This variable can be set with the `--bdb-logdir` option. This variable was added in MySQL 3.23.14.

- `bdb_max_lock`

The maximum number of locks that can be active for a `BDB` table (10,000 by default). You should increase this value if errors such as the following occur when you perform long transactions or when `mysqld` has to examine many rows to calculate a query:

```
bdb: Lock table is out of available locks
Got error 12 from ...
```

This variable was added in MySQL 3.23.29.

- `bdb_shared_data`

This is `ON` if you are using `--bdb-shared-data` to start Berkeley DB in multi-process mode. (Do not use `DB_PRIVATE` when initializing Berkeley DB.) This variable was added in MySQL 3.23.29.

- `bdb_tmpdir`

The `BDB` temporary file directory. This variable was added in MySQL 3.23.14.

- `bdb_version`

See the description for `version_bdb`.

- `binlog_cache_size`

The size of the cache to hold the SQL statements for the binary log during a transaction. A binary log cache is allocated for each client if the server supports any transactional storage engines and, starting from MySQL 4.1.2, if the server has the binary log enabled (`--log-bin` option). If you often use large, multiple-statement transactions, you can increase this cache size to get more performance. The `Binlog_cache_use` and `Binlog_cache_disk_use` status variables can be useful for tuning the size of this variable. This variable was added in MySQL 3.23.29. See [Section 5.11.4, “The Binary Log”](#).

- `bulk_insert_buffer_size`

`MyISAM` uses a special tree-like cache to make bulk inserts faster for `INSERT ... SELECT`, `INSERT ... VALUES (...), (...), ...`, and `LOAD DATA INFILE` when adding data to non-empty tables. This variable limits the size of the cache tree in bytes per thread. Setting it to 0 disables this optimization. The default value is 8MB. This variable was added in MySQL 4.0.3. This variable previously was named `myisam_bulk_insert_tree_size`.

- `character_set`

The default character set. This variable was added in MySQL 3.23.3, then removed in MySQL 4.1.1 and replaced by the various `character_set_xxx` variables.

- `character_set_client`

The character set for statements that arrive from the client. This variable was added in MySQL 4.1.1.

- `character_set_connection`

The character set used for literals that do not have a character set introducer and for number-to-string conversion. This variable was added in MySQL 4.1.1.

- `character_set_database`

The character set used by the default database. The server sets this variable whenever the default database changes. If there is no default database, the variable has the same value as `character_set_server`. This variable was added in MySQL 4.1.1.

- `character_set_results`

The character set used for returning query results to the client. This variable was added in MySQL 4.1.1.

- `character_set_server`

The server default character set. This variable was added in MySQL 4.1.1.

- `character_set_system`

The character set used by the server for storing identifiers. The value is always `utf8`. This variable was added in MySQL 4.1.1.
- `character_sets`

The supported character sets. This variable was added in MySQL 3.23.15 and removed in MySQL 4.1.1. (Use `SHOW CHARACTER SET` for a list of character sets.)
- `character_sets_dir`

The directory where character sets are installed. This variable was added in MySQL 4.1.2.
- `collation_connection`

The collation of the connection character set. This variable was added in MySQL 4.1.1.
- `collation_database`

The collation used by the default database. The server sets this variable whenever the default database changes. If there is no default database, the variable has the same value as `collation_server`. This variable was added in MySQL 4.1.1.
- `collation_server`

The server default collation. This variable was added in MySQL 4.1.1.
- `concurrent_insert`

If `ON` (the default), MySQL allows `INSERT` and `SELECT` statements to run concurrently for `MyISAM` tables that have no free blocks in the middle. You can turn this option off by starting `mysqld` with `--safe` or `--skip-new`. This variable was added in MySQL 3.23.7.

See also [Section 7.3.3, “Concurrent Inserts”](#).
- `connect_timeout`

The number of seconds that the `mysqld` server waits for a connect packet before responding with `Bad handshake`.
- `convert_character_set`

The current character set mapping that was set by `SET CHARACTER SET`. This variable was removed in MySQL 4.1.
- `datadir`

The MySQL data directory. This variable can be set with the `--datadir` option.
- `date_format`

This variable is not implemented.
- `datetime_format`

This variable is not implemented.
- `default_week_format`

The default mode value to use for the `WEEK()` function. See [Section 12.5, “Date and Time Functions”](#). This variable is available as of MySQL 4.0.14.

- `delay_key_write`

This option applies only to `MyISAM` tables. It can have one of the following values to affect handling of the `DELAY_KEY_WRITE` table option that can be used in `CREATE TABLE` statements.

Option	Description
OFF	<code>DELAY_KEY_WRITE</code> is ignored.
ON	MySQL honors any <code>DELAY_KEY_WRITE</code> option specified in <code>CREATE TABLE</code> statements. This is the default value.
ALL	All new opened tables are treated as if they were created with the <code>DELAY_KEY_WRITE</code> option enabled.

If `DELAY_KEY_WRITE` is enabled for a table, the key buffer is not flushed for the table on every index update, but only when the table is closed. This speeds up writes on keys a lot, but if you use this feature, you should add automatic checking of all `MyISAM` tables by starting the server with the `--myisam-recover` option (for example, `--myisam-recover=BACKUP, FORCE`). See [Section 5.2.1, “Command Options”](#), and [Section 14.1.1, “MyISAM Startup Options”](#).

Note that enabling external locking with `--external-locking` offers no protection against index corruption for tables that use delayed key writes.

This variable was added in MySQL 3.23.8.

- `delayed_insert_limit`

After inserting `delayed_insert_limit` delayed rows, the `INSERT DELAYED` handler thread checks whether there are any `SELECT` statements pending. If so, it allows them to execute before continuing to insert delayed rows.

- `delayed_insert_timeout`

How many seconds an `INSERT DELAYED` handler thread should wait for `INSERT` statements before terminating.

- `delayed_queue_size`

This is a per-table limit on the number of rows to queue when handling `INSERT DELAYED` statements. If the queue becomes full, any client that issues an `INSERT DELAYED` statement waits until there is room in the queue again.

- `expire_logs_days`

The number of days for automatic binary log removal. The default is 0, which means “no automatic removal.” Possible removals happen at startup and at binary log rotation. This variable was added in MySQL 4.1.0.

- `flush`

If `ON`, the server flushes (synchronizes) all changes to disk after each SQL statement. Normally, MySQL does a write of all changes to disk only after each SQL statement and lets the operating system handle the synchronizing to disk. See [Section A.4.2, “What to Do If MySQL Keeps Crashing”](#). This variable is set to `ON` if you start `mysqld` with the `--flush` option. This variable was added in

MySQL 3.22.9.

- `flush_time`

If this is set to a non-zero value, all tables are closed every `flush_time` seconds to free up resources and synchronize unflushed data to disk. We recommend that this option be used only on Windows 9x or Me, or on systems with minimal resources. This variable was added in MySQL 3.22.18.

- `ft_boolean_syntax`

The list of operators supported by boolean full-text searches performed using `IN BOOLEAN MODE`. See [Section 12.7.1, “Boolean Full-Text Searches”](#). This variable was added as a read-only variable in MySQL 4.0.1. It can be modified as of MySQL 4.1.2.

The default variable value is `'+ -><()~*:""&|'`. The rules for changing the value are as follows:

- Operator function is determined by position within the string.
- The replacement value must be 14 characters.
- Each character must be an ASCII non-alphanumeric character.
- Either the first or second character must be a space.
- No duplicates are allowed except the phrase quoting operators in positions 11 and 12. These two characters are not required to be the same, but they are the only two that may be.
- Positions 10, 13, and 14 (which by default are set to `':'`, `'&'`, and `'|'`) are reserved for future extensions.

- `ft_max_word_len`

The maximum length of the word to be included in a `FULLTEXT` index. This variable was added in MySQL 4.0.0.

**Note:** `FULLTEXT` indexes must be rebuilt after changing this variable. Use `REPAIR TABLE tbl_name QUICK`.

- `ft_min_word_len`

The minimum length of the word to be included in a `FULLTEXT` index. This variable was added in MySQL 4.0.0.

**Note:** `FULLTEXT` indexes must be rebuilt after changing this variable. Use `REPAIR TABLE tbl_name QUICK`.

- `ft_query_expansion_limit`

The number of top matches to use for full-text searches performed using `WITH QUERY EXPANSION`. This variable was added in MySQL 4.1.1.

- `ft_stopword_file`

The file from which to read the list of stopwords for full-text searches. All the words from the file are used; comments are *not* honored. By default, a built-in list of stopwords is used (as defined in the `mysam/ft_static.c` file). Setting this variable to the empty string (`' '`) disables stopword filtering. This variable was added in MySQL 4.0.10.

**Note:** `FULLTEXT` indexes must be rebuilt after changing this variable or the contents of the stop-word file. Use `REPAIR TABLE tbl_name QUICK`.

- `group_concat_max_len`

The maximum allowed result length for the `GROUP_CONCAT()` function. The default is 1024. This variable was added in MySQL 4.1.0.

- `have_archive`

`YES` if `mysqld` supports `ARCHIVE` tables, `NO` if not. This variable was added in MySQL 4.1.3.

- `have_bdb`

`YES` if `mysqld` supports `BDB` tables. `DISABLED` if `--skip-bdb` is used. This variable was added in MySQL 3.23.30.

- `have_blackhole_engine`

`YES` if `mysqld` supports `BLACKHOLE` tables, `NO` if not. This variable was added in MySQL 4.1.11.

- `have_compress`

`YES` if the `zlib` compression library is available to the server, `NO` if not. If not, the `COMPRESS()` and `UNCOMPRESS()` functions cannot be used. This variable was added in MySQL 4.1.1.

- `have_crypt`

`YES` if the `crypt()` system call is available to the server, `NO` if not. If not, the `ENCRYPT()` function cannot be used. This variable was added in MySQL 4.0.10.

- `have_csv`

`YES` if `mysqld` supports `ARCHIVE` tables, `NO` if not. This variable was added in MySQL 4.1.4.

- `have_example_engine`

`YES` if `mysqld` supports `EXAMPLE` tables, `NO` if not. This variable was added in MySQL 4.1.4.

- `have_geometry`

`YES` if the server supports spatial data types, `NO` if not. This variable was added in MySQL 4.1.3.

- `have_innodb`

`YES` if `mysqld` supports `InnoDB` tables. `DISABLED` if `--skip-innodb` is used. This variable was added in MySQL 3.23.37.

- `have_isam`

`YES` if `mysqld` supports `ISAM` tables. `DISABLED` if `--skip-isam` is used. This variable was added in MySQL 3.23.30.

- `have_ndbcluster`

`YES` if `mysqld` supports `NDB Cluster` tables. `DISABLED` if `--skip-ndbcluster` is used. This variable was added in MySQL 4.1.2.

- `have_openssl`



`YES` if `mysqld` supports SSL (encryption) connections, `NO` if not. This variable was added in MySQL 3.23.43.

- `have_query_cache`

`YES` if `mysqld` supports the query cache, `NO` if not. This variable was added in MySQL 4.0.2.

- `have_raid`

`YES` if `mysqld` supports the `RAID` option, `NO` if not. This variable was added in MySQL 3.23.30.

- `have_rtree_keys`

`YES` if `RTREE` indexes are available, `NO` if not. (These are used for spatial indexes in `MyISAM` tables.) This variable was added in MySQL 4.1.3.

- `have_symlink`

`YES` if symbolic link support is enabled, `NO` if not. This is required on Unix for support of the `DATA DIRECTORY` and `INDEX DIRECTORY` table options, and on Windows for support of data directory symlinks.

This variable was added in MySQL 4.0.0.

- `init_connect`

A string to be executed by the server for each client that connects. The string consists of one or more SQL statements. To specify multiple statements, separate them by semicolon characters. For example, each client begins by default with autocommit mode enabled. There is no global system variable to specify that autocommit should be disabled by default, but `init_connect` can be used to achieve the same effect:

```
SET GLOBAL init_connect='SET AUTOCOMMIT=0';
```

This variable can also be set on the command line or in an option file. To set the variable as just shown using an option file, include these lines:

```
[mysqld]
init_connect='SET AUTOCOMMIT=0'
```

Note that the content of `init_connect` is not executed for users that have the `SUPER` privilege. This is done so that an erroneous value for `init_connect` does not prevent all clients from connecting. For example, the value might contain a statement that has a syntax error, thus causing client connections to fail. Not executing `init_connect` for users that have the `SUPER` privilege enables them to open a connection and fix the `init_connect` value.

This variable was added in MySQL 4.1.2.

- `init_file`

The name of the file specified with the `--init-file` option when you start the server. This should be a file containing SQL statements that you want the server to execute when it starts. Each statement must be on a single line and should not include comments. This variable was added in MySQL 3.23.2.

- `init_slave`

This variable is similar to `init_connect`, but is a string to be executed by a slave server each

time the SQL thread starts. The format of the string is the same as for the `init_connect` variable. This variable was added in MySQL 4.1.2.

- `innodb_xxx`

InnoDB system variables are listed in [Section 14.2.5, “InnoDB Startup Options and System Variables”](#).

- `interactive_timeout`

The number of seconds the server waits for activity on an interactive connection before closing it. An interactive client is defined as a client that uses the `CLIENT_INTERACTIVE` option to `mysql_real_connect()`. See also `wait_timeout`.

- `join_buffer_size`

The size of the buffer that is used for joins that do not use indexes and thus perform full table scans. Normally, the best way to get fast joins is to add indexes. Increase the value of `join_buffer_size` to get a faster full join when adding indexes is not possible. One join buffer is allocated for each full join between two tables. For a complex join between several tables for which indexes are not used, multiple join buffers might be necessary.

- `key_buffer_size`

Index blocks for `MyISAM` and `ISAM` tables are buffered and are shared by all threads. `key_buffer_size` is the size of the buffer used for index blocks. The key buffer is also known as the key cache.

The maximum allowable setting for `key_buffer_size` is 4GB. The effective maximum size might be less, depending on your available physical RAM and per-process RAM limits imposed by your operating system or hardware platform.

Increase the value to get better index handling (for all reads and multiple writes) to as much as you can afford. Using a value that is 25% of total memory on a machine that mainly runs MySQL is quite common. However, if you make the value too large (for example, more than 50% of your total memory) your system might start to page and become extremely slow. MySQL relies on the operating system to perform filesystem caching for data reads, so you must leave some room for the filesystem cache. Consider also the memory requirements of other storage engines.

For even more speed when writing many rows at the same time, use `LOCK TABLES`. See [Section 7.2.13, “Speed of INSERT Statements”](#).

You can check the performance of the key buffer by issuing a `SHOW STATUS` statement and examining the `Key_read_requests`, `Key_reads`, `Key_write_requests`, and `Key_writes` status variables. (See [Section 13.5.4, “SHOW Syntax”](#).) The `Key_reads/Key_read_requests` ratio should normally be less than 0.01. The `Key_writes/Key_write_requests` ratio is usually near 1 if you are using mostly updates and deletes, but might be much smaller if you tend to do updates that affect many rows at the same time or if you are using the `DELAY_KEY_WRITE` table option.

The fraction of the key buffer in use can be determined using `key_buffer_size` in conjunction with the `Key_blocks_unused` status variable and the buffer block size. From MySQL 4.1.1 on, the buffer block size is available from the `key_cache_block_size` server variable. The fraction of the buffer in use is:

```
1 - ((Key_blocks_unused × key_cache_block_size) / key_buffer_size)
```

This value is an approximation because some space in the key buffer may be allocated internally for

administrative structures.

Before MySQL 4.1.1, key cache blocks are 1024 bytes, and before MySQL 4.1.2, `Key_blocks_unused` is unavailable. The `Key_blocks_used` variable can be used as follows to determine the fraction of the key buffer in use:

```
(Key_blocks_used × 1024) / key_buffer_size
```

However, `Key_blocks_used` indicates the maximum number of blocks that have ever been in use at once, so this formula does not necessarily represent the current fraction of the buffer that is in use.

As of MySQL 4.1, it is possible to create multiple `MyISAM` key caches. The size limit of 4GB applies to each cache individually, not as a group. See [Section 7.4.6, “The MyISAM Key Cache”](#).

- `key_cache_age_threshold`

This value controls the demotion of buffers from the hot sub-chain of a key cache to the warm sub-chain. Lower values cause demotion to happen more quickly. The minimum value is 100. The default value is 300. This variable was added in MySQL 4.1.1. See [Section 7.4.6, “The MyISAM Key Cache”](#).

- `key_cache_block_size`

The size in bytes of blocks in the key cache. The default value is 1024. This variable was added in MySQL 4.1.1. See [Section 7.4.6, “The MyISAM Key Cache”](#).

- `key_cache_division_limit`

The division point between the hot and warm sub-chains of the key cache buffer chain. The value is the percentage of the buffer chain to use for the warm sub-chain. Allowable values range from 1 to 100. The default value is 100. This variable was added in MySQL 4.1.1. See [Section 7.4.6, “The MyISAM Key Cache”](#).

- `language`

The language used for error messages.

- `large_file_support`

Whether `mysqld` was compiled with options for large file support. This variable was added in MySQL 3.23.28.

- `large_pages`

Whether large page support is enabled. This variable was added in MySQL 5.0.3.

- `license`

The type of license the server has. This variable was added in MySQL 4.0.19.

- `local_infile`

Whether `LOCAL` is supported for `LOAD DATA INFILE` statements. See [Section 5.6.4, “Security Issues with LOAD DATA LOCAL”](#). This variable was added in MySQL 4.0.3.

- `locked_in_memory`

Whether `mysqld` was locked in memory with `--memlock`. This variable was added in MySQL 3.23.25.

- `log`

Whether logging of all statements to the general query log is enabled. See [Section 5.11.2, “The General Query Log”](#).

- `log_bin`

Whether the binary log is enabled. This variable was added in MySQL 3.23.14. See [Section 5.11.4, “The Binary Log”](#).

- `log_error`

The location of the error log. This variable was added in MySQL 4.0.10.

- `log_slave_updates`

Whether updates received by a slave server from a master server should be logged to the slave's own binary log. Binary logging must be enabled on the slave for this variable to have any effect. This variable was added in MySQL 3.23.17. See [Section 6.8, “Replication Startup Options”](#).

- `log_slow_queries`

Whether slow queries should be logged. “Slow” is determined by the value of the `long_query_time` variable. This variable was added in MySQL 4.0.2. See [Section 5.11.5, “The Slow Query Log”](#).

- `log_update`

Whether the update log is enabled. This variable was added in MySQL 3.22.18. Note that the binary log is preferable to the update log, which is unavailable as of MySQL 5.0. See [Section 5.11.3, “The Update Log”](#).

- `log_warnings`

Whether to produce additional warning messages. This variable was added in MySQL 4.0.3. It is enabled by default as of MySQL 4.0.19 and 4.1.2. As of MySQL 4.0.21 and 4.1.3, the variable can take values greater than 1 and aborted connections are not logged to the error log unless the value is greater than 1.

- `long_query_time`

If a query takes longer than this many seconds, the server increments the `Slow_queries` status variable. If you are using the `--log-slow-queries` option, the query is logged to the slow query log file. This value is measured in real time, not CPU time, so a query that is under the threshold on a lightly loaded system might be above the threshold on a heavily loaded one. The minimum value is 1. The default is 10. See [Section 5.11.5, “The Slow Query Log”](#).

- `low_priority_updates`

If set to 1, all `INSERT`, `UPDATE`, `DELETE`, and `LOCK TABLE WRITE` statements wait until there is no pending `SELECT` or `LOCK TABLE READ` on the affected table. This variable previously was named `sql_low_priority_updates`. It was added in MySQL 3.22.5.

- `lower_case_file_system`

This variable describes the case sensitivity of filenames on the filesystem where the data directory is located. `OFF` means filenames are case sensitive, `ON` means they are not case sensitive. This variable was added in MySQL 4.0.19.

- `lower_case_table_names`

If set to 1 table names are stored in lowercase on disk and table name comparisons are not case sensitive. This variable was added in MySQL 3.23.6. If set to 2 (new in 4.0.18), table names are stored as given but compared in lowercase. From MySQL 4.0.2, this option also applies to database names. From 4.1.1, it also applies to table aliases. See [Section 9.2.2, “Identifier Case Sensitivity”](#).

**Note:** If you are using `InnoDB` tables, you should set this variable to 1 on all platforms to force names to be converted to lowercase.

You should *not* set this variable to 0 if you are running MySQL on a system that does not have case-sensitive filenames (such as Windows or Mac OS X). *New in 4.0.18:* If this variable is not set at start-up and the filesystem on which the data directory is located does not have case-sensitive filenames, MySQL automatically sets `lower_case_table_names` to 2.

- `max_allowed_packet`

The maximum size of one packet or any generated/intermediate string.

The packet message buffer is initialized to `net_buffer_length` bytes, but can grow up to `max_allowed_packet` bytes when needed. This value by default is small, to catch large (possibly incorrect) packets.

You must increase this value if you are using large `BLOB` columns or long strings. It should be as big as the largest `BLOB` you want to use. The protocol limit for `max_allowed_packet` is 16MB before MySQL 4.0 and 1GB thereafter.

- `max_binlog_cache_size`

If a multiple-statement transaction requires more than this amount of memory, the server generates a `Multi-statement transaction required more than 'max_binlog_cache_size' bytes of storage` error. This variable was added in MySQL 3.23.29.

- `max_binlog_size`

If a write to the binary log causes the current log file size to exceed the value of this variable, the server rotates the binary logs (closes the current file and opens the next one). You cannot set this variable to more than 1GB or to less than 4096 bytes. (The minimum before MySQL 4.0.14 is 1024 bytes.) The default value is 1GB. This variable was added in MySQL 3.23.33.

A transaction is written in one chunk to the binary log, so it is never split between several binary logs. Therefore, if you have big transactions, you might see binary logs larger than `max_binlog_size`.

If `max_relay_log_size` is 0, the value of `max_binlog_size` applies to relay logs as well. `max_relay_log_size` was added in MySQL 4.0.14.

- `max_connect_errors`

If there are more than this number of interrupted connections from a host, that host is blocked from further connections. You can unblock blocked hosts with the `FLUSH HOSTS` statement.

- `max_connections`

The number of simultaneous client connections allowed. Increasing this value increases the number of file descriptors that `mysqld` requires. See [Section 7.4.8, “How MySQL Opens and Closes Tables”](#), for comments on file descriptor limits. See also [Section A.2.6, “Too many connections”](#).

- `max_delayed_threads`

Do not start more than this number of threads to handle `INSERT DELAYED` statements. If you try to insert data into a new table after all `INSERT DELAYED` threads are in use, the row is inserted as if the `DELAYED` attribute wasn't specified. If you set this to 0, MySQL never creates a thread to handle `DELAYED` rows; in effect, doing so disables `DELAYED` entirely. This variable was added in MySQL 3.23.0.

- `max_error_count`

The maximum number of error, warning, and note messages to be stored for display by the `SHOW ERRORS` or `SHOW WARNINGS` statements. This variable was added in MySQL 4.1.0.

- `max_heap_table_size`

This variable sets the maximum size to which `MEMORY (HEAP)` tables are allowed to grow. The value of the variable is used to calculate `MEMORY` table `MAX_ROWS` values. Setting this variable has no effect on any existing `MEMORY` table, unless the table is re-created with a statement such as `CREATE TABLE`, or altered with `ALTER TABLE` or `TRUNCATE TABLE`. This variable was added in MySQL 3.23.0.

- `max_insert_delayed_threads`

This variable is a synonym for `max_delayed_threads`. It was added in MySQL 4.0.19.

- `max_join_size`

Do not allow `SELECT` statements that probably need to examine more than `max_join_size` rows (for single-table statements) or row combinations (for multiple-table statements) or that are likely to do more than `max_join_size` disk seeks. By setting this value, you can catch `SELECT` statements where keys are not used properly and that would probably take a long time. Set it if your users tend to perform joins that lack a `WHERE` clause, that take a long time, or that return millions of rows.

Setting this variable to a value other than `DEFAULT` resets the value of `SQL_BIG_SELECTS` to 0. If you set the `SQL_BIG_SELECTS` value again, the `max_join_size` variable is ignored.

If a query result is in the query cache, no result size check is performed, because the result has previously been computed and it does not burden the server to send it to the client.

This variable previously was named `sql_max_join_size`.

- `max_length_for_sort_data`

The cutoff on the size of index values that determines which `filesort` algorithm to use. See [Section 7.2.9, “ORDER BY Optimization”](#). This variable was added in MySQL 4.1.1

- `max_prepared_stmt_count`

This variable limits the total number of prepared statements in the server. It can be used in environments where there is the potential for denial-of-service attacks based on running the server out of memory by preparing huge numbers of statements. The default value is 16,382. The allowable range of values is from 0 to 1 million. If the value is set lower than the current number of prepared statements, existing statements are not affected and can be used, but no new statements can be prepared

until the current number drops below the limit. This variable was added in MySQL 4.1.19.

- `max_relay_log_size`

If a write by a replication slave to its relay log causes the current log file size to exceed the value of this variable, the slave rotates the relay logs (closes the current file and opens the next one). If `max_relay_log_size` is 0, the server uses `max_binlog_size` for both the binary log and the relay log. If `max_relay_log_size` is greater than 0, it constrains the size of the relay log, which enables you to have different sizes for the two logs. You must set `max_relay_log_size` to between 4096 bytes and 1GB (inclusive), or to 0. The default value is 0. This variable was added in MySQL 4.0.14. See [Section 6.3, “Replication Implementation Details”](#).

- `max_seeks_for_key`

Limit the assumed maximum number of seeks when looking up rows based on a key. The MySQL optimizer assumes that no more than this number of key seeks are required when searching for matching rows in a table by scanning an index, regardless of the actual cardinality of the index (see [Section 13.5.4.11, “SHOW INDEX Syntax”](#)). By setting this to a low value (say, 100), you can force MySQL to prefer indexes instead of table scans.

This variable was added in MySQL 4.0.14.

- `max_sort_length`

The number of bytes to use when sorting `BLOB` or `TEXT` values. Only the first `max_sort_length` bytes of each value are used; the rest are ignored.

- `max_tmp_tables`

The maximum number of temporary tables a client can keep open at the same time. (This option does not yet do anything.)

- `max_user_connections`

The maximum number of simultaneous connections allowed to any given MySQL account. A value of 0 means “no limit.” This variable was added in MySQL 3.23.34.

This variable has only a global form.

- `max_write_lock_count`

After this many write locks, allow some pending read lock requests to be processed in between. This variable was added in MySQL 3.23.7.

- `myisam_data_pointer_size`

The default pointer size in bytes, to be used by `CREATE TABLE` for `MyISAM` tables when no `MAX_ROWS` option is specified. This variable cannot be less than 2 or larger than 7. The default value is 4. This variable was added in MySQL 4.1.2. See [Section A.2.11, “The table is full”](#).

- `myisam_max_extra_sort_file_size`

If the temporary file used for fast `MyISAM` index creation would be larger than using the key cache by the amount specified here, prefer the key cache method. This is mainly used to force long character keys in large tables to use the slower key cache method to create the index. This variable was added in MySQL 3.23.37. **Note:** The value is given in megabytes before 4.0.3 and in bytes thereafter.

- `myisam_max_sort_file_size`

The maximum size of the temporary file that MySQL is allowed to use while re-creating a `MyISAM` index (during `REPAIR TABLE`, `ALTER TABLE`, or `LOAD DATA INFILE`). If the file size would be larger than this value, the index is created using the key cache instead, which is slower. This variable was added in MySQL 3.23.37. **Note:** The value is given in megabytes before 4.0.3 and in bytes thereafter.

- `myisam_recover_options`

The value of the `--myisam-recover` option. See [Section 5.2.1, “Command Options”](#). This variable was added in MySQL 3.23.36.

- `myisam_repair_threads`

If this value is greater than 1, `MyISAM` table indexes are created in parallel (each index in its own thread) during the `Repair by sorting` process. The default value is 1.

**Note:** Multi-threaded repair is still *beta-quality* code. This variable was added in MySQL 4.0.13.

- `myisam_sort_buffer_size`

The size of the buffer that is allocated when sorting `MyISAM` indexes during a `REPAIR TABLE` or when creating indexes with `CREATE INDEX` or `ALTER TABLE`. This variable was added in MySQL 3.23.16.

- `myisam_stats_method`

How the server treats `NULL` values when collecting statistics about the distribution of index values for `MyISAM` tables. This variable has two possible values, `nulls_equal` and `nulls_unequal`. For `nulls_equal`, all `NULL` index values are considered equal and form a single value group that has a size equal to the number of `NULL` values. For `nulls_unequal`, `NULL` values are considered unequal, and each `NULL` forms a distinct value group of size 1.

The method that is used for generating table statistics influences how the optimizer chooses indexes for query execution, as described in [Section 7.4.7, “MyISAM Index Statistics Collection”](#).

This variable was added in MySQL 4.1.15/5.0.14. For older versions, the statistics collection method is equivalent to `nulls_equal`.

- `named_pipe`

On Windows, indicates whether the server supports connections over named pipes. This variable was added in MySQL 3.23.50.

- `ndb_autoincrement_prefetch_sz`

Determines the probability of gaps in an autoincremented column. Set to 1 to minimize this. Set to a high value for optimization — makes inserts faster, but decreases the likelihood that consecutive autoincrement numbers will be used in a batch of inserts. Default value: 32. Minimum value: 1.

- `ndb_cache_check_time`

The number of milliseconds to wait before checking the `NDB` query cache. Setting this to 0 (the default and minimum value) means that the `NDB` query cache will be checked for validation on every query.

The recommended maximum value for this variable is 1000, which means that the query cache is checked once per second. A larger value means the `NDB` query cache is less often checked and inval-



idated due to updates on a different `mysqld`. It is generally not desirable to set this to a value greater than `2000`.

- `ndb_force_send`

Forces sending of buffers to `NDB` immediately, without waiting for other threads. Defaults to `ON`.

- `ndb_index_stat_cache_entries`

Sets the granularity of the statistics by determining the number of starting and ending keys to store in the statistics memory cache. Zero means no caching takes place; in this case, the data nodes are always queries directly. Default value: `32`.

- `ndb_index_stat_enable`

Use `NDB` index statistics in query optimization. Defaults to `ON`.

- `ndb_index_stat_update_freq`

How often to query data nodes instead of the statistics cache. For example, a value of `20` (the default) means to direct every 20<sup>th</sup> query to the data nodes.

- `ndb_report_thresh_binlog_epoch_slip`

This is a threshold on the number of epochs to be behind before reporting binlog status. For example, a value of `3` (the default) means that if the difference between which epoch has been received from the storage nodes and which epoch has been applied to the binlog is 3 or more, a status message will be sent to the cluster log.

- `ndb_report_thresh_binlog_mem_usage`

This is a threshold on the percentage of free memory remaining before reporting binlog status. For example, a value of `10` (the default) means that if the amount of available memory for receiving binlog data from the data nodes falls below 10%, a status message will be sent to the cluster log.

- `ndb_use_exact_count`

Forces `NDB` to use a count of records during `SELECT COUNT(*)` query planning to speed up this type of query. The default value is `ON`. For faster queries overall, disable this feature by setting the value of `ndb_use_exact_count` to `OFF`.

- `ndb_use_transactions`

You can disable `NDB` transaction support by setting this variable's values to `OFF` (not recommended). The default is `ON`.

- `net_buffer_length`

Each client thread is associated with a connection buffer and result buffer. Both begin with a size given by `net_buffer_length` but are dynamically enlarged up to `max_allowed_packet` bytes as needed. The result buffer shrinks to `net_buffer_length` after each SQL statement.

This variable should not normally be changed, but if you have very little memory, you can set it to the expected length of statements sent by clients. If statements exceed this length, the connection buffer is automatically enlarged.

- `net_read_timeout`

The number of seconds to wait for more data from a connection before aborting the read. This

timeout applies only to TCP/IP connections, not to connections made via Unix socket files, named pipes, or shared memory. When the server is reading from the client, `net_read_timeout` is the timeout value controlling when to abort. When the server is writing to the client, `net_write_timeout` is the timeout value controlling when to abort. See also `slave_net_timeout`. This variable was added in MySQL 3.23.20.

- `net_retry_count`

If a read on a communication port is interrupted, retry this many times before giving up. This value should be set quite high on FreeBSD because internal interrupts are sent to all threads. This variable was added in MySQL 3.23.7.

- `net_write_timeout`

The number of seconds to wait for a block to be written to a connection before aborting the write. This timeout applies only to TCP/IP connections, not to connections made via Unix socket files, named pipes, or shared memory. See also `net_read_timeout`. This variable was added in MySQL 3.23.20.

- `new`

This variable is used in MySQL 4.0 to turn on some 4.1 behaviors. This variable was added in MySQL 4.0.12.

- `old_passwords`

Whether the server should use pre-4.1-style passwords for MySQL user accounts. This variable was added in MySQL 4.1.1.

- `one_shot`

This is not a variable, but it can be used when setting some variables. It is described in [Section 13.5.3, “SET Syntax”](#).

- `open_files_limit`

The number of files that the operating system allows `mysqld` to open. This is the real value allowed by the system and might be different from the value you gave using the `--open-files-limit` option to `mysqld` or `mysqld_safe`. The value is 0 on systems where MySQL can't change the number of open files. This variable was added in MySQL 3.23.20.

- `pid_file`

The pathname of the process ID (PID) file. This variable can be set with the `--pid-file` option. This variable was added in MySQL 3.23.23.

- `port`

The number of the port on which the server listens for TCP/IP connections. This variable can be set with the `--port` option.

- `preload_buffer_size`

The size of the buffer that is allocated when preloading indexes. This variable was added in MySQL 4.1.1.

- `prepared_stmt_count`

The current number of prepared statements. (The maximum number of statements is given by the `max_prepared_stmt_count` system variable.) This variable was added in MySQL 4.1.19.

- `protocol_version`

The version of the client/server protocol used by the MySQL server. This variable was added in MySQL 3.23.18.

- `query_alloc_block_size`

The allocation size of memory blocks that are allocated for objects created during statement parsing and execution. If you have problems with memory fragmentation, it might help to increase this a bit. This variable was added in MySQL 4.0.16.

- `query_cache_limit`

Don't cache results that are larger than this number of bytes. The default value is 1MB. This variable was added in MySQL 4.0.1.

- `query_cache_min_res_unit`

The minimum size for blocks allocated by the query cache. The default value is 4KB. Tuning information for this variable is given in [Section 5.13.3, “Query Cache Configuration”](#). This variable is present from MySQL 4.1.

- `query_cache_size`

The amount of memory allocated for caching query results. The default value is 0, which disables the query cache. The allowable values are multiples of 1024; other values are rounded down to the nearest multiple. Note that `query_cache_size` bytes of memory are allocated even if `query_cache_type` is set to 0. This variable was added in MySQL 4.0.1.

- `query_cache_type`

Set the query cache type. Setting the `GLOBAL` value sets the type for all clients that connect thereafter. Individual clients can set the `SESSION` value to affect their own use of the query cache.

Option	Description
0 or OFF	Don't cache results in or retrieve results from the query cache. Note that this does not deallocate the query cache buffer. To do that, you should set <code>query_cache_size</code> to 0.
1 or ON	Cache all query results except for those that begin with <code>SELECT SQL_NO_CACHE</code> .
2 or DEMAND	Cache results only for queries that begin with <code>SELECT SQL_CACHE</code> .

This variable was added in MySQL 4.0.3.

- `query_cache_wlock_invalidate`

Normally, when one client acquires a `WRITE` lock on a `MyISAM` table, other clients are not blocked from issuing statements that read from the table if the query results are present in the query cache. Setting this variable to 1 causes acquisition of a `WRITE` lock for a table to invalidate any queries in the query cache that refer to the table. This forces other clients that attempt to access the table to wait while the lock is in effect. This variable was added in MySQL 4.0.19.

- `query_prealloc_size`

The size of the persistent buffer used for statement parsing and execution. This buffer is not freed between statements. If you are running complex queries, a larger `query_prealloc_size` value might be helpful in improving performance, because it can reduce the need for the server to perform memory allocation during query execution operations.

This variable was added in MySQL 4.0.16.

- `range_alloc_block_size`

The size of blocks that are allocated when doing range optimization. This variable was added in MySQL 4.0.16.

- `read_buffer_size`

Each thread that does a sequential scan allocates a buffer of this size for each table it scans. If you do many sequential scans, you might want to increase this value. This variable was added in MySQL 4.0.3. Previously, it was named `record_buffer`.

- `read_only`

When the variable is set to `ON` for a replication slave server, it causes the slave to allow no updates except from slave threads or from users that have the `SUPER` privilege. This can be useful to ensure that a slave server accepts updates only from its master server and not from clients. This variable was added in MySQL 4.0.14.

- `read_rnd_buffer_size`

When reading rows in sorted order following a key-sorting operation, the rows are read through this buffer to avoid disk seeks. Setting the variable to a large value can improve `ORDER BY` performance by a lot. However, this is a buffer allocated for each client, so you should not set the global variable to a large value. Instead, change the session variable only from within those clients that need to run large queries. This variable was added in MySQL 4.0.3. Previously, it was named `record_rnd_buffer`.

- `relay_log_purge`

Disables or enables automatic purging of relay logs as soon as they are not needed any more. The default value is 1 (`ON`). This variable was added in MySQL 4.1.1.

- `safe_show_database`

Do not show databases for which the user has no database or table privileges. This can improve security if you are concerned about people being able to see what databases other users have. See also `skip_show_database`.

This variable was removed in MySQL 4.0.5. Beginning with this version, you should instead use the `SHOW DATABASES` privilege to control access by MySQL accounts to databases.

- `rpl_recovery_rank`

This variable is unused.

- `secure_auth`

If the MySQL server has been started with the `--secure-auth` option, it blocks connections from all accounts that have passwords stored in the old (pre-4.1) format. In that case, the value of this variable is `ON`, otherwise it is `OFF`.

You should enable this option if you want to prevent all use of passwords in the old format (and hence insecure communication over the network). This variable was added in MySQL 4.1.1.

Server startup fails with an error if this option is enabled and the privilege tables are in pre-4.1 format.

- `server_id`

The server ID. This value is set by the `--server-id` option. It is used for replication to enable master and slave servers to identify themselves uniquely. This variable was added in MySQL 3.23.26.

- `shared_memory`

(Windows only.) Whether the server allows shared-memory connections. This variable was added in MySQL 4.1.1.

- `shared_memory_base_name`

(Windows only.) The name of shared memory to use for shared-memory connections. This is useful when running multiple MySQL instances on a single physical machine. This variable was added in MySQL 4.1.0.

- `skip_external_locking`

This is `OFF` if `mysqld` uses external locking, `ON` if external locking is disabled. This variable was added in MySQL 4.0.3. Previously, it was named `skip_locking`.

- `skip_networking`

This is `ON` if the server allows only local (non-TCP/IP) connections. On Unix, local connections use a Unix socket file. On Windows, local connections use a named pipe or shared memory. On NetWare, only TCP/IP connections are supported, so do not set this variable to `ON`. This variable can be set to `ON` with the `--skip-networking` option. This variable was added in MySQL 3.22.23.

- `skip_show_database`

This prevents people from using the `SHOW DATABASES` statement if they do not have the `SHOW DATABASES` privilege. This can improve security if you are concerned about people being able to see what databases other users have. See also `safe_show_database`. This variable was added in MySQL 3.23.4. As of MySQL 4.0.2, its effect also depends on the `SHOW DATABASES` privilege: If the variable value is `ON`, the `SHOW DATABASES` statement is allowed only to users who have the `SHOW DATABASES` privilege, and the statement displays all database names. If the value is `OFF`, `SHOW DATABASES` is allowed to all users, but displays each database name only if the user has the `SHOW DATABASES` privilege or some privilege for the database. Note that any global privilege is a privilege for the database.

- `slave_compressed_protocol`

Whether to use compression of the master/slave protocol if both the slave and the master support it. This variable was added in MySQL 4.0.3.

- `slave_load_tmpdir`

The name of the directory where the slave creates temporary files for replicating `LOAD DATA INFILE` statements. This variable was added in MySQL 4.0.0.

- `slave_net_timeout`

The number of seconds to wait for more data from a master/slave connection before aborting the read. This timeout applies only to TCP/IP connections, not to connections made via Unix socket files, named pipes, or shared memory. This variable was added in MySQL 3.23.40.

- `slave_skip_errors`

The replication errors that the slave should skip (ignore). This variable was added in MySQL 3.23.47.

- `slave_transaction_retries`

If a replication slave SQL thread fails to execute a transaction because of an `InnoDB` deadlock or `InnoDB`'s `innodb_lock_wait_timeout` or `NDB Cluster`'s `TransactionDeadlock-DetectionTimeout` or `TransactionInactiveTimeout` was exceeded, it automatically retries `slave_transaction_retries` times before stopping with an error. The default in MySQL 4.1 is 0. You must explicitly set the value to greater than 0 to enable the “retry” behavior, which is probably a good idea.

- `slow_launch_time`

If creating a thread takes longer than this many seconds, the server increments the `Slow_launch_threads` status variable. This variable was added in MySQL 3.23.15.

- `socket`

On Unix platforms, this variable is the name of the socket file that is used for local client connections. The default is `/tmp/mysql.sock`. (For some distribution formats, the directory might be different, such as `/var/lib/mysql` for RPMs.)

On Windows, this variable is the name of the named pipe that is used for local client connections. The default value is `MySQL` (not case sensitive).

- `sort_buffer_size`

Each thread that needs to do a sort allocates a buffer of this size. Increase this value for faster `ORDER BY` or `GROUP BY` operations. See [Section A.4.4, “Where MySQL Stores Temporary Files”](#).

- `sql_mode`

The current server SQL mode. This variable was added in MySQL 3.23.41. It can be set dynamically as of MySQL 4.1.1. See [Section 5.2.5, “SQL Modes”](#).

- `sql_slave_skip_counter`

The number of events from the master that a slave server should skip. See [Section 13.6.2.6, “SET GLOBAL SQL\\_SLAVE\\_SKIP\\_COUNTER Syntax”](#). This variable was added in MySQL 3.23.33.

- `storage_engine`

This variable is a synonym for `table_type`. It was added in MySQL 4.1.2.

- `sync_binlog`

If the value of this variable is positive, the MySQL server synchronizes its binary log to disk (using `fsync()`) after every `sync_binlog` writes to the binary log. Note that there is one write to the binary log per statement if autocommit is enabled, and one write per transaction otherwise. The default value is 0, which does no synchronizing to disk. A value of 1 is the safest choice, because in the event of a crash you lose at most one statement or transaction from the binary log. However, it is

also the slowest choice (unless the disk has a battery-backed cache, which makes synchronization very fast). This variable was added in MySQL 4.1.3.

If the value of `sync_binlog` is 0 (the default), no extra flushing is done. The server relies on the operating system to flush the file contents occasionally as for any other file.

- `sync_frm`

If this variable is set to 1, when any non-temporary table is created its `.frm` file is synchronized to disk (using `fdatasync()`). This is slower but safer in case of a crash. The default is 1. This was added as a command-line option in MySQL 4.0.18. It is also a settable global variable as of MySQL 4.1.3.

- `system_time_zone`

The server system time zone. When the server begins executing, it inherits a time zone setting from the machine defaults, possibly modified by the environment of the account used for running the server or the startup script. The value is used to set `system_time_zone`. Typically the time zone is specified by the `TZ` environment variable. It also can be specified using the `--timezone` option of the `mysqld_safe` script. This variable was added in MySQL 4.1.3.

- `table_cache`

The number of open tables for all threads. Increasing this value increases the number of file descriptors that `mysqld` requires. You can check whether you need to increase the table cache by checking the `Opened_tables` status variable. See [Section 5.2.4, “Status Variables”](#). If the value of `Opened_tables` is large and you do not do `FLUSH TABLES` often (which just forces all tables to be closed and reopened), then you should increase the value of the `table_cache` variable. For more information about the table cache, see [Section 7.4.8, “How MySQL Opens and Closes Tables”](#).

- `table_type`

The default table type (storage engine). To set the table type at server startup, use the `--default-table-type` option. This variable was added in MySQL 3.23.0. See [Section 5.2.1, “Command Options”](#).

- `thread_cache_size`

How many threads the server should cache for reuse. When a client disconnects, the client's threads are put in the cache if there are fewer than `thread_cache_size` threads there. Requests for threads are satisfied by reusing threads taken from the cache if possible, and only when the cache is empty is a new thread created. This variable can be increased to improve performance if you have a lot of new connections. (Normally, this doesn't provide a notable performance improvement if you have a good thread implementation.) By examining the difference between the `Connections` and `Threads_created` status variables, you can see how efficient the thread cache is. For details, see [Section 5.2.4, “Status Variables”](#). This variable was added in MySQL 3.23.16.

- `thread_concurrency`

On Solaris, `mysqld` calls `thr_setconcurrency()` with this value. This function enables applications to give the threads system a hint about the desired number of threads that should be run at the same time. This variable was added in MySQL 3.23.7.

- `thread_stack`

The stack size for each thread. Many of the limits detected by the `crash-me` test are dependent on this value. The default is large enough for normal operation. See [Section 7.1.4, “The MySQL Benchmark Suite”](#). The default is 64KB before MySQL 4.0.10 and 192KB thereafter.

- `time_format`

This variable is not implemented.

- `time_zone`

The current time zone. This variable is used to initialize the time zone for each client that connects. By default, the initial value of this is `'SYSTEM'` (which means, “use the value of `system_time_zone`”). The value can be specified explicitly at server startup with the `--default-time-zone` option. See [Section 5.10.8, “MySQL Server Time Zone Support”](#). This variable was added in MySQL 4.1.3.

- `timezone`

The time zone for the server. This is set from the `TZ` environment variable when `mysqld` is started. The time zone also can be set by giving a `--timezone` argument to `mysqld_safe`. This variable was added in MySQL 3.23.15. As of MySQL 4.1.3, it is obsolete and has been replaced by the `system_time_zone` variable. See [Section A.4.6, “Time Zone Problems”](#).

- `tmp_table_size`

The maximum size of in-memory temporary tables. (The actual limit is determined as the smaller of `max_heap_table_size` and `tmp_table_size`.) If an in-memory temporary table exceeds the limit, MySQL automatically converts it to an on-disk `MyISAM` table. Increase the value of `tmp_table_size` (and `max_heap_table_size` if necessary) if you do many advanced `GROUP BY` queries and you have lots of memory.

- `tmpdir`

The directory used for temporary files and temporary tables. Starting from MySQL 4.1, this variable can be set to a list of several paths that are used in round-robin fashion. Paths should be separated by colon characters (`:`) on Unix and semicolon characters (`;`) on Windows, NetWare, and OS/2.

The multiple-directory feature can be used to spread the load between several physical disks. If the MySQL server is acting as a replication slave, you should not set `tmpdir` to point to a directory on a memory-based filesystem or to a directory that is cleared when the server host restarts. A replication slave needs some of its temporary files to survive a machine restart so that it can replicate temporary tables or `LOAD DATA INFILE` operations. If files in the temporary file directory are lost when the server restarts, replication fails. However, if you are using MySQL 4.0.0 or later, you can set the slave's temporary directory using the `slave_load_tmpdir` variable. In that case, the slave won't use the general `tmpdir` value and you can set `tmpdir` to a non-permanent location.

This variable was added in MySQL 3.22.4.

- `transaction_alloc_block_size`

The amount in bytes by which to increase a per-transaction memory pool which needs memory. See the description of `transaction_prealloc_size`. This variable was added in MySQL 4.0.16.

- `transaction_prealloc_size`

There is a per-transaction memory pool from which various transaction-related allocations take memory. The initial size of the pool in bytes is `transaction_prealloc_size`. For every allocation that cannot be satisfied from the pool because it has insufficient memory available, the pool is increased by `transaction_alloc_block_size` bytes. When the transaction ends, the pool is truncated to `transaction_prealloc_size` bytes.

By making `transaction_prealloc_size` sufficiently large to contain all statements within a single transaction, you can avoid many `malloc()` calls. This variable was added in MySQL 4.0.16.



The `system_time_zone` variable differs from `time_zone`. Although they might have the same value, the latter variable is used to initialize the time zone for each client that connects. See [Section 5.10.8, “MySQL Server Time Zone Support”](#).

- `tx_isolation`

The default transaction isolation level. This variable was added in MySQL 4.0.3.

This variable is set by the `SET TRANSACTION ISOLATION LEVEL` statement. See [Section 13.4.6, “SET TRANSACTION Syntax”](#). If you set `tx_isolation` directly to an isolation level name that contains a space, the name should be enclosed within quotes, with the space replaced by a dash. For example:

```
SET tx_isolation = 'READ-COMMITTED';
```

- `version`

The version number for the server.

- `version_bdb`

The BDB storage engine version. This variable was added in MySQL 3.23.31 with the name `bdb_version` and renamed to `version_bdb` in MySQL 4.1.1.

- `version_comment`

The `configure` script has a `--with-comment` option that allows a comment to be specified when building MySQL. This variable contains the value of that comment. This variable was added in MySQL 4.0.17.

- `version_compile_machine`

The type of machine or architecture on which MySQL was built. This variable was added in MySQL 4.1.1.

- `version_compile_os`

The type of operating system on which MySQL was built. This variable was added in MySQL 4.0.19.

- `wait_timeout`

The number of seconds the server waits for activity on a non-interactive connection before closing it. This timeout applies only to TCP/IP connections, not to connections made via Unix socket files, named pipes, or shared memory.

On thread startup, the session `wait_timeout` value is initialized from the global `wait_timeout` value or from the global `interactive_timeout` value, depending on the type of client (as defined by the `CLIENT_INTERACTIVE` connect option to `mysql_real_connect()`). See also [interactive\\_timeout](#).

## 5.2.3. Using System Variables

The `mysql` server maintains many system variables that indicate how it is configured. [Section 5.2.2, “System Variables”](#), describes the meaning of these variables. Each system variable has a default value. System variables can be set at server startup using options on the command line or in an option file. As

of MySQL 4.0.3, most of them can be changed dynamically while the server is running by means of the `SET` statement, which enables you to modify operation of the server without having to stop and restart it. You can refer to system variable values in expressions.

Beginning with MySQL 4.0.3, the server maintains two kinds of system variables. Global variables affect the overall operation of the server. Session variables affect its operation for individual client connections. A given system variable can have both a global and a session value. Global and session system variables are related as follows:

- When the server starts, it initializes all global variables to their default values. These defaults can be changed by options specified on the command line or in an option file. (See [Section 4.3, “Specifying Program Options”](#).)
- The server also maintains a set of session variables for each client that connects. The client's session variables are initialized at connect time using the current values of the corresponding global variables. For example, the client's SQL mode is controlled by the session `sql_mode` value, which is initialized when the client connects to the value of the global `sql_mode` value.

System variable values can be set globally at server startup by using options on the command line or in an option file. When you use a startup option to set a variable that takes a numeric value, the value can be given with a suffix of `K`, `M`, or `G` (either uppercase or lowercase) to indicate a multiplier of 1024, 1024<sup>2</sup> or 1024<sup>3</sup>; that is, units of kilobytes, megabytes, or gigabytes, respectively. Thus, the following command starts the server with a query cache size of 16 megabytes and a maximum packet size of one gigabyte:

```
mysqld --query_cache_size=16M --max_allowed_packet=1G
```

Before MySQL 4.0.2, use this syntax instead:

```
mysqld --set-variable=query_cache_size=16M \
  --set-variable=max_allowed_packet=1G
```

Within an option file, those variables are set like this:

```
[mysqld]
query_cache_size=16M
max_allowed_packet=1G
```

Or like this before MySQL 4.0.2:

```
[mysqld]
set-variable=query_cache_size=16M
set-variable=max_allowed_packet=1G
```

The lettercase of suffix letters does not matter; `16M` and `16m` are equivalent, as are `1G` and `1g`.

If you want to restrict the maximum value to which a system variable can be set at runtime with the `SET` statement, you can specify this maximum by using an option of the form

`--maximum-var_name=value` at server startup. For example, to prevent the value of `query_cache_size` from being increased to more than 32MB at runtime, use the option `--maximum-query_cache_size=32M`. This feature is available as of MySQL 4.0.2.

Many system variables are dynamic and can be changed while the server runs by using the `SET` statement. For a list, see [Section 5.2.3.2, “Dynamic System Variables”](#). To change a system variable with `SET`, refer to it as `var_name`, optionally preceded by a modifier:

- To indicate explicitly that a variable is a global variable, precede its name by `GLOBAL` or

`@@global.` The `SUPER` privilege is required to set global variables.

- To indicate explicitly that a variable is a session variable, precede its name by `SESSION`, `@@session.`, or `@@`. Setting a session variable requires no special privilege, but a client can change only its own session variables, not those of any other client.
- `LOCAL` and `@@local.` are synonyms for `SESSION` and `@@session.`
- If no modifier is present, `SET` changes the session variable.

A `SET` statement can contain multiple variable assignments, separated by commas. If you set several system variables, the most recent `GLOBAL` or `SESSION` modifier in the statement is used for following variables that have no modifier specified.

Examples:

```
SET sort_buffer_size=10000;
SET @@local.sort_buffer_size=10000;
SET GLOBAL sort_buffer_size=1000000, SESSION sort_buffer_size=1000000;
SET @@sort_buffer_size=1000000;
SET @@global.sort_buffer_size=1000000, @@local.sort_buffer_size=1000000;
```

When you assign a value to a system variable with `SET`, you cannot use suffix letters in the value (as can be done with startup options). However, the value can take the form of an expression:

```
SET sort_buffer_size = 10 * 1024 * 1024;
```

The `@@var_name` syntax for system variables is supported for compatibility with some other database systems.

If you change a session system variable, the value remains in effect until your session ends or until you change the variable to a different value. The change is not visible to other clients.

If you change a global system variable, the value is remembered and used for new connections until the server restarts. (To make a global system variable setting permanent, you should set it in an option file.) The change is visible to any client that accesses that global variable. However, the change affects the corresponding session variable only for clients that connect after the change. The global variable change does not affect the session variable for any client that is currently connected (not even that of the client that issues the `SET GLOBAL` statement).

To prevent incorrect usage, MySQL produces an error if you use `SET GLOBAL` with a variable that can only be used with `SET SESSION` or if you do not specify `GLOBAL` (or `@@global.`) when setting a global variable.

To set a `SESSION` variable to the `GLOBAL` value or a `GLOBAL` value to the compiled-in MySQL default value, use the `DEFAULT` keyword. For example, the following two statements are identical in setting the session value of `max_join_size` to the global value:

```
SET max_join_size=DEFAULT;
SET @@session.max_join_size=@@global.max_join_size;
```

Not all system variables can be set to `DEFAULT`. In such cases, use of `DEFAULT` results in an error.

You can refer to the values of specific global or session system variables in expressions by using one of the `@@`-modifiers. For example, you can retrieve values in a `SELECT` statement like this:

```
SELECT @@global.sql_mode, @@session.sql_mode, @@sql_mode;
```

When you refer to a system variable in an expression as `@@var_name` (that is, when you do not specify `@@global.` or `@@session.`), MySQL returns the session value if it exists and the global value otherwise. (This differs from `SET @@var_name = value`, which always refers to the session value.)

*Note:* Some system variables can be enabled with the `SET` statement by setting them to `ON` or `1`, or disabled by setting them to `OFF` or `0`. However, to set such a variable on the command line or in an option file, you must set it to `1` or `0`; setting it to `ON` or `OFF` will not work. For example, on the command line, `--delay_key_write=1` works but `--delay_key_write=ON` does not.

To display system variable names and values, use the `SHOW VARIABLES` statement:

```
mysql> SHOW VARIABLES;
```

Variable_name	Value
back_log	50
basedir	/usr/local/mysql
bdb_cache_size	8388600
bdb_home	/usr/local/mysql
bdb_log_buffer_size	32768
bdb_logdir	
bdb_max_lock	10000
bdb_shared_data	OFF
bdb_tmpdir	/tmp/
binlog_cache_size	32768
bulk_insert_buffer_size	8388608
character_set_client	latin1
character_set_connection	latin1
character_set_database	latin1
character_set_results	latin1
character_set_server	latin1
character_set_system	utf8
character_sets_dir	/usr/local/mysql/share/charsets/
collation_connection	latin1_swedish_ci
collation_database	latin1_swedish_ci
collation_server	latin1_swedish_ci
..	
innodb_additional_mem_pool_size	1048576
innodb_autoextend_increment	8
innodb_buffer_pool_ave_mem_mb	0
innodb_buffer_pool_size	8388608
innodb_data_file_path	ibdata1:10M:autoextend
innodb_data_home_dir	
..	
version	4.1.18-max-log
version_comment	MySQL Community Edition - Max (GPL)
version_compile_machine	i686
version_compile_os	pc-linux-gnu
wait_timeout	28800

With a `LIKE` clause, the statement displays only those variables that match the pattern. To obtain a specific variable name, use a `LIKE` clause as shown:

```
SHOW VARIABLES LIKE 'max_join_size';
SHOW SESSION VARIABLES LIKE 'max_join_size';
```

To get a list of variables whose name match a pattern, use the `'%'` wildcard character in a `LIKE` clause:

```
SHOW VARIABLES LIKE '%size%';
SHOW GLOBAL VARIABLES LIKE '%size%';
```

Wildcard characters can be used in any position within the pattern to be matched. Strictly speaking, because `'_'` is a wildcard that matches any single character, you should escape it as `'\_'` to match it literally. In practice, this is rarely necessary.

For `SHOW VARIABLES`, if you specify neither `GLOBAL` nor `SESSION`, MySQL returns `SESSION` values.

The reason for requiring the `GLOBAL` keyword when setting `GLOBAL`-only variables but not when retrieving them is to prevent problems in the future. If we were to remove a `SESSION` variable that has the same name as a `GLOBAL` variable, a client with the `SUPER` privilege might accidentally change the `GLOBAL` variable rather than just the `SESSION` variable for its own connection. If we add a `SESSION` variable with the same name as a `GLOBAL` variable, a client that intends to change the `GLOBAL` variable might find only its own `SESSION` variable changed.

### 5.2.3.1. Structured System Variables

Structured system variables are supported beginning with MySQL 4.1.1. A structured variable differs from a regular system variable in two respects:

- Its value is a structure with components that specify server parameters considered to be closely related.
- There might be several instances of a given type of structured variable. Each one has a different name and refers to a different resource maintained by the server.

In MySQL 4.1 (4.1.1 and above), MySQL supports one structured variable type. It specifies parameters that govern the operation of key caches. A key cache structured variable has these components:

- `key_buffer_size`
- `key_cache_block_size`
- `key_cache_division_limit`
- `key_cache_age_threshold`

The purpose of this section is to describe the syntax for referring to structured variables. Key cache variables are used for syntax examples, but specific details about how key caches operate are found elsewhere, in [Section 7.4.6, “The MyISAM Key Cache”](#).

To refer to a component of a structured variable instance, you can use a compound name in *instance\_name.component\_name* format. Examples:

```
hot_cache.key_buffer_size
hot_cache.key_cache_block_size
cold_cache.key_cache_block_size
```

For each structured system variable, an instance with the name of `default` is always predefined. If you refer to a component of a structured variable without any instance name, the `default` instance is used. Thus, `default.key_buffer_size` and `key_buffer_size` both refer to the same system variable.

Structured variable instances and components follow these naming rules:

- For a given type of structured variable, each instance must have a name that is unique *within* variables of that type. However, instance names need not be unique *across* structured variable types. For example, each structured variable has an instance named `default`, so `default` is not unique across variable types.
- The names of the components of each structured variable type must be unique across all system variable names. If this were not true (that is, if two different types of structured variables could share

component member names), it would not be clear which default structured variable to use for references to member names that are not qualified by an instance name.

- If a structured variable instance name is not legal as an unquoted identifier, refer to it as a quoted identifier using backticks. For example, `hot-cache` is not legal, but ``hot-cache`` is.
- `global`, `session`, and `local` are not legal instance names. This avoids a conflict with notation such as `@global.var_name` for referring to non-structured system variables.

At the moment, the first two rules have no possibility of being violated because the only structured variable type is the one for key caches. These rules will assume greater significance if some other type of structured variable is created in the future.

With one exception, it is allowable to refer to structured variable components using compound names in any context where simple variable names can occur. For example, you can assign a value to a structured variable using a command-line option:

```
shell> mysql --hot_cache.key_buffer_size=64K
```

In an option file, use this syntax:

```
[mysqld]
hot_cache.key_buffer_size=64K
```

If you start the server with such an option, it creates a key cache named `hot_cache` with a size of 64KB in addition to the default key cache that has a default size of 8MB.

Suppose that you start the server as follows:

```
shell> mysql --key_buffer_size=256K \
--extra_cache.key_buffer_size=128K \
--extra_cache.key_cache_block_size=2048
```

In this case, the server sets the size of the default key cache to 256KB. (You could also have written `--default.key_buffer_size=256K`.) In addition, the server creates a second key cache named `extra_cache` that has a size of 128KB, with the size of block buffers for caching table index blocks set to 2048 bytes.

The following example starts the server with three different key caches having sizes in a 3:1:1 ratio:

```
shell> mysql --key_buffer_size=6M \
--hot_cache.key_buffer_size=2M \
--cold_cache.key_buffer_size=2M
```

Structured variable values may be set and retrieved at runtime as well. For example, to set a key cache named `hot_cache` to a size of 10MB, use either of these statements:

```
mysql> SET GLOBAL hot_cache.key_buffer_size = 10*1024*1024;
mysql> SET @@global.hot_cache.key_buffer_size = 10*1024*1024;
```

To retrieve the cache size, do this:

```
mysql> SELECT @@global.hot_cache.key_buffer_size;
```

However, the following statement does not work. The variable is not interpreted as a compound name, but as a simple string for a `LIKE` pattern-matching operation:

```
mysql> SHOW GLOBAL VARIABLES LIKE 'hot_cache.key_buffer_size';
```

This is the exception to being able to use structured variable names anywhere a simple variable name may occur.

### 5.2.3.2. Dynamic System Variables

Beginning with MySQL 4.0.3, many server system variables are dynamic and can be set at runtime using `SET GLOBAL` or `SET SESSION`. You can also select their values using `SELECT`. See [Section 5.2.3, “Using System Variables”](#).

The following table shows the full list of all dynamic system variables. The last column indicates for each variable whether `GLOBAL` or `SESSION` (or both) apply. The table also lists session options that can be set with the `SET` statement. [Section 13.5.3, “SET Syntax”](#), discusses these options.

Variables that have a type of “string” take a string value. Variables that have a type of “numeric” take a numeric value. Variables that have a type of “boolean” can be set to 0, 1, `ON` or `OFF`. (If you set them on the command line or in an option file, use the numeric values.) Variables that are marked as “enumeration” normally should be set to one of the available values for the variable, but can also be set to the number that corresponds to the desired enumeration value. For enumerated system variables, the first enumeration value corresponds to 0. This differs from `ENUM` columns, for which the first enumeration value corresponds to 1.

Variable Name	Value Type	Type
<code>autocommit</code>	boolean	SESSION
<code>big_tables</code>	boolean	SESSION
<code>binlog_cache_size</code>	numeric	GLOBAL
<code>bulk_insert_buffer_size</code>	numeric	GLOBAL   SESSION
<code>character_set_client</code>	string	GLOBAL   SESSION
<code>character_set_connection</code>	string	GLOBAL   SESSION
<code>character_set_results</code>	string	GLOBAL   SESSION
<code>character_set_server</code>	string	GLOBAL   SESSION
<code>collation_connection</code>	string	GLOBAL   SESSION
<code>collation_server</code>	string	GLOBAL   SESSION
<code>concurrent_insert</code>	boolean	GLOBAL
<code>connect_timeout</code>	numeric	GLOBAL
<code>convert_character_set</code>	string	GLOBAL   SESSION
<code>default_week_format</code>	numeric	GLOBAL   SESSION
<code>delay_key_write</code>	OFF   ON   ALL	GLOBAL
<code>delayed_insert_limit</code>	numeric	GLOBAL
<code>delayed_insert_timeout</code>	numeric	GLOBAL
<code>delayed_queue_size</code>	numeric	GLOBAL
<code>error_count</code>	numeric	SESSION
<code>expire_logs_days</code>	numeric	GLOBAL
<code>flush</code>	boolean	GLOBAL
<code>flush_time</code>	numeric	GLOBAL
<code>foreign_key_checks</code>	boolean	SESSION
<code>ft_boolean_syntax</code>	numeric	GLOBAL
<code>group_concat_max_len</code>	numeric	GLOBAL   SESSION

identity	numeric	SESSION
innodb_autoextend_increment	numeric	GLOBAL
innodb_concurrency_tickets	numeric	GLOBAL
innodb_max_dirty_pages_pct	numeric	GLOBAL
innodb_max_purge_lag	numeric	GLOBAL
innodb_sync_spin_loops	numeric	GLOBAL
innodb_table_locks	boolean	GLOBAL   SESSION
innodb_thread_sleep_delay	numeric	GLOBAL
insert_id	numeric	SESSION
interactive_timeout	numeric	GLOBAL   SESSION
join_buffer_size	numeric	GLOBAL   SESSION
key_buffer_size	numeric	GLOBAL
last_insert_id	numeric	SESSION
local_infile	boolean	GLOBAL
log_warnings	numeric	GLOBAL
long_query_time	numeric	GLOBAL   SESSION
low_priority_updates	boolean	GLOBAL   SESSION
max_allowed_packet	numeric	GLOBAL   SESSION
max_binlog_cache_size	numeric	GLOBAL
max_binlog_size	numeric	GLOBAL
max_connect_errors	numeric	GLOBAL
max_connections	numeric	GLOBAL
max_delayed_threads	numeric	GLOBAL
max_error_count	numeric	GLOBAL   SESSION
max_heap_table_size	numeric	GLOBAL   SESSION
max_insert_delayed_threads	numeric	GLOBAL
max_join_size	numeric	GLOBAL   SESSION
max_prepared_stmt_count	numeric	GLOBAL
max_relay_log_size	numeric	GLOBAL
max_seeks_for_key	numeric	GLOBAL   SESSION
max_sort_length	numeric	GLOBAL   SESSION
max_tmp_tables	numeric	GLOBAL   SESSION
max_user_connections	numeric	GLOBAL
max_write_lock_count	numeric	GLOBAL
multi_read_range	numeric	GLOBAL   SESSION
myisam_data_pointer_size	numeric	GLOBAL
myisam_max_sort_file_size	numeric	GLOBAL   SESSION
myisam_repair_threads	numeric	GLOBAL   SESSION
myisam_sort_buffer_size	numeric	GLOBAL   SESSION
myisam_stats_method	enum	GLOBAL   SESSION
net_buffer_length	numeric	GLOBAL   SESSION
net_read_timeout	numeric	GLOBAL   SESSION



net_retry_count	numeric	GLOBAL   SESSION
net_write_timeout	numeric	GLOBAL   SESSION
old_passwords	numeric	GLOBAL   SESSION
optimizer_prune_level	numeric	GLOBAL   SESSION
optimizer_search_depth	numeric	GLOBAL   SESSION
preload_buffer_size	numeric	GLOBAL   SESSION
query_alloc_block_size	numeric	GLOBAL   SESSION
query_cache_limit	numeric	GLOBAL
query_cache_size	numeric	GLOBAL
query_cache_type	enumeration	GLOBAL   SESSION
query_cache_wlock_invalidate	boolean	GLOBAL   SESSION
query_prealloc_size	numeric	GLOBAL   SESSION
range_alloc_block_size	numeric	GLOBAL   SESSION
read_buffer_size	numeric	GLOBAL   SESSION
read_only	numeric	GLOBAL
read_rnd_buffer_size	numeric	GLOBAL   SESSION
rpl_recovery_rank	numeric	GLOBAL
safe_show_database	boolean	GLOBAL
secure_auth	boolean	GLOBAL
server_id	numeric	GLOBAL
slave_compressed_protocol	boolean	GLOBAL
slave_net_timeout	numeric	GLOBAL
slave_transaction_retries	numeric	GLOBAL
slow_launch_time	numeric	GLOBAL
sort_buffer_size	numeric	GLOBAL   SESSION
sql_auto_is_null	boolean	SESSION
sql_big_selects	boolean	SESSION
sql_big_tables	boolean	SESSION
sql_buffer_result	boolean	SESSION
sql_log_bin	boolean	SESSION
sql_log_off	boolean	SESSION
sql_log_update	boolean	SESSION
sql_low_priority_updates	boolean	GLOBAL   SESSION
sql_max_join_size	numeric	GLOBAL   SESSION
sql_mode	enumeration	GLOBAL   SESSION
sql_notes	boolean	SESSION
sql_quote_show_create	boolean	SESSION
sql_safe_updates	boolean	SESSION
sql_select_limit	numeric	SESSION
sql_slave_skip_counter	numeric	GLOBAL
updatable_views_with_limit	enumeration	GLOBAL   SESSION
sql_warnings	boolean	SESSION

<code>sync_binlog</code>	numeric	GLOBAL
<code>sync_frm</code>	boolean	GLOBAL
<code>storage_engine</code>	enumeration	GLOBAL   SESSION
<code>table_cache</code>	numeric	GLOBAL
<code>table_type</code>	enumeration	GLOBAL   SESSION
<code>thread_cache_size</code>	numeric	GLOBAL
<code>time_zone</code>	string	GLOBAL   SESSION
<code>timestamp</code>	boolean	SESSION
<code>tmp_table_size</code>	enumeration	GLOBAL   SESSION
<code>transaction_alloc_block_size</code>	numeric	GLOBAL   SESSION
<code>transaction_prealloc_size</code>	numeric	GLOBAL   SESSION
<code>tx_isolation</code>	enumeration	GLOBAL   SESSION
<code>unique_checks</code>	boolean	SESSION
<code>wait_timeout</code>	numeric	GLOBAL   SESSION
<code>warning_count</code>	numeric	SESSION

## 5.2.4. Status Variables

The server maintains many status variables that provide information about its operation. You can view these variables and their values by using the `SHOW STATUS` statement:

```
mysql> SHOW STATUS;
```

Variable_name	Value
Aborted_clients	0
Aborted_connects	0
Bytes_received	155372598
Bytes_sent	1176560426
Connections	30023
...	

Many status variables are reset to 0 by the `FLUSH STATUS` statement.

The status variables have the following meanings. The `Com_xxx` statement counter variables were added beginning with MySQL 3.23.47. The `Qcache_xxx` query cache variables were added beginning with MySQL 4.0.1. Otherwise, variables with no version indicated have been present since at least MySQL 3.22.

- `Aborted_clients`

The number of connections that were aborted because the client died without closing the connection properly. See [Section A.2.10, “Communication Errors and Aborted Connections”](#).

- `Aborted_connects`

The number of failed attempts to connect to the MySQL server. See [Section A.2.10, “Communication Errors and Aborted Connections”](#).

- `Binlog_cache_disk_use`

The number of transactions that used the temporary binary log cache but that exceeded the value of

`binlog_cache_size` and used a temporary file to store statements from the transaction. This variable was added in MySQL 4.1.2.

- `Binlog_cache_use`

The number of transactions that used the temporary binary log cache. This variable was added in MySQL 4.1.2.

- `Bytes_received`

The number of bytes received from all clients. This variable was added in MySQL 3.23.7.

- `Bytes_sent`

The number of bytes sent to all clients. This variable was added in MySQL 3.23.7.

- `Com_xxx`

The `Com_xxx` statement counter variables were added beginning with MySQL 3.23.47. They indicate the number of times each `xxx` statement has been executed. There is one status variable for each type of statement. For example, `Com_delete` and `Com_insert` count `DELETE` and `INSERT` statements, respectively. However, if a query result is returned from query cache, the server increments the `Qcache_hits` status variable, not `Com_select`. See Section 5.13.4, “Query Cache Status and Maintenance”.

New `Com_stmt_xxx` status variables have been added in MySQL 4.1.13:

- `Com_stmt_prepare`
- `Com_stmt_execute`
- `Com_stmt_send_long_data`
- `Com_stmt_reset`
- `Com_stmt_close`

Those variables stand for prepared statement commands. Their names refer to the `COM_xxx` command set used in the network layer. In other words, their values increase whenever prepared statement API calls such as `mysql_stmt_prepare()`, `mysql_stmt_execute()`, and so forth are executed. However, `Com_stmt_prepare`, `Com_stmt_execute` and `Com_stmt_close` also increase for `PREPARE`, `EXECUTE`, or `DEALLOCATE PREPARE`, respectively. Additionally, the values of the older (available since MySQL 4.1.3) statement counter variables `Com_prepare_sql`, `Com_execute_sql`, and `Com_dealloc_sql` increase for the `PREPARE`, `EXECUTE`, and `DEALLOCATE PREPARE` statements.

All of the `Com_stmt_xxx` variables are increased even if their argument (a prepared statement) is unknown or an error occurred during execution; in other words, their values correspond to the number of requests issued, not to the number of requests successfully completed.

- `Connections`

The number of connection attempts (successful or not) to the MySQL server.

- `Created_tmp_disk_tables`

The number of temporary tables on disk created automatically by the server while executing statements. This variable was added in MySQL 3.23.24.

- `Created_tmp_files`

How many temporary files `mysqld` has created. This variable was added in MySQL 3.23.28.
- `Created_tmp_tables`

The number of in-memory temporary tables created automatically by the server while executing statements. If `Created_tmp_disk_tables` is large, you may want to increase the `tmp_table_size` value to cause temporary tables to be memory-based instead of disk-based.
- `Delayed_errors`

The number of rows written with `INSERT DELAYED` for which some error occurred (probably `duplicate key`).
- `Delayed_insert_threads`

The number of `INSERT DELAYED` handler threads in use.
- `Delayed_writes`

The number of `INSERT DELAYED` rows written.
- `Flush_commands`

The number of executed `FLUSH` statements.
- `Handler_commit`

The number of internal `COMMIT` statements. This variable was added in MySQL 4.0.2.
- `Handler_discover`

The MySQL server can ask the `NDB Cluster` storage engine if it knows about a table with a given name. This is called discovery. `Handler_discover` indicates the number of times that tables have been discovered. This variable was added in MySQL 4.1.2.
- `Handler_delete`

The number of times a row was deleted from a table.
- `Handler_read_first`

The number of times the first entry was read from an index. If this value is high, it suggests that the server is doing a lot of full index scans; for example, `SELECT col1 FROM foo`, assuming that `col1` is indexed.
- `Handler_read_key`

The number of requests to read a row based on a key. If this value is high, it is a good indication that your tables are properly indexed for your queries.
- `Handler_read_next`

The number of requests to read the next row in key order. This value is incremented if you are querying an index column with a range constraint or if you are doing an index scan.
- `Handler_read_prev`

The number of requests to read the previous row in key order. This read method is mainly used to

optimize `ORDER BY ... DESC`. This variable was added in MySQL 3.23.6.

- `Handler_read_rnd`

The number of requests to read a row based on a fixed position. This value is high if you are doing a lot of queries that require sorting of the result. You probably have a lot of queries that require MySQL to scan entire tables or you have joins that don't use keys properly.

- `Handler_read_rnd_next`

The number of requests to read the next row in the data file. This value is high if you are doing a lot of table scans. Generally this suggests that your tables are not properly indexed or that your queries are not written to take advantage of the indexes you have.

- `Handler_rollback`

The number of internal `ROLLBACK` statements. This variable was added in MySQL 4.0.2.

- `Handler_update`

The number of requests to update a row in a table.

- `Handler_write`

The number of requests to insert a row in a table.

- `Key_blocks_not_flushed`

The number of key blocks in the key cache that have changed but have not yet been flushed to disk. This variable was added in MySQL 4.1.1. It used to be known as `Not_flushed_key_blocks`.

- `Key_blocks_unused`

The number of unused blocks in the key cache. You can use this value to determine how much of the key cache is in use; see the discussion of `key_buffer_size` in [Section 5.2.2, "System Variables"](#). This variable was added in MySQL 4.1.2.

- `Key_blocks_used`

The number of used blocks in the key cache. This value is a high-water mark that indicates the maximum number of blocks that have ever been in use at one time.

- `Key_read_requests`

The number of requests to read a key block from the cache.

- `Key_reads`

The number of physical reads of a key block from disk. If `Key_reads` is large, then your `key_buffer_size` value is probably too small. The cache miss rate can be calculated as `Key_reads/Key_read_requests`.

- `Key_write_requests`

The number of requests to write a key block to the cache.

- `Key_writes`

The number of physical writes of a key block to disk.

- `Max_used_connections`  
The maximum number of connections that have been in use simultaneously since the server started.
- `Not_flushed_delayed_rows`  
The number of rows waiting to be written in `INSERT DELAY` queues.
- `Not_flushed_key_blocks`  
The old name for `Key_blocks_not_flushed` before MySQL 4.1.1.
- `Open_files`  
The number of files that are open.
- `Open_streams`  
The number of streams that are open (used mainly for logging).
- `Open_tables`  
The number of tables that are open.
- `Opened_tables`  
The number of tables that have been opened. If `Opened_tables` is big, your `table_cache` value is probably too small.
- `Qcache_free_blocks`  
The number of free memory blocks in the query cache.
- `Qcache_free_memory`  
The amount of free memory for the query cache.
- `Qcache_hits`  
The number of query cache hits.
- `Qcache_inserts`  
The number of queries added to the query cache.
- `Qcache_lowmem_prunes`  
The number of queries that were deleted from the query cache because of low memory.
- `Qcache_not_cached`  
The number of non-cached queries (not cacheable, or not cached due to the `query_cache_type` setting).
- `Qcache_queries_in_cache`  
The number of queries registered in the query cache.
- `Qcache_total_blocks`

The total number of blocks in the query cache.

- `Questions`

The number of statements that clients have sent to the server.

- `Rpl_status`

The status of fail-safe replication (not yet implemented).

- `Select_full_join`

The number of joins that perform table scans because they do not use indexes. If this value is not 0, you should carefully check the indexes of your tables. This variable was added in MySQL 3.23.25.

- `Select_full_range_join`

The number of joins that used a range search on a reference table. This variable was added in MySQL 3.23.25.

- `Select_range`

The number of joins that used ranges on the first table. This is normally not critical issue even if the value is quite large. This variable was added in MySQL 3.23.25.

- `Select_range_check`

The number of joins without keys that check for key usage after each row. (If this is not equal to 0, you should very carefully check the indexes of your tables.) This variable was added in MySQL 3.23.25.

- `Select_scan`

The number of joins that did a full scan of the first table. This variable was added in MySQL 3.23.25.

- `Slave_open_temp_tables`

The number of temporary tables that the slave SQL thread currently has open. This variable was added in MySQL 3.23.29.

- `Slave_running`

This is `ON` if this server is a slave that is connected to a master. This variable was added in MySQL 3.23.16.

- `Slave_retried_transactions`

Total (since startup) number of times the replication slave SQL thread has retried transactions. This variable was added in MySQL 4.1.11.

- `Slow_launch_threads`

The number of threads that have taken more than `slow_launch_time` seconds to create. This variable was added in MySQL 3.23.15.

- `Slow_queries`

The number of queries that have taken more than `long_query_time` seconds. See [Sec-](#)

tion 5.11.5, “The Slow Query Log”.

- `Sort_merge_passes`

The number of merge passes that the sort algorithm has had to do. If this value is large, you should consider increasing the value of the `sort_buffer_size` system variable. This variable was added in MySQL 3.23.28.

- `Sort_range`

The number of sorts that were done with ranges. This variable was added in MySQL 3.23.25.

- `Sort_rows`

The number of sorted rows. This variable was added in MySQL 3.23.25.

- `Sort_scan`

The number of sorts that were done by scanning the table. This variable was added in MySQL 3.23.25.

- `Ssl_xxx`

Variables used for SSL connections. These variables were added in MySQL 4.0.0.

- `Table_locks_immediate`

The number of times that a table lock was acquired immediately. This variable was added in MySQL 3.23.33.

- `Table_locks_waited`

The number of times that a table lock could not be acquired immediately and a wait was needed. If this is high and you have performance problems, you should first optimize your queries, and then either split your table or tables or use replication. This variable was added in MySQL 3.23.33.

- `Threads_cached`

The number of threads in the thread cache. This variable was added in MySQL 3.23.17.

- `Threads_connected`

The number of currently open connections.

- `Threads_created`

The number of threads created to handle connections. If `Threads_created` is big, you may want to increase the `thread_cache_size` value. The cache miss rate can be calculated as `Threads_created` divided by `Connections`. This variable was added in MySQL 3.23.31.

- `Threads_running`

The number of threads that are not sleeping.

- `Uptime`

The number of seconds that the server has been up.



## 5.2.5. SQL Modes

The MySQL server can operate in different SQL modes, and (as of MySQL 4.1) can apply these modes differentially for different clients. This capability enables each application to tailor the server's operating mode to its own requirements.

Modes define what SQL syntax MySQL should support and what kind of data validation checks it should perform. This makes it easier to use MySQL in different environments and to use MySQL together with other database servers.

You can set the default SQL mode by starting `mysqld` with the `--sql-mode="modes"` option. *modes* is a list of different modes separated by comma (',' ) characters. The default value is empty (no modes set). The *modes* value also can be empty (`--sql-mode= "`) if you want to clear it explicitly.

Beginning with MySQL 4.1, you can change the SQL mode at runtime by using a `SET [GLOBAL|SESSION] sql_mode='modes'` statement to set the `sql_mode` system value. Setting the `GLOBAL` variable requires the `SUPER` privilege and affects the operation of all clients that connect from that time on. Setting the `SESSION` variable affects only the current client. Any client can change its own session `sql_mode` value at any time.

You can retrieve the current global or session `sql_mode` value with the following statements:

```
SELECT @@global.sql_mode;
SELECT @@session.sql_mode;
```

The most important `sql_mode` value is `ANSI`, which changes syntax and behavior to be more conformant to standard SQL. This mode is available beginning in MySQL 4.1.1.

The following list describes all supported modes:

- `ANSI_QUOTES`

Treat `"` as an identifier quote character (like the ``` quote character) and not as a string quote character. You can still use ``` to quote identifiers with this mode enabled. With `ANSI_QUOTES` enabled, you cannot use double quotes to quote literal strings, because it is interpreted as an identifier. (Added in MySQL 4.0.0)

- `IGNORE_SPACE`

Allow spaces between a function name and the `(` character. This forces all function names to be treated as reserved words. As a result, if you want to access any database, table, or column name that is a reserved word, you must quote it. For example, because there is a `USER()` function, the name of the `user` table in the `mysql` database and the `User` column in that table become reserved, so you must quote them:

```
SELECT "User" FROM mysql."user";
```

(Added in MySQL 4.0.0)

See [Section 14.1.1, “MyISAM Startup Options”](#).

- `NO_AUTO_VALUE_ON_ZERO`

`NO_AUTO_VALUE_ON_ZERO` affects handling of `AUTO_INCREMENT` columns. Normally, you generate the next sequence number for the column by inserting either `NULL` or `0` into it. `NO_AUTO_VALUE_ON_ZERO` suppresses this behavior for `0` so that only `NULL` generates the next sequence number. (Added in MySQL 4.1.1)

This mode can be useful if 0 has been stored in a table's `AUTO_INCREMENT` column. (Storing 0 is not a recommended practice, by the way.) For example, if you dump the table with `mysqldump` and then reload it, MySQL normally generates new sequence numbers when it encounters the 0 values, resulting in a table with contents different from the one that was dumped. Enabling `NO_AUTO_VALUE_ON_ZERO` before reloading the dump file solves this problem. As of MySQL 4.1.1, `mysqldump` automatically includes a statement in the dump output that enables `NO_AUTO_VALUE_ON_ZERO` to avoid this problem.

- `NO_DIR_IN_CREATE`

When creating a table, ignore all `INDEX DIRECTORY` and `DATA DIRECTORY` directives. This option is useful on slave replication servers. (Added in MySQL 4.0.15)

- `NO_FIELD_OPTIONS`

Do not print MySQL-specific column options in the output of `SHOW CREATE TABLE`. This mode is used by `mysqldump` in portability mode. (Added in MySQL 4.1.1)

- `NO_KEY_OPTIONS`

Do not print MySQL-specific index options in the output of `SHOW CREATE TABLE`. This mode is used by `mysqldump` in portability mode. (Added in MySQL 4.1.1)

- `NO_TABLE_OPTIONS`

Do not print MySQL-specific table options (such as `ENGINE`) in the output of `SHOW CREATE TABLE`. This mode is used by `mysqldump` in portability mode. (Added in MySQL 4.1.1)

- `NO_UNSIGNED_SUBTRACTION`

In integer subtraction operations, do not mark the result as `UNSIGNED` if one of the operands is unsigned. In other words, *the result of a subtraction is always signed whenever this mode is in effect, even if one of the operands is unsigned*. For example, compare the type of column `c2` in table `t1` with that of column `c2` in table `t2`:

```
mysql> SET SQL_MODE='';
mysql> CREATE TABLE test (c1 BIGINT UNSIGNED NOT NULL);
mysql> CREATE TABLE t1 SELECT c1 - 1 AS c2 FROM test;
mysql> DESCRIBE t1;
+-----+-----+-----+-----+-----+-----+
| Field | Type                | Null | Key | Default | Extra |
+-----+-----+-----+-----+-----+-----+
| c2    | bigint(21) unsigned |      |     | 0        |       |
+-----+-----+-----+-----+-----+-----+

mysql> SET SQL_MODE='NO_UNSIGNED_SUBTRACTION';
mysql> CREATE TABLE t2 SELECT c1 - 1 AS c2 FROM test;
mysql> DESCRIBE t2;
+-----+-----+-----+-----+-----+-----+
| Field | Type                | Null | Key | Default | Extra |
+-----+-----+-----+-----+-----+-----+
| c2    | bigint(21)          |      |     | 0        |       |
+-----+-----+-----+-----+-----+-----+
```

Note that this means that `BIGINT UNSIGNED` is not 100% usable in all contexts. See [Section 12.8, “Cast Functions and Operators”](#). (Added in MySQL 4.0.2)

```
mysql> SET SQL_MODE = '';
mysql> SELECT CAST(0 AS UNSIGNED) - 1;
+-----+
| CAST(0 AS UNSIGNED) - 1 |
+-----+
| 18446744073709551615   |
+-----+
```

```
mysql> SET SQL_MODE = 'NO_UNSIGNED_SUBTRACTION';
mysql> SELECT CAST(0 AS UNSIGNED) - 1;
+-----+
| CAST(0 AS UNSIGNED) - 1 |
+-----+
| -1 |
+-----+
```

- `ONLY_FULL_GROUP_BY`

Do not allow queries for which the `SELECT` list refers to non-aggregated columns that are not named in the `GROUP BY` clause. (Added in MySQL 4.0.0) The following query is invalid with this mode enabled because `address` is not named in the `GROUP BY` clause:

```
SELECT name, address, MAX(age) FROM t GROUP BY name;
```

- `PIPES_AS_CONCAT`

Treat `||` as a string concatenation operator (same as `CONCAT( )`) rather than as a synonym for `OR`. (Added in MySQL 4.0.0)

- `REAL_AS_FLOAT`

Treat `REAL` as a synonym for `FLOAT`. By default, MySQL treats `REAL` as a synonym for `DOUBLE`. (Added in MySQL 4.0.0)

The following special modes are provided as shorthand for combinations of mode values from the preceding list. All are available as of MySQL 4.1.1.

The descriptions include all mode values that are available in the most recent version of MySQL. For older versions, a combination mode does not include individual mode values that are not available except in newer versions.

- `ANSI`

Equivalent to `REAL_AS_FLOAT`, `PIPES_AS_CONCAT`, `ANSI_QUOTES`, `IGNORE_SPACE`. Before MySQL 4.1.11, `ANSI` also includes `ONLY_FULL_GROUP_BY`. See [Section 1.9.3, “Running MySQL in ANSI Mode”](#).

- `DB2`

Equivalent to `PIPES_AS_CONCAT`, `ANSI_QUOTES`, `IGNORE_SPACE`, `NO_KEY_OPTIONS`, `NO_TABLE_OPTIONS`, `NO_FIELD_OPTIONS`.

- `MAXDB`

Equivalent to `PIPES_AS_CONCAT`, `ANSI_QUOTES`, `IGNORE_SPACE`, `NO_KEY_OPTIONS`, `NO_TABLE_OPTIONS`, `NO_FIELD_OPTIONS`.

- `MSSQL`

Equivalent to `PIPES_AS_CONCAT`, `ANSI_QUOTES`, `IGNORE_SPACE`, `NO_KEY_OPTIONS`, `NO_TABLE_OPTIONS`, `NO_FIELD_OPTIONS`.

- `MYSQL323`

Equivalent to `NO_FIELD_OPTIONS`.

- `MYSQL40`

Equivalent to `NO_FIELD_OPTIONS`.

- `ORACLE`

Equivalent to `PIPES_AS_CONCAT`, `ANSI_QUOTES`, `IGNORE_SPACE`, `NO_KEY_OPTIONS`, `NO_TABLE_OPTIONS`, `NO_FIELD_OPTIONS`.

- `POSTGRESQL`

Equivalent to `PIPES_AS_CONCAT`, `ANSI_QUOTES`, `IGNORE_SPACE`, `NO_KEY_OPTIONS`, `NO_TABLE_OPTIONS`, `NO_FIELD_OPTIONS`.

## 5.2.6. The Shutdown Process

The server shutdown process takes place as follows:

1. The shutdown process is initiated.

Server shutdown can be initiated several ways. For example, a user with the `SHUTDOWN` privilege can execute a `mysqladmin shutdown` command. `mysqladmin` can be used on any platform supported by MySQL. Other operating system-specific shutdown initiation methods are possible as well: The server shuts down on Unix when it receives a `SIGTERM` signal. A server running as a service on Windows shuts down when the services manager tells it to. (On Windows, a user with Administrator rights can also shut down the server using `NET STOP service_name`, where `service_name` is the name of the MySQL service. By default, this is `MySQL`.)

2. The server creates a shutdown thread if necessary.

Depending on how shutdown was initiated, the server might create a thread to handle the shutdown process. If shutdown was requested by a client, a shutdown thread is created. If shutdown is the result of receiving a `SIGTERM` signal, the signal thread might handle shutdown itself, or it might create a separate thread to do so. If the server tries to create a shutdown thread and cannot (for example, if memory is exhausted), it issues a diagnostic message that appears in the error log:

```
Error: Can't create thread to kill server
```

3. The server stops accepting new connections.

To prevent new activity from being initiated during shutdown, the server stops accepting new client connections. It does this by closing the network connections to which it normally listens for connections: the TCP/IP port, the Unix socket file, the Windows named pipe, and shared memory on Windows.

4. The server terminates current activity.

For each thread that is associated with a client connection, the connection to the client is broken and the thread is marked as killed. Threads die when they notice that they are so marked. Threads for idle connections die quickly. Threads that currently are processing statements check their state periodically and take longer to die. For additional information about thread termination, see [Section 13.5.5.3, “KILL Syntax”](#), in particular for the instructions about killed `REPAIR TABLE` or `OPTIMIZE TABLE` operations on `MyISAM` tables.

For threads that have an open transaction, the transaction is rolled back. Note that if a thread is updating a non-transactional table, an operation such as a multiple-row `UPDATE` or `INSERT` may

leave the table partially updated, because the operation can terminate before completion.

If the server is a master replication server, threads associated with currently connected slaves are treated like other client threads. That is, each one is marked as killed and exits when it next checks its state.

If the server is a slave replication server, the I/O and SQL threads, if active, are stopped before client threads are marked as killed. The SQL thread is allowed to finish its current statement (to avoid causing replication problems), and then stops. If the SQL thread was in the middle of a transaction at this point, the transaction is rolled back.

5. Storage engines are shut down or closed.

At this stage, the table cache is flushed and all open tables are closed.

Each storage engine performs any actions necessary for tables that it manages. For example, [MyISAM](#) flushes any pending index writes for a table. [InnoDB](#) flushes its buffer pool to disk, writes the current LSN to the tablespace, and terminates its own internal threads.

6. The server exits.

## 5.2.7. Server-Side Help

As of MySQL 4.1, MySQL Server supports a [HELP](#) statement that returns online information from the MySQL Reference manual (see [Section 13.3.2](#), “[HELP Syntax](#)”). The proper operation of this statement requires that the help tables in the `mysql` database be initialized with help topic information, which is done by processing the contents of the `fill_help_tables.sql` script.

For a MySQL binary distribution on Unix, help table setup occurs when you run `mysql_install_db`. For an RPM distribution on Linux or binary distribution on Windows, help table setup occurs as part of the MySQL installation process.

For a MySQL source distribution, you can find the `fill_help_tables.sql` file in the `scripts` directory. To load the file manually, make sure that you have initialized the `mysql` database by running `mysql_install_db`, and then process the file with the `mysql` client as follows:

```
shell> mysql -u root mysql < fill_help_tables.sql
```

If you are working with BitKeeper and a MySQL development source tree, the tree doesn't contain `fill_help_tables.sql`. You can download the proper file for your version of MySQL from <http://dev.mysql.com/doc/>. After downloading and uncompressing the file, process it with `mysql` as just described.

## 5.3. The `mysqld-max` Extended MySQL Server

A MySQL-Max server is a version of the `mysqld` MySQL server that has been built to include additional features. The MySQL-Max distribution to use depends on your platform:

- For Windows, MySQL binary distributions include both the standard server (`mysqld.exe`) and the MySQL-Max server (`mysqld-max.exe`), so no special distribution is needed. Just use a regular Windows distribution. See [Section 2.3](#), “[Installing MySQL on Windows](#)”.
- For Linux, if you install MySQL using RPM distributions, the `MySQL-Max` RPM presupposes that you have already installed the regular server RPM. Use the regular `MySQL-server` RPM first to install a standard server named `mysqld`, and then use the `MySQL-Max` RPM to install a server

named `mysqld-max`. See [Section 2.4, “Installing MySQL on Linux”](#), for more information on the Linux RPM packages.

- All other MySQL-Max distributions contain a single server that is named `mysqld` but that has the additional features included.

You can find the MySQL-Max binaries on the MySQL AB Web site at <http://dev.mysql.com/downloads/>.

MySQL AB builds the MySQL-Max servers by using the following `configure` options:

- `--with-server-suffix=-max`

This option adds a `-max` suffix to the `mysqld` version string.

- `--with-innodb`

This option enables support for the `InnoDB` storage engine. MySQL-Max servers always include `InnoDB` support, but this option actually is needed only for MySQL 3.23. From MySQL 4.0 onward, `InnoDB` is included by default in all binary distributions, so a MySQL-Max server is not needed to obtain `InnoDB` support.

- `--with-bdb`

This option enables support for the Berkeley DB (`BDB`) storage engine on those platforms for which `BDB` is available. (See notes in the following discussion.)

- `--with-blackhole-storage-engine`

This option enables support for the `BLACKHOLE` storage engine in MySQL 4.1.11 and newer.

- `--with-example-storage-engine`

This option enables support for the `EXAMPLE` storage engine in MySQL 4.1.10 and newer.

- `--with-ndbcluster`

As of MySQL 4.1.2, this option enables support for the `NDB Cluster` storage engine on those platforms for which Cluster is available. (See notes in the following discussion.)

- `USE_SYMDIR`

This define is enabled to turn on database symbolic link support for Windows. This applies only before MySQL 4.0. From MySQL 4.0 onward, symbolic link support is enabled for all Windows servers, so a MySQL-Max server is not needed to take advantage of this feature.

MySQL-Max binary distributions are a convenience for those who wish to install precompiled programs. If you build MySQL using a source distribution, you can build your own Max-like server by enabling the same features at configuration time that the MySQL-Max binary distributions are built with.

MySQL-Max servers include the BerkeleyDB (`BDB`) storage engine whenever possible, but not all platforms support `BDB`.

The following table shows on which platforms allow MySQL-Max binaries include support for `BDB` and `NDB Cluster`:

As of MySQL 4.1.2, MySQL Cluster is supported on Linux (on most platforms), Solaris, Mac OS X,

and HP-UX only. Some users have reported success in using MySQL Cluster built from source on BSD operating systems, but these are not officially supported at this time. Note that, even for servers compiled with Cluster support, the `NDB Cluster` storage engine is not enabled by default. You must start the server with the `--ndbcluster` option to use it as part of a MySQL Cluster. (For details, see [Section 15.4, “MySQL Cluster Configuration”](#).)

The following table shows the platforms for which MySQL-Max binaries include support for `BDB` and `NDB Cluster`.

System	BDB Support	NDB Support
AIX 5.2	N	N
HP-UX	Y	Y
Linux-Alpha	N	N
Linux-IA-64	N	Y
Linux-Intel	Y	Y
Mac OS X	N	Y
NetWare	N	N
SCO 6	N	N
Solaris-SPARC	Y	Y
Solaris-Intel	N	Y
Solaris-AMD 64	Y	Y
Windows NT/2000/XP	Y	N

To find out which storage engines your server supports, use the `SHOW ENGINES` statement. (See [Section 13.5.4.8, “SHOW ENGINES Syntax”](#).) For example:

```
mysql> SHOW ENGINES\G
***** 1. row *****
Engine: MyISAM
Support: DEFAULT
Comment: Default engine as of MySQL 3.23 with great performance
***** 2. row *****
Engine: HEAP
Support: YES
Comment: Alias for MEMORY
***** 3. row *****
Engine: MEMORY
Support: YES
Comment: Hash based, stored in memory, useful for temporary tables
***** 4. row *****
Engine: MERGE
Support: YES
Comment: Collection of identical MyISAM tables
...

```

Before MySQL 4.1.2, `SHOW ENGINES` is unavailable. Use the following statement instead and check the value of the variable for the storage engine in which you are interested:

```
mysql> SHOW VARIABLES LIKE 'have%';
+-----+-----+
| Variable_name | Value |
+-----+-----+
| have_archive  | YES   |
| have_bdb      | YES   |
| have_blackhole_engine | YES   |
| have_compress | YES   |
| have_crypt    | YES   |
| have_csv      | YES   |
| have_example_engine | YES   |
| have_geometry | YES   |
+-----+-----+

```

```

have_innodb      YES
have_isam       NO
have_ndbcluster NO
have_openssl    YES
have_query_cache YES
have_raid       NO
have_rtree_keys YES
have_symlink    YES
+-----+-----+
16 rows in set (0.00 sec)

```

The precise output from these statements may vary according to the MySQL version used (and the features that are enabled). The values of the second column of the output indicate the server's level of support for each feature, as shown here:

Value	Meaning
YES	The feature is supported and is active.
NO	The feature is not supported.
DISABLED	The feature is supported but has been disabled.

A value of **NO** means that the server was compiled without support for the feature, so it cannot be activated at runtime.

A value of **DISABLED** occurs either because the server was started with an option that disables the feature, or because not all options required to enable it were given. In the latter case, the error log file should contain a reason indicating why the option is disabled. See [Section 5.11.1, “The Error Log”](#).

One situation in which you might see **DISABLED** occurs with MySQL 3.23 when the **InnoDB** storage engine is compiled in. In MySQL 3.23, you must supply at least the `innodb_data_file_path` option at runtime to set up the **InnoDB** tablespace. Without this option, **InnoDB** disables itself. See [Section 14.2.3, “InnoDB in MySQL 3.23”](#). You can specify configuration options for the **BDB** storage engine, too, but **BDB** does not disable itself if you do not provide them. See [Section 14.5.3, “BDB Startup Options”](#).

You might also see **DISABLED** for a storage engine if the server was compiled to support it, but was started with a `--skip-engine` option. For example, `--skip-innodb` disables the **InnoDB** engine. For the **NDB Cluster** storage engine, **DISABLED** means the server was compiled with support for MySQL Cluster, but was not started with the `--ndb-cluster` option.

As of version 3.23, all MySQL servers support **MyISAM** tables, because **MyISAM** is the default storage engine.

## 5.4. MySQL Server Startup Programs

This section describes several programs that are used to start `mysqld`, the MySQL server.

### 5.4.1. `mysqld_safe` — MySQL Server Startup Script

`mysqld_safe` is the recommended way to start a `mysqld` server on Unix and NetWare. `mysqld_safe` adds some safety features such as restarting the server when an error occurs and logging runtime information to an error log file. NetWare-specific behaviors are listed later in this section.

**Note:** Before MySQL 4.0, `mysqld_safe` is named `safe_mysqld`. To preserve backward compatibility, MySQL binary distributions include `safe_mysqld` as a symbolic link to `mysqld_safe` until MySQL 5.1.

By default, `mysqld_safe` tries to start an executable named `mysqld-max` if it exists, and `mysqld`



otherwise. Be aware of the implications of this behavior:

- On Linux, the `MySQL-Max` RPM relies on this `mysqld_safe` behavior. The RPM installs an executable named `mysqld-max`, which causes `mysqld_safe` to automatically use that executable rather than `mysqld` from that point on.
- If you install a MySQL-Max distribution that includes a server named `mysqld-max`, and then upgrade later to a non-Max version of MySQL, `mysqld_safe` will still attempt to run the old `mysqld-max` server. If you perform such an upgrade, you should manually remove the old `mysqld-max` server to ensure that `mysqld_safe` runs the new `mysqld` server.

To override the default behavior and specify explicitly the name of the server you want to run, specify a `--mysqld` or `--mysqld-version` option to `mysqld_safe`. You can also use `--ledir` to indicate the directory where `mysqld_safe` should look for the server.

Many of the options to `mysqld_safe` are the same as the options to `mysqld`. See [Section 5.2.1, “Command Options”](#).

All options specified to `mysqld_safe` on the command line are passed to `mysqld`. If you want to use any options that are specific to `mysqld_safe` and that `mysqld` does not support, do not specify them on the command line. Instead, list them in the `[mysqld_safe]` group of an option file. See [Section 4.3.2, “Using Option Files”](#).

`mysqld_safe` reads all options from the `[mysqld]`, `[server]`, and `[mysqld_safe]` sections in option files. For backward compatibility, it also reads `[safe_mysqld]` sections, although you should rename such sections to `[mysqld_safe]` when you begin using MySQL 4.0 or later.

`mysqld_safe` supports the following options:

- `--autoclose`

(NetWare only) On NetWare, `mysqld_safe` provides a screen presence. When you unload (shut down) the `mysqld_safe` NLM, the screen does not by default go away. Instead, it prompts for user input:

```
*<NLM has terminated; Press any key to close the screen>*
```

If you want NetWare to close the screen automatically instead, use the `--autoclose` option to `mysqld_safe`.

- `--basedir=path`

The path to the MySQL installation directory.

- `--core-file-size=size`

The size of the core file that `mysqld` should be able to create. The option value is passed to `ulimit -c`.

- `--datadir=path`

The path to the data directory.

- `--defaults-extra-file=path`

The name of an option file to be read in addition to the usual option files. This must be the first option on the command line if it is used.

- `--defaults-file=file_name`

The name of an option file to be read instead of the usual option files. This must be the first option on the command line if it is used.

- `--err-log=file_name`

The old form of the `--log-error` option, to be used before MySQL 4.0.

- `--ledir=path`

If `mysqld_safe` cannot find the server, use this option to indicate the pathname to the directory where the server is located.

- `--log-error=file_name`

Write the error log to the given file. See [Section 5.11.1](#), “The Error Log”.

- `--mysqld=prog_name`

The name of the server program (in the `ledir` directory) that you want to start. This option is needed if you use the MySQL binary distribution but have the data directory outside of the binary distribution. If `mysqld_safe` cannot find the server, use the `--ledir` option to indicate the pathname to the directory where the server is located.

- `--mysqld-version=suffix`

This option is similar to the `--mysqld` option, but you specify only the suffix for the server program name. The basename is assumed to be `mysqld`. For example, if you use `--mysqld-version=max`, `mysqld_safe` starts the `mysqld-max` program in the `ledir` directory. If the argument to `--mysqld-version` is empty, `mysqld_safe` uses `mysqld` in the `ledir` directory.

- `--nice=priority`

Use the `nice` program to set the server's scheduling priority to the given value. This option was added in MySQL 4.0.14.

- `--no-defaults`

Do not read any option files. This must be the first option on the command line if it is used.

- `--open-files-limit=count`

The number of files that `mysqld` should be able to open. The option value is passed to `ulimit -n`. Note that you need to start `mysqld_safe` as `root` for this to work properly.

- `--pid-file=file_name`

The pathname of the process ID file.

- `--port=port_num`

The port number that the server should use when listening for TCP/IP connections. The port number must be 1024 or higher unless the server is started by the `root` system user.

- `--socket=path`

The Unix socket file that the server should use when listening for local connections.

- `--timezone=timezone`

Set the `TZ` time zone environment variable to the given option value. Consult your operating system documentation for legal time zone specification formats.

- `--user={user_name|user_id}`

Run the `mysqld` server as the user having the name `user_name` or the numeric user ID `user_id`. (“User” in this context refers to a system login account, not a MySQL user listed in the grant tables.)

If you execute `mysqld_safe` with the `--defaults-file` or `--defaults-extra-option` option to name an option file, the option must be the first one given on the command line or the option file will not be used. For example, this command will not use the named option file:

```
mysql> mysqld_safe --port=port_num --defaults-file=file_name
```

Instead, use the following command:

```
mysql> mysqld_safe --defaults-file=file_name --port=port_num
```

The `mysqld_safe` script is written so that it normally can start a server that was installed from either a source or a binary distribution of MySQL, even though these types of distributions typically install the server in slightly different locations. (See [Section 2.1.5, “Installation Layouts”](#).) `mysqld_safe` expects one of the following conditions to be true:

- The server and databases can be found relative to the working directory (the directory from which `mysqld_safe` is invoked). For binary distributions, `mysqld_safe` looks under its working directory for `bin` and `data` directories. For source distributions, it looks for `libexec` and `var` directories. This condition should be met if you execute `mysqld_safe` from your MySQL installation directory (for example, `/usr/local/mysql` for a binary distribution).
- If the server and databases cannot be found relative to the working directory, `mysqld_safe` attempts to locate them by absolute pathnames. Typical locations are `/usr/local/libexec` and `/usr/local/var`. The actual locations are determined from the values configured into the distribution at the time it was built. They should be correct if MySQL is installed in the location specified at configuration time.

Because `mysqld_safe` tries to find the server and databases relative to its own working directory, you can install a binary distribution of MySQL anywhere, as long as you run `mysqld_safe` from the MySQL installation directory:

```
shell> cd mysql_installation_directory
shell> bin/mysqld_safe &
```

If `mysqld_safe` fails, even when invoked from the MySQL installation directory, you can specify the `--ledir` and `--datadir` options to indicate the directories in which the server and databases are located on your system.

Normally, you should not edit the `mysqld_safe` script. Instead, configure `mysqld_safe` by using command-line options or options in the `[mysqld_safe]` section of a `my.cnf` option file. In rare cases, it might be necessary to edit `mysqld_safe` to get it to start the server properly. However, if you do this, your modified version of `mysqld_safe` might be overwritten if you upgrade MySQL in the future, so you should make a copy of your edited version that you can reinstall.

On NetWare, `mysqld_safe` is a NetWare Loadable Module (NLM) that is ported from the original Unix shell script. It starts the server as follows:

1. Runs a number of system and option checks.
2. Runs a check on `MyISAM` and `ISAM` tables.
3. Provides a screen presence for the MySQL server.
4. Starts `mysqld`, monitors it, and restarts it if it terminates in error.
5. Sends error messages from `mysqld` to the `host_name.err` file in the data directory.
6. Sends `mysqld_safe` screen output to the `host_name.safe` file in the data directory.

## 5.4.2. `mysql.server` — MySQL Server Startup Script

MySQL distributions on Unix include a script named `mysql.server`. It can be used on systems such as Linux and Solaris that use System V-style run directories to start and stop system services. It is also used by the Mac OS X Startup Item for MySQL.

`mysql.server` can be found in the `support-files` directory under your MySQL installation directory or in a MySQL source distribution.

If you use the Linux server RPM package (`MySQL-server-VERSION.rpm`), the `mysql.server` script will be installed in the `/etc/init.d` directory with the name `mysql`. You need not install it manually. See [Section 2.4, “Installing MySQL on Linux”](#), for more information on the Linux RPM packages.

Some vendors provide RPM packages that install a startup script under a different name such as `mysqld`.

If you install MySQL from a source distribution or using a binary distribution format that does not install `mysql.server` automatically, you can install it manually. Instructions are provided in [Section 2.10.2.2, “Starting and Stopping MySQL Automatically”](#).

`mysql.server` reads options from the `[mysql.server]` and `[mysqld]` sections of option files. For backward compatibility, it also reads `[mysql_server]` sections, although you should rename such sections to `[mysql.server]` when you begin using MySQL 4.0 or later.

## 5.4.3. `mysqld_multi` — Manage Multiple MySQL Servers

`mysqld_multi` is designed to manage several `mysqld` processes that listen for connections on different Unix socket files and TCP/IP ports. It can start or stop servers, or report their current status.

`mysqld_multi` searches for groups named `[mysqldN]` in `my.cnf` (or in the file named by the `--config-file` option). `N` can be any positive integer. This number is referred to in the following discussion as the option group number, or *GNR*. Group numbers distinguish option groups from one another and are used as arguments to `mysqld_multi` to specify which servers you want to start, stop, or obtain a status report for. Options listed in these groups are the same that you would use in the `[mysqld]` group used for starting `mysqld`. (See, for example, [Section 2.10.2.2, “Starting and Stopping MySQL Automatically”](#).) However, when using multiple servers, it is necessary that each one use its own value for options such as the Unix socket file and TCP/IP port number. For more information on which options must be unique per server in a multiple-server environment, see [Section 5.12, “Running Multiple MySQL Servers on the Same Machine”](#).

To invoke `mysqld_multi`, use the following syntax:

```
shell> mysqld_multi [options] {start|stop|report} [GNR[,GNR] ...]
```

`start`, `stop`, and `report` indicate which operation to perform. You can perform the designated operation for a single server or multiple servers, depending on the `GNR` list that follows the option name. If there is no list, `mysqld_multi` performs the operation for all servers in the option file.

Each `GNR` value represents an option group number or range of group numbers. The value should be the number at the end of the group name in the option file. For example, the `GNR` for a group named `[mysqld17]` is 17. To specify a range of numbers, separate the first and last numbers by a dash. The `GNR` value `10-13` represents groups `[mysqld10]` through `[mysqld13]`. Multiple groups or group ranges can be specified on the command line, separated by commas. There must be no whitespace characters (spaces or tabs) in the `GNR` list; anything after a whitespace character is ignored.

This command starts a single server using option group `[mysqld17]`:

```
shell> mysqld_multi start 17
```

This command stops several servers, using option groups `[mysqld8]` and `[mysqld10]` through `[mysqld13]`:

```
shell> mysqld_multi stop 8,10-13
```

For an example of how you might set up an option file, use this command:

```
shell> mysqld_multi --example
```

`mysqld_multi` supports the following options:

- `- -help`

Display a help message and exit.

- `--config-file=file_name`

Specify the name of an alternative option file. This affects where `mysqld_multi` looks for `[mysqldN]` option groups. Without this option, all options are read from the usual `my.cnf` file. The option does not affect where `mysqld_multi` reads its own options, which are always taken from the `[mysqld_multi]` group in the usual `my.cnf` file.

- `--example`

Display a sample option file.

- `--log=file_name`

Specify the name of the log file. If the file exists, log output is appended to it.

- `--mysqladmin=prog_name`

The `mysqladmin` binary to be used to stop servers.

- `--mysqld=prog_name`

The `mysqld` binary to be used. Note that you can specify `mysqld_safe` as the value for this option also. If you use `mysqld_safe` to start the server, you can include the `mysqld` or `ledir` op-

tions in the corresponding `[mysqldN]` option group. These options indicate the name of the server that `mysqld_safe` should start and the pathname of the directory where the server is located. (See the descriptions for these options in [Section 5.4.1](#), “`mysqld_safe` — MySQL Server Startup Script”). Example:

```
[mysqld38]
mysqld = mysqld-max
ledir = /opt/local/mysql/libexec
```

- `--no-log`  
Print log information to `stdout` rather than to the log file. By default, output goes to the log file.
- `--password=password`  
The password of the MySQL account to use when invoking `mysqladmin`. Note that the password value is not optional for this option, unlike for other MySQL programs.
- `--silent`  
Silent mode; disable warnings. This option was added in MySQL 4.1.6.
- `--tcp-ip`  
Connect to each MySQL server via the TCP/IP port instead of the Unix socket file. (If a socket file is missing, the server might still be running, but accessible only via the TCP/IP port.) By default, connections are made using the Unix socket file. This option affects `stop` and `report` operations.
- `--user=user_name`  
The username of the MySQL account to use when invoking `mysqladmin`.
- `--verbose`  
Be more verbose. This option was added in MySQL 4.1.6.
- `--version`  
Display version information and exit.

Some notes about `mysqld_multi`:

- **Most important:** Before using `mysqld_multi` be sure that you understand the meanings of the options that are passed to the `mysqld` servers and *why* you would want to have separate `mysqld` processes. Beware of the dangers of using multiple `mysqld` servers with the same data directory. Use separate data directories, unless you *know* what you are doing. Starting multiple servers with the same data directory does *not* give you extra performance in a threaded system. See [Section 5.12](#), “Running Multiple MySQL Servers on the Same Machine”.
- **Important:** Make sure that the data directory for each server is fully accessible to the Unix account that the specific `mysqld` process is started as. *Do not* use the Unix `root` account for this, unless you *know* what you are doing. See [Section 5.6.5](#), “How to Run MySQL as a Normal User”.
- Make sure that the MySQL account used for stopping the `mysqld` servers (with the `mysqladmin` program) has the same username and password for each server. Also, make sure that the account has the `SHUTDOWN` privilege. If the servers that you want to manage have different usernames or passwords for the administrative accounts, you might want to create an account on each server that has

the same username and password. For example, you might set up a common `multi_admin` account by executing the following commands for each server:

```
shell> mysql -u root -S /tmp/mysql.sock -p
Enter password:
mysql> GRANT SHUTDOWN ON *.*
-> TO 'multi_admin'@'localhost' IDENTIFIED BY 'multipass';
```

See [Section 5.7.2, “How the Privilege System Works”](#). You have to do this for each `mysqld` server. Change the connection parameters appropriately when connecting to each one. Note that the host-name part of the account name must allow you to connect as `multi_admin` from the host where you want to run `mysqld_multi`.

- The Unix socket file and the TCP/IP port number must be different for every `mysqld`.
- The `--pid-file` option is very important if you are using `mysqld_safe` to start `mysqld` (for example, `--mysqld=mysqld_safe`) Every `mysqld` should have its own process ID file. The advantage of using `mysqld_safe` instead of `mysqld` is that `mysqld_safe` monitors its `mysqld` process and restarts it if the process terminates due to a signal sent using `kill -9` or for other reasons, such as a segmentation fault. Please note that the `mysqld_safe` script might require that you start it from a certain place. This means that you might have to change location to a certain directory before running `mysqld_multi`. If you have problems starting, please see the `mysqld_safe` script. Check especially the lines:

```
-----
MY_PWD=`pwd`
# Check if we are starting this relative (for the binary release)
if test -d $MY_PWD/data/mysql -a -f ./share/mysql/english/errmsg.sys -a \
-x ./bin/mysqld
-----
```

The test performed by these lines should be successful, or you might encounter problems. See [Section 5.4.1, “mysqld\\_safe — MySQL Server Startup Script”](#).

- You might want to use the `--user` option for `mysqld`, but to do this you need to run the `mysqld_multi` script as the Unix `root` user. Having the option in the option file does not matter; you merely get a warning if you are not the superuser and the `mysqld` processes are started under your own Unix account.

The following example shows how you might set up an option file for use with `mysqld_multi`. The order in which the `mysqld` programs are started or stopped depends on the order in which they appear in the option file. Group numbers need not form an unbroken sequence. The first and fifth `[mysqldN]` groups were intentionally omitted from the example to illustrate that you can have “gaps” in the option file. This gives you more flexibility.

```
# This file should probably be in your home dir (~/.my.cnf)
# or /etc/my.cnf
# Version 2.1 by Jani Tolonen

[mysqld_multi]
mysqld      = /usr/local/bin/mysqld_safe
mysqldadmin = /usr/local/bin/mysqldadmin
user        = multi_admin
password    = multipass

[mysqld2]
socket      = /tmp/mysql.sock2
port        = 3307
pid-file    = /usr/local/mysql/var2/hostname.pid2
datadir     = /usr/local/mysql/var2
language    = /usr/local/share/mysql/english
user        = john

[mysqld3]
```

```

socket      = /tmp/mysql.sock3
port        = 3308
pid-file    = /usr/local/mysql/var3/hostname.pid3
datadir     = /usr/local/mysql/var3
language    = /usr/local/share/mysql/swedish
user        = monty

[mysqld4]
socket      = /tmp/mysql.sock4
port        = 3309
pid-file    = /usr/local/mysql/var4/hostname.pid4
datadir     = /usr/local/mysql/var4
language    = /usr/local/share/mysql/estonia
user        = tonu

[mysqld6]
socket      = /tmp/mysql.sock6
port        = 3311
pid-file    = /usr/local/mysql/var6/hostname.pid6
datadir     = /usr/local/mysql/var6
language    = /usr/local/share/mysql/japanese
user        = jani

```

See [Section 4.3.2, “Using Option Files”](#).

## 5.5. Installation-Related Programs

### 5.5.1. `mysql_fix_privilege_tables` — Upgrade MySQL System Tables

Some releases of MySQL introduce changes to the structure of the system tables in the `mysql` database to add new privileges or support new features. When you update to a new version of MySQL, you should update your system tables as well to make sure that their structure is up to date. Otherwise, there might be capabilities that you cannot take advantage of. First, make a backup of your `mysql` database, and then use the following procedure.

On Unix or Unix-like systems, update the system tables by running the `mysql_fix_privilege_tables` script:

```
shell> mysql_fix_privilege_tables
```

You must run this script while the server is running. It attempts to connect to the server running on the local host as `root`. If your `root` account requires a password, indicate the password on the command line. For MySQL 4.1 and up, specify the password like this:

```
shell> mysql_fix_privilege_tables --password=root_password
```

Prior to MySQL 4.1, specify the password like this:

```
shell> mysql_fix_privilege_tables root_password
```

The `mysql_fix_privilege_tables` script performs any actions necessary to convert your system tables to the current format. You might see some `Duplicate column name` warnings as it runs; you can ignore them.

After running the script, stop the server and restart it.

On Windows systems, there isn't an easy way to update the system tables until MySQL 4.0.15. From version 4.0.15 on, MySQL distributions include a `mysql_fix_privilege_tables.sql` SQL script that you can run using the `mysql` client. For example, if your MySQL installation is located at `C:\Program Files\MySQL\MySQL Server 4.1`, the commands look like this:



```
C:\> cd "C:\Program Files\MySQL\MySQL Server 4.1"
C:\> bin\mysql -u root -p mysql
mysql> SOURCE scripts/mysql_fix_privilege_tables.sql
```

The `mysql` command will prompt you for the `root` password; enter it when prompted.

If your installation is located in some other directory, adjust the pathnames appropriately.

As with the Unix procedure, you might see some `Duplicate column name` warnings as `mysql` processes the statements in the `mysql_fix_privilege_tables.sql` script; you can ignore them.

After running the script, stop the server and restart it.

## 5.6. General Security Issues

This section describes some general security issues to be aware of and what you can do to make your MySQL installation more secure against attack or misuse. For information specifically about the access control system that MySQL uses for setting up user accounts and checking database access, see [Section 5.7, “The MySQL Access Privilege System”](#).

### 5.6.1. General Security Guidelines

Anyone using MySQL on a computer connected to the Internet should read this section to avoid the most common security mistakes.

In discussing security, we emphasize the necessity of fully protecting the entire server host (not just the MySQL server) against all types of applicable attacks: eavesdropping, altering, playback, and denial of service. We do not cover all aspects of availability and fault tolerance here.

MySQL uses security based on Access Control Lists (ACLs) for all connections, queries, and other operations that users can attempt to perform. There is also support for SSL-encrypted connections between MySQL clients and servers. Many of the concepts discussed here are not specific to MySQL at all; the same general ideas apply to almost all applications.

When running MySQL, follow these guidelines whenever possible:

- **Do not ever give anyone (except MySQL `root` accounts) access to the `user` table in the `mysql` database!** This is critical, particularly before MySQL 4.1, when *the encrypted password is the real password in MySQL*: Anyone who knows the password that is listed in the `mysql.user` table and who has access to the host listed for the account *can easily log in as that user*. In MySQL 4.1, the password hashing algorithm was changed so that this is no longer true.
- Learn the MySQL access privilege system. The `GRANT` and `REVOKE` statements are used for controlling access to MySQL. Do not grant more privileges than necessary. Never grant privileges to all hosts.

Checklist:

- Try `mysql -u root`. If you are able to connect successfully to the server without being asked for a password, anyone can connect to your MySQL server as the MySQL `root` user with full privileges! Review the MySQL installation instructions, paying particular attention to the information about setting a `root` password. See [Section 2.10.3, “Securing the Initial MySQL Accounts”](#).
- Use the `SHOW GRANTS` statement to check which accounts have access to what. Then use the `REVOKE` statement to remove those privileges that are not necessary.

- Do not store any plain-text passwords in your database. If your computer becomes compromised, the intruder can take the full list of passwords and use them. Instead, use `MD5()`, `SHA1()`, or some other one-way hashing function and store the hash value.
- Do not choose passwords from dictionaries. Special programs exist to break passwords. Even passwords like “xfish98” are very bad. Much better is “duag98” which contains the same word “fish” but typed one key to the left on a standard QWERTY keyboard. Another method is to use a password that is taken from the first characters of each word in a sentence (for example, “Mary had a little lamb” results in a password of “Mhall”). The password is easy to remember and type, but difficult to guess for someone who does not know the sentence.
- Invest in a firewall. This protects you from at least 50% of all types of exploits in any software. Put MySQL behind the firewall or in a demilitarized zone (DMZ).

Checklist:

- Try to scan your ports from the Internet using a tool such as `nmap`. MySQL uses port 3306 by default. This port should not be accessible from untrusted hosts. Another simple way to check whether or not your MySQL port is open is to try the following command from some remote machine, where `server_host` is the hostname or IP number of the host on which your MySQL server runs:

```
shell> telnet server_host 3306
```

If you get a connection and some garbage characters, the port is open, and should be closed on your firewall or router, unless you really have a good reason to keep it open. If `telnet` hangs or the connection is refused, the port is blocked, which is how you want it to be.

- Do not trust any data entered by users of your applications. They can try to trick your code by entering special or escaped character sequences in Web forms, URLs, or whatever application you have built. Be sure that your application remains secure if a user enters something like “`;  
DROP DATABASE mysql;`”. This is an extreme example, but large security leaks and data loss might occur as a result of hackers using similar techniques, if you do not prepare for them.

A common mistake is to protect only string data values. Remember to check numeric data as well. If an application generates a query such as `SELECT * FROM table WHERE ID=234` when a user enters the value `234`, the user can enter the value `234 OR 1=1` to cause the application to generate the query `SELECT * FROM table WHERE ID=234 OR 1=1`. As a result, the server retrieves every row in the table. This exposes every row and causes excessive server load. The simplest way to protect from this type of attack is to use single quotes around the numeric constants: `SELECT * FROM table WHERE ID='234'`. If the user enters extra information, it all becomes part of the string. In a numeric context, MySQL automatically converts this string to a number and strips any trailing non-numeric characters from it.

Sometimes people think that if a database contains only publicly available data, it need not be protected. This is incorrect. Even if it is allowable to display any row in the database, you should still protect against denial of service attacks (for example, those that are based on the technique in the preceding paragraph that causes the server to waste resources). Otherwise, your server becomes unresponsive to legitimate users.

Checklist:

- Try to enter single and double quote marks (‘`'`’ and ‘`"`’) in all of your Web forms. If you get any kind of MySQL error, investigate the problem right away.
- Try to modify dynamic URLs by adding `%22` (‘`"`’), `%23` (‘`#`’), and `%27` (‘`'`’) to them.
- Try to modify data types in dynamic URLs from numeric to character types using the characters

shown in the previous examples. Your application should be safe against these and similar attacks.

- Try to enter characters, spaces, and special symbols rather than numbers in numeric fields. Your application should remove them before passing them to MySQL or else generate an error. Passing unchecked values to MySQL is very dangerous!
- Check the size of data before passing it to MySQL.
- Have your application connect to the database using a username different from the one you use for administrative purposes. Do not give your applications any access privileges they do not need.
- Many application programming interfaces provide a means of escaping special characters in data values. Properly used, this prevents application users from entering values that cause the application to generate statements that have a different effect than you intend:
  - MySQL C API: Use the `mysql_real_escape_string()` API call.
  - MySQL++: Use the `escape` and `quote` modifiers for query streams.
  - PHP: Use the `mysql_real_escape_string()` function (available as of PHP 4.3.0, prior to that PHP version use `mysql_escape_string()`, and prior to PHP 4.0.3, use `addslashes()`). Note that only `mysql_real_escape_string()` is character set-aware; the other functions can be “bypassed” when using (invalid) multi-byte character sets. In PHP 5 (and as of MySQL 4.1), you can use the `mysqli` extension, which supports the improved MySQL authentication protocol and passwords, as well as prepared statements with placeholders.
  - Perl DBI: Use placeholders or the `quote()` method.
  - Ruby DBI: Use placeholders or the `quote()` method.
  - Java JDBC: Use a `PreparedStatement` object and placeholders.

Other programming interfaces might have similar capabilities.

- Do not transmit plain (unencrypted) data over the Internet. This information is accessible to everyone who has the time and ability to intercept it and use it for their own purposes. Instead, use an encrypted protocol such as SSL or SSH. MySQL supports internal SSL connections as of version 4.0. Another technique is to use SSH port-forwarding to create an encrypted (and compressed) tunnel for the communication.
- Learn to use the `tcpdump` and `strings` utilities. In most cases, you can check whether MySQL data streams are unencrypted by issuing a command like the following:

```
shell> tcpdump -l -i eth0 -w - - src or dst port 3306 | strings
```

(This works under Linux and should work with small modifications under other systems.) **Warning:** If you do not see plaintext data, this does not always mean that the information actually is encrypted. If you need high security, you should consult with a security expert.

## 5.6.2. Making MySQL Secure Against Attackers

When you connect to a MySQL server, you should use a password. The password is not transmitted in clear text over the connection. Password handling during the client connection sequence was upgraded in MySQL 4.1.1 to be very secure. If you are still using pre-4.1.1-style passwords, the encryption al-

gorithm is not as strong as the newer algorithm. With some effort, a clever attacker who can sniff the traffic between the client and the server can crack the password. (See [Section 5.7.9, “Password Hashing as of MySQL 4.1”](#), for a discussion of the different password handling methods.)

All other information is transferred as text, and can be read by anyone who is able to watch the connection. If the connection between the client and the server goes through an untrusted network, and you are concerned about this, you can use the compressed protocol (in MySQL 3.22 and above) to make traffic much more difficult to decipher. You can also use MySQL's internal SSL support to make the connection even more secure in MySQL 4.0 and up. See [Section 5.8.7, “Using Secure Connections”](#). Alternatively, use SSH to get an encrypted TCP/IP connection between a MySQL server and a MySQL client. You can find an Open Source SSH client at <http://www.openssh.org/>, and a commercial SSH client at <http://www.ssh.com/>.

To make a MySQL system secure, you should strongly consider the following suggestions:

- Require all MySQL accounts to have a password. A client program does not necessarily know the identity of the person running it. It is common for client/server applications that the user can specify any username to the client program. For example, anyone can use the `mysql` program to connect as any other person simply by invoking it as `mysql -u other_user db_name` if `other_user` has no password. If all accounts have a password, connecting using another user's account becomes much more difficult.

For a discussion of methods for setting passwords, see [Section 5.8.5, “Assigning Account Passwords”](#).

- Never run the MySQL server as the Unix `root` user. This is extremely dangerous, because any user with the `FILE` privilege is able to cause the server to create files as `root` (for example, `~root/.bashrc`). To prevent this, `mysqld` refuses to run as `root` unless that is specified explicitly using the `--user=root` option.

`mysqld` can (and should) be run as an ordinary, unprivileged user instead. You can create a separate Unix account named `mysql` to make everything even more secure. Use this account only for administering MySQL. To start `mysqld` as a different Unix user, add a `user` option that specifies the username in the `[mysqld]` group of the `my.cnf` option file where you specify server options. For example:

```
[mysqld]
user=mysql
```

This causes the server to start as the designated user whether you start it manually or by using `mysqld_safe` or `mysql.server`. For more details, see [Section 5.6.5, “How to Run MySQL as a Normal User”](#).

Running `mysqld` as a Unix user other than `root` does not mean that you need to change the `root` username in the `user` table. *Usernames for MySQL accounts have nothing to do with usernames for Unix accounts.*

- Do not allow the use of symlinks to tables. (This capability can be disabled with the `--skip-symbolic-links` option.) This is especially important if you run `mysqld` as `root`, because anyone that has write access to the server's data directory then could delete any file in the system! See [Section 7.6.1.2, “Using Symbolic Links for Tables on Unix”](#).
- Make sure that the only Unix user with read or write privileges in the database directories is the user that `mysqld` runs as.
- Do not grant the `PROCESS` or `SUPER` privilege to non-administrative users. The output of `mysql-admin processlist` and `SHOW PROCESSLIST` shows the text of any statements currently be-

ing executed, so any user who is allowed to see the server process list might be able to see statements issued by other users such as `UPDATE user SET password=PASSWORD('not_secure')`.

`mysqld` reserves an extra connection for users who have the `SUPER` privilege (`PROCESS` before MySQL 4.0.2), so that a MySQL `root` user can log in and check server activity even if all normal connections are in use.

The `SUPER` privilege can be used to terminate client connections, change server operation by changing the value of system variables, and control replication servers.

- Do not grant the `FILE` privilege to non-administrative users. Any user that has this privilege can write a file anywhere in the filesystem with the privileges of the `mysqld` daemon. To make this a bit safer, files generated with `SELECT ... INTO OUTFILE` do not overwrite existing files and are writable by everyone.

The `FILE` privilege may also be used to read any file that is world-readable or accessible to the Unix user that the server runs as. With this privilege, you can read any file into a database table. This could be abused, for example, by using `LOAD DATA` to load `/etc/passwd` into a table, which then can be displayed with `SELECT`.

- If you do not trust your DNS, you should use IP numbers rather than hostnames in the grant tables. In any case, you should be very careful about creating grant table entries using hostname values that contain wildcards.
- If you want to restrict the number of connections allowed to a single account, you can do so by setting the `max_user_connections` variable in `mysqld`. The `GRANT` statement also supports resource control options for limiting the extent of server use allowed to an account. See [Section 13.5.1.2, “GRANT Syntax”](#).

### 5.6.3. Security-Related `mysqld` Options

The following `mysqld` options affect security:

- `--allow-suspicious-udfs`

This option controls whether user-defined functions that have only an `xxx` symbol for the main function can be loaded. By default, the option is turned off and only UDFs that have at least one auxiliary symbol can be loaded; this prevents attempts at loading functions from shared object files other than those containing legitimate UDFs. This option was added in MySQL 4.0.24 and 4.1.10a. See [Section 19.2.4.6, “User-Defined Function Security Precautions”](#).

- `--local-infile[={0|1}]`

If you start the server with `--local-infile=0`, clients cannot use `LOCAL` in `LOAD DATA` statements. See [Section 5.6.4, “Security Issues with LOAD DATA LOCAL”](#).

- `--old-passwords`

Force the server to generate short (pre-4.1) password hashes for new passwords. This is useful for compatibility when the server must support older client programs. See [Section 5.7.9, “Password Hashing as of MySQL 4.1”](#).

- `--safe-show-database`

With this option, the `SHOW DATABASES` statement displays the names of only those databases for

which the user has some kind of privilege. As of MySQL 4.0.2, this option is deprecated and does not do anything (it is enabled by default), because there is a `SHOW DATABASES` privilege that can be used to control access to database names on a per-account basis. See [Section 13.5.1.2, “GRANT Syntax”](#).

- `--safe-user-create`

If this option is enabled, a user cannot create new MySQL users by using the `GRANT` statement unless the user has the `INSERT` privilege for the `mysql.user` table. If you want a user to have the ability to create new users that have those privileges that the user has right to grant, you should grant the user the following privilege:

```
GRANT INSERT(user) ON mysql.user TO 'user_name'@'host_name';
```

This ensures that the user cannot change any privilege columns directly, but has to use the `GRANT` statement to give privileges to other users.

- `--secure-auth`

Disallow authentication for accounts that have old (pre-4.1) passwords. This option is available as of MySQL 4.1.1.

The `mysql` client also has a `--secure-auth` option, which prevents connections to a server if the server requires a password in old format for the client account.

- `--skip-grant-tables`

This option causes the server not to use the privilege system at all. This gives anyone with access to the server *unrestricted access to all databases*. You can cause a running server to start using the grant tables again by executing `mysqladmin flush-privileges` or `mysqladmin reload` command from a system shell, or by issuing a MySQL `FLUSH PRIVILEGES` statement. This option also suppresses loading of user-defined functions (UDFs).

- `--skip-name-resolve`

Hostnames are not resolved. All `Host` column values in the grant tables must be IP numbers or `localhost`.

- `--skip-networking`

Do not allow TCP/IP connections over the network. All connections to `mysqld` must be made via Unix socket files. This option is unsuitable when using a MySQL version prior to 3.23.27 with the MIT-pthreads package, because Unix socket files were not supported by MIT-pthreads at that time.

- `--skip-show-database`

With this option, the `SHOW DATABASES` statement is allowed only to users who have the `SHOW DATABASES` privilege, and the statement displays all database names. Without this option, `SHOW DATABASES` is allowed to all users, but displays each database name only if the user has the `SHOW DATABASES` privilege or some privilege for the database. Note that any global privilege is a privilege for the database.

## 5.6.4. Security Issues with `LOAD DATA LOCAL`

The `LOAD DATA` statement can load a file that is located on the server host, or it can load a file that is located on the client host when the `LOCAL` keyword is specified.

There are two potential security issues with supporting the `LOCAL` version of `LOAD DATA` statements:

- The transfer of the file from the client host to the server host is initiated by the MySQL server. In theory, a patched server could be built that would tell the client program to transfer a file of the server's choosing rather than the file named by the client in the `LOAD DATA` statement. Such a server could access any file on the client host to which the client user has read access.
- In a Web environment where the clients are connecting from a Web server, a user could use `LOAD DATA LOCAL` to read any files that the Web server process has read access to (assuming that a user could run any command against the SQL server). In this environment, the client with respect to the MySQL server actually is the Web server, not the remote program being run by the user who connects to the Web server.

To deal with these problems, we changed how `LOAD DATA LOCAL` is handled as of MySQL 3.23.49 and MySQL 4.0.2 (4.0.13 on Windows):

- By default, all MySQL clients and libraries in binary distributions are compiled with the `-enable-local-infile` option, to be compatible with MySQL 3.23.48 and before.
- If you build MySQL from source but do not invoke `configure` with the `-enable-local-infile` option, `LOAD DATA LOCAL` cannot be used by any client unless it is written explicitly to invoke `mysql_options(... MYSQL_OPT_LOCAL_INFILE, 0)`. See [Section 17.2.3.47](#), “`mysql_options()`”.
- You can disable all `LOAD DATA LOCAL` commands from the server side by starting `mysqld` with the `--local-infile=0` option.
- For the `mysql` command-line client, `LOAD DATA LOCAL` can be enabled by specifying the `--local-infile[=1]` option, or disabled with the `--local-infile=0` option. Similarly, for `mysqlimport`, the `--local` or `-L` option enables local data file loading. In any case, successful use of a local loading operation requires that the server is enabled to allow it.
- If you use `LOAD DATA LOCAL` in Perl scripts or other programs that read the `[client]` group from option files, you can add the `local-infile=1` option to that group. However, to keep this from causing problems for programs that do not understand `local-infile`, specify it using the `loose-` prefix:

```
[client]
loose-local-infile=1
```

The `loose-` prefix can be used as of MySQL 4.0.2.

- If `LOAD DATA LOCAL INFILE` is disabled, either in the server or the client, a client that attempts to issue such a statement receives the following error message:

```
ERROR 1148: The used command is not allowed with this MySQL version
```

## 5.6.5. How to Run MySQL as a Normal User

On Windows, you can run the server as a Windows service using a normal user account beginning with MySQL 4.0.17 and 4.1.2. (Older MySQL versions required you to have administrator rights. This was a bug introduced in MySQL 3.23.54.)

On Unix, the MySQL server `mysqld` can be started and run by any user. However, you should avoid

running the server as the Unix `root` user for security reasons. To change `mysqld` to run as a normal unprivileged Unix user `user_name`, you must do the following:

1. Stop the server if it's running (use `mysqldadmin shutdown`).
2. Change the database directories and files so that `user_name` has privileges to read and write files in them (you might need to do this as the Unix `root` user):

```
shell> chown -R user_name /path/to/mysql/datadir
```

If you do not do this, the server will not be able to access databases or tables when it runs as `user_name`.

If directories or files within the MySQL data directory are symbolic links, you'll also need to follow those links and change the directories and files they point to. `chown -R` might not follow symbolic links for you.

3. Start the server as user `user_name`. If you are using MySQL 3.22 or later, another alternative is to start `mysqld` as the Unix `root` user and use the `--user=user_name` option. `mysqld` starts up, then switches to run as the Unix user `user_name` before accepting any connections.
4. To start the server as the given user automatically at system startup time, specify the username by adding a `user` option to the `[mysqld]` group of the `/etc/my.cnf` option file or the `my.cnf` option file in the server's data directory. For example:

```
[mysqld]
user=user_name
```

If your Unix machine itself isn't secured, you should assign passwords to the MySQL `root` accounts in the grant tables. Otherwise, any user with a login account on that machine can run the `mysql` client with a `--user=root` option and perform any operation. (It is a good idea to assign passwords to MySQL accounts in any case, but especially so when other login accounts exist on the server host.) See [Section 2.10, "Post-Installation Setup and Testing"](#).

## 5.7. The MySQL Access Privilege System

MySQL has an advanced but non-standard security and privilege system. The following discussion describes how it works.

### 5.7.1. What the Privilege System Does

The primary function of the MySQL privilege system is to authenticate a user who connects from a given host and to associate that user with privileges on a database such as `SELECT`, `INSERT`, `UPDATE`, and `DELETE`.

Additional functionality includes the ability to have anonymous users and to grant privileges for MySQL-specific functions such as `LOAD DATA INFILE` and administrative operations.

### 5.7.2. How the Privilege System Works

The MySQL privilege system ensures that all users may perform only the operations allowed to them. As a user, when you connect to a MySQL server, your identity is determined by *the host from which you connect* and *the username you specify*. When you issue requests after connecting, the system grants privileges according to your identity and *what you want to do*.



MySQL considers both your hostname and username in identifying you because there is little reason to assume that a given username belongs to the same person everywhere on the Internet. For example, the user `joe` who connects from `office.example.com` need not be the same person as the user `joe` who connects from `home.example.com`. MySQL handles this by allowing you to distinguish users on different hosts that happen to have the same name: You can grant one set of privileges for connections by `joe` from `office.example.com`, and a different set of privileges for connections by `joe` from `home.example.com`.

MySQL access control involves two stages when you run a client program that connects to the server:

- Stage 1: The server checks whether it should allow you to connect.
- Stage 2: Assuming that you can connect, the server checks each statement you issue to determine whether you have sufficient privileges to perform it. For example, if you try to select rows from a table in a database or drop a table from the database, the server verifies that you have the `SELECT` privilege for the table or the `DROP` privilege for the database.

If your privileges are changed (either by yourself or someone else) while you are connected, those changes do not necessarily take effect immediately for the next statement that you issue. See [Section 5.7.7, “When Privilege Changes Take Effect”](#), for details.

The server stores privilege information in the grant tables of the `mysql` database (that is, in the database named `mysql`). The MySQL server reads the contents of these tables into memory when it starts and re-reads them under the circumstances indicated in [Section 5.7.7, “When Privilege Changes Take Effect”](#). Access-control decisions are based on the in-memory copies of the grant tables.

Normally, you manipulate the contents of the grant tables indirectly by using statements such as `GRANT` and `REVOKE` to set up accounts and control the privileges available to each one. See [Section 13.5.1, “Account Management Statements”](#). The discussion here describes the underlying structure of the grant tables and how the server uses their contents when interacting with clients.

The server uses the `user`, `db`, and `host` tables in the `mysql` database at both stages of access control. The columns in the `user` and `db` tables are shown here. The `host` table is similar to the `db` table but has a specialized use as described in [Section 5.7.6, “Access Control, Stage 2: Request Verification”](#).

Table Name	<code>user</code>	<code>db</code>
<b>Scope columns</b>	<code>Host</code>	<code>Host</code>
	<code>User</code>	<code>Db</code>
	<code>Password</code>	<code>User</code>
<b>Privilege columns</b>	<code>Select_priv</code>	<code>Select_priv</code>
	<code>Insert_priv</code>	<code>Insert_priv</code>
	<code>Update_priv</code>	<code>Update_priv</code>
	<code>Delete_priv</code>	<code>Delete_priv</code>
	<code>Index_priv</code>	<code>Index_priv</code>
	<code>Alter_priv</code>	<code>Alter_priv</code>
	<code>Create_priv</code>	<code>Create_priv</code>
	<code>Drop_priv</code>	<code>Drop_priv</code>
	<code>Grant_priv</code>	<code>Grant_priv</code>
	<code>References_priv</code>	<code>References_priv</code>
	<code>Execute_priv</code>	
<code>Reload_priv</code>		

	Shutdown_priv	
	Process_priv	
	File_priv	
	Show_db_priv	
	Super_priv	
	Create_tmp_table_priv	Create_tmp_table_priv
	Lock_tables_priv	Lock_tables_priv
	Repl_slave_priv	
	Repl_client_priv	
<b>Security columns</b>	ssl_type	
	ssl_cipher	
	x509_issuer	
	x509_subject	
<b>Resource control columns</b>	max_questions	
	max_updates	
	max_connections	
	max_user_connections	

The `ssl_type`, `ssl_cipher`, `x509_issuer`, and `x509_subject` columns were added in MySQL 4.0.0.

The `Create_tmp_table_priv`, `Execute_priv`, `Lock_tables_priv`, `Repl_client_priv`, `Repl_slave_priv`, `Show_db_priv`, `Super_priv`, `max_questions`, `max_updates`, and `max_connections` columns were added in MySQL 4.0.2. `Execute_priv` is not operational through MySQL 4.1.

During the second stage of access control, the server performs request verification to make sure that each client has sufficient privileges for each request that it issues. In addition to the `user`, `db`, and `host` grant tables, the server may also consult the `tables_priv` and `columns_priv` tables for requests that involve tables. The `tables_priv` and `columns_priv` tables provide finer privilege control at the table and column levels. They have the following columns:

Table Name	tables_priv	columns_priv
<b>Scope columns</b>	Host	Host
	Db	Db
	User	User
	Table_name	Table_name
		Column_name
<b>Privilege columns</b>	Table_priv	Column_priv
	Column_priv	
<b>Other columns</b>	Timestamp	Timestamp
	Grantor	

The `Timestamp` and `Grantor` columns currently are unused and are discussed no further here.

Each grant table contains scope columns and privilege columns:

- Scope columns determine the scope of each row (entry) in the tables; that is, the context in which the row applies. For example, a `user` table row with `Host` and `User` values of `'thomas.loc.gov'` and `'bob'` would be used for authenticating connections made to the server from the host `thomas.loc.gov` by a client that specifies a username of `bob`. Similarly, a `db` table row with `Host`, `User`, and `Db` column values of `'thomas.loc.gov'`, `'bob'` and `'reports'` would be used when `bob` connects from the host `thomas.loc.gov` to access the `reports` database. The `tables_priv` and `columns_priv` tables contain scope columns indicating tables or table/column combinations to which each row applies. The `procs_priv` scope columns indicate the stored routine to which each row applies.
- Privilege columns indicate which privileges are granted by a table row; that is, what operations can be performed. The server combines the information in the various grant tables to form a complete description of a user's privileges. [Section 5.7.6, “Access Control, Stage 2: Request Verification”](#), describes the rules that are used to do this.

Scope columns contain strings. They are declared as shown here; the default value for each is the empty string:

Column Name	Type
<code>Host</code>	<code>CHAR(60)</code>
<code>User</code>	<code>CHAR(16)</code>
<code>Password</code>	<code>CHAR(16)</code>
<code>Db</code>	<code>CHAR(64)</code>
<code>Table_name</code>	<code>CHAR(64)</code>
<code>Column_name</code>	<code>CHAR(64)</code>
<code>Routine_name</code>	<code>CHAR(64)</code>

Before MySQL 3.23, the `Db` column is `CHAR(32)` in some tables and `CHAR(60)` in others.

For access-checking purposes, comparisons of `Host` values are case-insensitive. `User`, `Password`, `Db`, and `Table_name` values are case sensitive. `Column_name` values are case insensitive in MySQL 3.22.12 or later.

In the `user`, `db`, and `host` tables, each privilege is listed in a separate column that is declared as `ENUM('N','Y') DEFAULT 'N'`. In other words, each privilege can be disabled or enabled, with the default being disabled.

In the `tables_priv`, `columns_priv`, and `procs_priv` tables, the privilege columns are declared as `SET` columns. Values in these columns can contain any combination of the privileges controlled by the table:

Table Name	Column Name	Possible Set Elements
<code>tables_priv</code>	<code>Table_priv</code>	<code>'Select', 'Insert', 'Update', 'Delete', 'Create', 'Drop', 'Grant', 'References', 'Index', 'Alter'</code>
<code>tables_priv</code>	<code>Column_priv</code>	<code>'Select', 'Insert', 'Update', 'References'</code>
<code>columns_priv</code>	<code>Column_priv</code>	<code>'Select', 'Insert', 'Update', 'References'</code>

Briefly, the server uses the grant tables in the following manner:

- The `user` table scope columns determine whether to reject or allow incoming connections. For allowed connections, any privileges granted in the `user` table indicate the user's global (superuser) privileges. Any privilege granted in this table applies to *all* databases on the server.

**Note:** Because any global privilege is considered a privilege for all databases, any global privilege enables a user to see all database names with `SHOW DATABASES`.

- The `db` table scope columns determine which users can access which databases from which hosts. The privilege columns determine which operations are allowed. A privilege granted at the database level applies to the database and to all its tables.
- The `host` table is used in conjunction with the `db` table when you want a given `db` table row to apply to several hosts. For example, if you want a user to be able to use a database from several hosts in your network, leave the `Host` value empty in the user's `db` table row, then populate the `host` table with a row for each of those hosts. This mechanism is described more detail in [Section 5.7.6, “Access Control, Stage 2: Request Verification”](#).

**Note:** The `host` table must be modified directly with statements such as `INSERT`, `UPDATE`, and `DELETE`. It is not affected by statements such as `GRANT` and `REVOKE` that modify the grant tables indirectly. Most MySQL installations need not use this table at all.

- The `tables_priv` and `columns_priv` tables are similar to the `db` table, but are more fine-grained: They apply at the table and column levels rather than at the database level. A privilege granted at the table level applies to the table and to all its columns. A privilege granted at the column level applies only to a specific column.

Administrative privileges (such as `RELOAD` or `SHUTDOWN`) are specified only in the `user` table. The reason for this is that administrative operations are operations on the server itself and are not database-specific, so there is no reason to list these privileges in the other grant tables. In fact, to determine whether you can perform an administrative operation, the server need consult only the `user` table.

The `FILE` privilege also is specified only in the `user` table. It is not an administrative privilege as such, but your ability to read or write files on the server host is independent of the database you are accessing.

The `mysqld` server reads the contents of the grant tables into memory when it starts. You can tell it to re-read the tables by issuing a `FLUSH PRIVILEGES` statement or executing a `mysqladmin flush-privileges` or `mysqladmin reload` command. Changes to the grant tables take effect as indicated in [Section 5.7.7, “When Privilege Changes Take Effect”](#).

When you modify the contents of the grant tables, it is a good idea to make sure that your changes set up privileges the way you want. To check the privileges for a given account, use the `SHOW GRANTS` statement. (See [Section 13.5.4.10, “SHOW GRANTS Syntax”](#).) For example, to determine the privileges that are granted to an account with `Host` and `User` values of `pc84.example.com` and `bob`, issue this statement:

```
SHOW GRANTS FOR 'bob'@'pc84.example.com';
```

For additional help in diagnosing privilege-related problems, see [Section 5.7.8, “Causes of Access denied Errors”](#). For general advice on security issues, see [Section 5.6, “General Security Issues”](#).

### 5.7.3. Privileges Provided by MySQL

Information about account privileges is stored in the `user`, `db`, `host`, `tables_priv`, and `columns_priv` tables in the `mysql` database. The MySQL server reads the contents of these tables

into memory when it starts and re-reads them under the circumstances indicated in [Section 5.7.7, “When Privilege Changes Take Effect”](#). Access-control decisions are based on the in-memory copies of the grant tables.

The names used in the `GRANT` and `REVOKE` statements to refer to privileges are shown in the following table, along with the column name associated with each privilege in the grant tables and the context in which the privilege applies. Further information about the meaning of each privilege may be found at [Section 13.5.1.2, “GRANT Syntax”](#).

Privilege	Column	Context
<code>CREATE</code>	<code>Create_priv</code>	databases, tables, or indexes
<code>DROP</code>	<code>Drop_priv</code>	databases or tables
<code>GRANT OPTION</code>	<code>Grant_priv</code>	databases, tables, or stored routines
<code>REFERENCES</code>	<code>References_priv</code>	databases or tables
<code>ALTER</code>	<code>Alter_priv</code>	tables
<code>DELETE</code>	<code>Delete_priv</code>	tables
<code>INDEX</code>	<code>Index_priv</code>	tables
<code>INSERT</code>	<code>Insert_priv</code>	tables
<code>SELECT</code>	<code>Select_priv</code>	tables
<code>UPDATE</code>	<code>Update_priv</code>	tables
<code>FILE</code>	<code>File_priv</code>	file access on server host
<code>CREATE TEMPORARY TABLES</code>	<code>Create_tmp_table_priv</code>	server administration
<code>LOCK TABLES</code>	<code>Lock_tables_priv</code>	server administration
<code>PROCESS</code>	<code>Process_priv</code>	server administration
<code>RELOAD</code>	<code>Reload_priv</code>	server administration
<code>REPLICATION CLIENT</code>	<code>Repl_client_priv</code>	server administration
<code>REPLICATION SLAVE</code>	<code>Repl_slave_priv</code>	server administration
<code>SHOW DATABASES</code>	<code>Show_db_priv</code>	server administration
<code>SHUTDOWN</code>	<code>Shutdown_priv</code>	server administration
<code>SUPER</code>	<code>Super_priv</code>	server administration

Some releases of MySQL introduce changes to the structure of the grant tables to add new privileges or features. Whenever you update to a new version of MySQL, you should update your grant tables to make sure that they have the current structure so that you can take advantage of any new capabilities. See [Section 5.5.1, “mysql\\_fix\\_privilege\\_tables — Upgrade MySQL System Tables”](#).

The `CREATE TEMPORARY TABLES`, `EXECUTE`, `LOCK TABLES`, `REPLICATION CLIENT`, `REPLICATION SLAVE`, `SHOW DATABASES`, and `SUPER` privileges were added in MySQL 4.0.2. (`EXECUTE` is not used in any MySQL version through the 4.1 release series.)

The `CREATE` and `DROP` privileges allow you to create new databases and tables, or to drop (remove) existing databases and tables. If you grant the `DROP` privilege for the `mysql` database to a user, that user can drop the database in which the MySQL access privileges are stored!

The `SELECT`, `INSERT`, `UPDATE`, and `DELETE` privileges allow you to perform operations on rows in existing tables in a database. `INSERT` is also required for the `ANALYZE TABLE`, `OPTIMIZE TABLE`, and `REPAIR TABLE` table-maintenance statements.

`SELECT` statements require the `SELECT` privilege only if they actually retrieve rows from a table. Some

`SELECT` statements do not access tables and can be executed without permission for any database. For example, you can use the `mysql` client as a simple calculator to evaluate expressions that make no reference to tables:

```
SELECT 1+1;
SELECT PI()*2;
```

The `INDEX` privilege enables you to create or drop (remove) indexes. `INDEX` applies to existing tables. If you have the `CREATE` privilege for a table, you can include index definitions in the `CREATE TABLE` statement.

The `ALTER` privilege enables you to use `ALTER TABLE` to change the structure of or rename tables.

The `GRANT` privilege enables you to give to other users those privileges that you yourself possess. It can be used for databases, tables, and stored routines.

The `FILE` privilege gives you permission to read and write files on the server host using the `LOAD DATA INFILE` and `SELECT ... INTO OUTFILE` statements. A user who has the `FILE` privilege can read any file on the server host that is either world-readable or readable by the MySQL server. (This implies the user can read any file in any database directory, because the server can access any of those files.) The `FILE` privilege also enables the user to create new files in any directory where the MySQL server has write access. As a security measure, the server will not overwrite existing files.

The remaining privileges are used for administrative operations. Many of them can be performed by using the `mysqladmin` program or by issuing SQL statements. The following table shows which `mysqladmin` commands each administrative privilege enables you to execute:

Privilege	Commands Permitted to Privilege Holders
<code>RELOAD</code>	<code>flush-hosts</code> , <code>flush-logs</code> , <code>flush-privileges</code> , <code>flush-status</code> , <code>flush-tables</code> , <code>flush-threads</code> , <code>refresh</code> , <code>reload</code>
<code>SHUTDOWN</code>	<code>shutdown</code>
<code>PROCESS</code>	<code>processlist</code>
<code>SUPER</code>	<code>kill</code>

The `reload` command tells the server to re-read the grant tables into memory. `flush-privileges` is a synonym for `reload`. The `refresh` command closes and reopens the log files and flushes all tables. The other `flush-xxx` commands perform functions similar to `refresh`, but are more specific and may be preferable in some instances. For example, if you want to flush just the log files, `flush-logs` is a better choice than `refresh`.

The `shutdown` command shuts down the server. There is no corresponding SQL statement.

The `processlist` command displays information about the threads executing within the server (that is, information about the statements being executed by clients). The `kill` command terminates server threads. You can always display or kill your own threads, but you need the `PROCESS` privilege to display threads initiated by other users and the `SUPER` privilege to kill them. See [Section 13.5.5.3, “KILL Syntax”](#). Prior to MySQL 4.0.2 when `SUPER` was introduced, the `PROCESS` privilege controls the ability to both see and terminate threads for other clients.

The `CREATE TEMPORARY TABLES` privilege enables the use of the keyword `TEMPORARY` in `CREATE TABLE` statements.

The `LOCK TABLES` privilege enables the use of explicit `LOCK TABLES` statements to lock tables for which you have the `SELECT` privilege. This includes the use of write locks, which prevents anyone else from reading the locked table.

The `REPLICATION CLIENT` privilege enables the use of `SHOW MASTER STATUS` and `SHOW SLAVE STATUS`.

The `REPLICATION SLAVE` privilege should be granted to accounts that are used by slave servers to connect to the current server as their master. Without this privilege, the slave cannot request updates that have been made to databases on the master server.

The `SHOW DATABASES` privilege allows the account to see database names by issuing the `SHOW DATABASE` statement. Accounts that do not have this privilege see only databases for which they have some privileges, and cannot use the statement at all if the server was started with the `--skip-show-database` option. Note that *any* global privilege is a privilege for the database.

It is a good idea to grant to an account only those privileges that it needs. You should exercise particular caution in granting the `FILE` and administrative privileges:

- The `FILE` privilege can be abused to read into a database table any files that the MySQL server can read on the server host. This includes all world-readable files and files in the server's data directory. The table can then be accessed using `SELECT` to transfer its contents to the client host.
- The `GRANT` privilege enables users to give their privileges to other users. Two users that have different privileges and with the `GRANT` privilege are able to combine privileges.
- The `ALTER` privilege may be used to subvert the privilege system by renaming tables.
- The `SHUTDOWN` privilege can be abused to deny service to other users entirely by terminating the server.
- The `PROCESS` privilege can be used to view the plain text of currently executing statements, including statements that set or change passwords.
- The `SUPER` privilege can be used to terminate other clients or change how the server operates.
- Privileges granted for the `mysql` database itself can be used to change passwords and other access privilege information. Passwords are stored encrypted, so a malicious user cannot simply read them to know the plain text password. However, a user with write access to the `user` table `Password` column can change an account's password, and then connect to the MySQL server using that account.

There are some things that you cannot do with the MySQL privilege system:

- You cannot explicitly specify that a given user should be denied access. That is, you cannot explicitly match a user and then refuse the connection.
- You cannot specify that a user has privileges to create or drop tables in a database but not to create or drop the database itself.
- A password applies globally to an account. You cannot associate a password with a specific object such as a database or table.

## 5.7.4. Connecting to the MySQL Server

MySQL client programs generally expect you to specify certain connection parameters when you want to access a MySQL server:

- The name of the host where the MySQL server is running
- Your username
- Your password

For example, the `mysql` client can be started as follows from a command-line prompt (indicated here by `shell>`):

```
shell> mysql -h host_name -u user_name -p your_pass
```

Alternative forms of the `-h`, `-u`, and `-p` options are `--host=host_name`, `--user=user_name`, and `--password=your_pass`. Note that there is *no space* between `-p` or `--password=` and the password following it.

If you use a `-p` or `--password` option but do not specify the password value, the client program prompts you to enter the password. The password is not displayed as you enter it. This is more secure than giving the password on the command line. Any user on your system may be able to see a password specified on the command line by executing a command such as `ps auxw`. See [Section 5.8.6, “Keeping Your Password Secure”](#).

MySQL client programs use default values for any connection parameter option that you do not specify:

- The default hostname is `localhost`.
- The default username is `ODBC` on Windows and your Unix login name on Unix.
- No password is supplied if neither `-p` nor `--password` is given.

Thus, for a Unix user with a login name of `joe`, all of the following commands are equivalent:

```
shell> mysql -h localhost -u joe
shell> mysql -h localhost
shell> mysql -u joe
shell> mysql
```

Other MySQL clients behave similarly.

You can specify different default values to be used when you make a connection so that you need not enter them on the command line each time you invoke a client program. This can be done in a couple of ways:

- You can specify connection parameters in the `[client]` section of an option file. The relevant section of the file might look like this:

```
[client]
host=host_name
user=user_name
password=your_pass
```

[Section 4.3.2, “Using Option Files”](#), discusses option files further.

- You can specify some connection parameters using environment variables. The host can be specified for `mysql` using `MYSQL_HOST`. The MySQL username can be specified using `USER` (this is for Windows and NetWare only). The password can be specified using `MYSQL_PWD`, although this is insecure; see [Section 5.8.6, “Keeping Your Password Secure”](#). For a list of variables, see [Ap-](#)



pendix F, *Environment Variables*.

- `--ssl*`

Options that begin with `--ssl` specify whether to allow clients to connect via SSL and indicate where to find SSL keys and certificates. See [Section 5.8.7.3, “SSL Command Options”](#).

## 5.7.5. Access Control, Stage 1: Connection Verification

When you attempt to connect to a MySQL server, the server accepts or rejects the connection based on your identity and whether you can verify your identity by supplying the correct password. If not, the server denies access to you completely. Otherwise, the server accepts the connection, and then enters Stage 2 and waits for requests.

Your identity is based on two pieces of information:

- The client host from which you connect
- Your MySQL username

Identity checking is performed using the three `user` table scope columns (`Host`, `User`, and `Password`). The server accepts the connection only if the `Host` and `User` columns in some `user` table row match the client hostname and username and the client supplies the password specified in that row.

`Host` values in the `user` table may be specified as follows:

- A `Host` value may be a hostname or an IP number, or `'localhost'` to indicate the local host.
- You can use the wildcard characters `'%'` and `'_'` in `Host` column values. These have the same meaning as for pattern-matching operations performed with the `LIKE` operator. For example, a `Host` value of `'%'` matches any hostname, whereas a value of `'%.mysql.com'` matches any host in the `mysql.com` domain.
- As of MySQL 3.23, for `Host` values specified as IP numbers, you can specify a netmask indicating how many address bits to use for the network number. For example:

```
GRANT ALL PRIVILEGES ON db.* TO david@'192.58.197.0/255.255.255.0';
```

This allows `david` to connect from any client host having an IP number `client_ip` for which the following condition is true:

```
client_ip & netmask = host_ip
```

That is, for the `GRANT` statement just shown:

```
client_ip & 255.255.255.0 = 192.58.197.0
```

IP numbers that satisfy this condition and can connect to the MySQL server are those in the range from `192.58.197.0` to `192.58.197.255`.

Note: The netmask can only be used to tell the server to use 8, 16, 24, or 32 bits of the address. Examples:

- `192.0.0.0/255.0.0.0`: anything on the 192 class A network

- `192.168.0.0/255.255.0.0`: anything on the 192.168 class B network
- `192.168.1.0/255.255.255.0`: anything on the 192.168.1 class C network
- `192.168.1.1`: only this specific IP

The following netmask (28 bits) will not work:

```
192.168.0.1/255.255.255.240
```

- A blank `Host` value in a `db` table row means that its privileges should be combined with those in the row in the `host` table that matches the client hostname. The privileges are combined using an AND (intersection) operation, not OR (union). [Section 5.7.6, “Access Control, Stage 2: Request Verification”](#), discusses use of the `host` table further.

A blank `Host` value in the other grant tables is the same as `'%'`.

Because you can use IP wildcard values in the `Host` column (for example, `'144.155.166.%'` to match every host on a subnet), someone could try to exploit this capability by naming a host `144.155.166.somewhere.com`. To foil such attempts, MySQL disallows matching on hostnames that start with digits and a dot. Thus, if you have a host named something like `1.2.foo.com`, its name never matches the `Host` column of the grant tables. An IP wildcard value can match only IP numbers, not hostnames.

In the `User` column, wildcard characters are not allowed, but you can specify a blank value, which matches any name. If the `user` table row that matches an incoming connection has a blank username, the user is considered to be an anonymous user with no name, not a user with the name that the client actually specified. This means that a blank username is used for all further access checking for the duration of the connection (that is, during Stage 2).

The `Password` column can be blank. This is not a wildcard and does not mean that any password matches. It means that the user must connect without specifying a password.

Non-blank `Password` values in the `user` table represent encrypted passwords. MySQL does not store passwords in plaintext form for anyone to see. Rather, the password supplied by a user who is attempting to connect is encrypted (using the `PASSWORD()` function). The encrypted password then is used during the connection process when checking whether the password is correct. (This is done without the encrypted password ever traveling over the connection.) From MySQL's point of view, the encrypted password is the *real* password, so you should never give anyone access to it. In particular, *do not give non-administrative users read access to tables in the `mysql` database.*

From version 4.1 on, MySQL employs a stronger authentication method that has better password protection during the connection process than in earlier versions. It is secure even if TCP/IP packets are sniffed or the `mysql` database is captured. [Section 5.7.9, “Password Hashing as of MySQL 4.1”](#), discusses password encryption further.

The following table shows how various combinations of `Host` and `User` values in the `user` table apply to incoming connections.

Host Value	User Value	Allowable Connections
<code>'thomas.loc.gov'</code>	<code>'fred'</code>	<code>fred</code> , connecting from <code>thomas.loc.gov</code>
<code>'thomas.loc.gov'</code>	<code>''</code>	Any user, connecting from <code>thomas.loc.gov</code>
<code>'%'</code>	<code>'fred'</code>	<code>fred</code> , connecting from any host
<code>'%'</code>	<code>''</code>	Any user, connecting from any host

'%.loc.gov'	'fred'	fred, connecting from any host in the loc.gov domain
'x.y.%'	'fred'	fred, connecting from x.y.net, x.y.com, x.y.edu, and so on (this is probably not useful)
'144.155.166.177'	'fred'	fred, connecting from the host with IP address 144.155.166.177
'144.155.166.%'	'fred'	fred, connecting from any host in the 144.155.166 class C subnet
'144.155.166.0/255.255.255.0'	'fred'	Same as previous example

It is possible for the client hostname and username of an incoming connection to match more than one row in the `user` table. The preceding set of examples demonstrates this: Several of the entries shown match a connection from `thomas.loc.gov` by `fred`.

When multiple matches are possible, the server must determine which of them to use. It resolves this issue as follows:

- Whenever the server reads the `user` table into memory, it sorts the rows.
- When a client attempts to connect, the server looks through the rows in sorted order.
- The server uses the first row that matches the client hostname and username.

To see how this works, suppose that the `user` table looks like this:

```

+-----+-----+
| Host   | User   | ...
+-----+-----+
| %      | root   | ...
| %      | jeffrey| ...
| localhost | root   | ...
| localhost |       | ...
+-----+-----+

```

When the server reads the table into memory, it orders the rows with the most-specific `Host` values first. Literal hostnames and IP numbers are the most specific. The pattern `'%'` means “any host” and is least specific. Rows with the same `Host` value are ordered with the most-specific `User` values first (a blank `User` value means “any user” and is least specific). For the `user` table just shown, the result after sorting looks like this:

```

+-----+-----+
| Host   | User   | ...
+-----+-----+
| localhost | root   | ...
| localhost |       | ...
| %      | jeffrey| ...
| %      | root   | ...
+-----+-----+

```

When a client attempts to connect, the server looks through the sorted rows and uses the first match found. For a connection from `localhost` by `jeffrey`, two of the rows from the table match: the one with `Host` and `User` values of `'localhost'` and `''`, and the one with values of `'%'` and `'jeffrey'`. The `'localhost'` row appears first in sorted order, so that is the one the server uses.

Here is another example. Suppose that the `user` table looks like this:

Host	User	...
%	jeffrey	...
thomas.loc.gov		...

The sorted table looks like this:

Host	User	...
thomas.loc.gov		...
%	jeffrey	...

A connection by `jeffrey` from `thomas.loc.gov` is matched by the first row, whereas a connection by `jeffrey` from `whitehouse.gov` is matched by the second.

It is a common misconception to think that, for a given username, all rows that explicitly name that user are used first when the server attempts to find a match for the connection. This is simply not true. The previous example illustrates this, where a connection from `thomas.loc.gov` by `jeffrey` is first matched not by the row containing '`jeffrey`' as the `User` column value, but by the row with no username. As a result, `jeffrey` is authenticated as an anonymous user, even though he specified a username when connecting.

If you are able to connect to the server, but your privileges are not what you expect, you probably are being authenticated as some other account. To find out what account the server used to authenticate you, use the `CURRENT_USER()` function. (See [Section 12.9.3, “Information Functions”](#).) It returns a value in `user_name@host_name` format that indicates the `User` and `Host` values from the matching `user` table row. Suppose that `jeffrey` connects and issues the following query:

```
mysql> SELECT CURRENT_USER();
+-----+
| CURRENT_USER() |
+-----+
| @localhost     |
+-----+
```

The result shown here indicates that the matching `user` table row had a blank `User` column value. In other words, the server is treating `jeffrey` as an anonymous user.

The `CURRENT_USER()` function is available as of MySQL 4.0.6. See [Section 12.9.3, “Information Functions”](#). Another way to diagnose authentication problems is to print out the `user` table and sort it by hand to see where the first match is being made.

## 5.7.6. Access Control, Stage 2: Request Verification

After you establish a connection, the server enters Stage 2 of access control. For each request that you issue via that connection, the server determines what operation you want to perform, then checks whether you have sufficient privileges to do so. This is where the privilege columns in the grant tables come into play. These privileges can come from any of the `user`, `db`, `host`, `tables_priv`, or `columns_priv` tables. (You may find it helpful to refer to [Section 5.7.2, “How the Privilege System Works”](#), which lists the columns present in each of the grant tables.)

The `user` table grants privileges that are assigned to you on a global basis and that apply no matter what the default database is. For example, if the `user` table grants you the `DELETE` privilege, you can delete rows from any table in any database on the server host! In other words, `user` table privileges are superuser privileges. It is wise to grant privileges in the `user` table only to superusers such as database administrators. For other users, you should leave all privileges in the `user` table set to '`N`' and grant privileges at more specific levels only. You can grant privileges for particular databases, tables, or

columns.

The `db` and `host` tables grant database-specific privileges. Values in the scope columns of these tables can take the following forms:

- The wildcard characters ‘%’ and ‘\_’ can be used in the `Host` and `Db` columns of either table. These have the same meaning as for pattern-matching operations performed with the `LIKE` operator. If you want to use either character literally when granting privileges, you must escape it with a backslash. For example, to include ‘\_’ character as part of a database name, specify it as ‘\\_’ in the `GRANT` statement.
- A ‘%’ `Host` value in the `db` table means “any host.” A blank `Host` value in the `db` table means “consult the `host` table for further information” (a process that is described later in this section).
- A ‘%’ or blank `Host` value in the `host` table means “any host.”
- A ‘%’ or blank `Db` value in either table means “any database.”
- A blank `User` value in either table matches the anonymous user.

The server reads the `db` and `host` tables into memory and sorts them at the same time that it reads the `user` table. The server sorts the `db` table based on the `Host`, `Db`, and `User` scope columns, and sorts the `host` table based on the `Host` and `Db` scope columns. As with the `user` table, sorting puts the most-specific values first and least-specific values last, and when the server looks for matching entries, it uses the first match that it finds.

The `tables_priv` and `columns_priv` tables grant table-specific and column-specific privileges. Values in the scope columns of these tables can take the following forms:

- The wildcard characters ‘%’ and ‘\_’ can be used in the `Host` column. These have the same meaning as for pattern-matching operations performed with the `LIKE` operator.
- A ‘%’ or blank `Host` value means “any host.”
- The `Db`, `Table_name`, and `Column_name` columns cannot contain wildcards or be blank.

The server sorts the `tables_priv` and `columns_priv` tables based on the `Host`, `Db`, and `User` columns. This is similar to `db` table sorting, but simpler because only the `Host` column can contain wildcards.

The server uses the sorted tables to verify each request that it receives. For requests that require administrative privileges such as `SHUTDOWN` or `RELOAD`, the server checks only the `user` table row because that is the only table that specifies administrative privileges. The server grants access if the row allows the requested operation and denies access otherwise. For example, if you want to execute `mysqladmin shutdown` but your `user` table row does not grant the `SHUTDOWN` privilege to you, the server denies access without even checking the `db` or `host` tables. (They contain no `Shutdown_priv` column, so there is no need to do so.)

For database-related requests (`INSERT`, `UPDATE`, and so on), the server first checks the user's global (superuser) privileges by looking in the `user` table row. If the row allows the requested operation, access is granted. If the global privileges in the `user` table are insufficient, the server determines the user's database-specific privileges by checking the `db` and `host` tables:

1. The server looks in the `db` table for a match on the `Host`, `Db`, and `User` columns. The `Host` and `User` columns are matched to the connecting user's hostname and MySQL username. The `Db`

column is matched to the database that the user wants to access. If there is no row for the `Host` and `User`, access is denied.

2. If there is a matching `db` table row and its `Host` column is not blank, that row defines the user's database-specific privileges.
3. If the matching `db` table row's `Host` column is blank, it signifies that the `host` table enumerates which hosts should be allowed access to the database. In this case, a further lookup is done in the `host` table to find a match on the `Host` and `Db` columns. If no `host` table row matches, access is denied. If there is a match, the user's database-specific privileges are computed as the intersection (not the union!) of the privileges in the `db` and `host` table entries; that is, the privileges that are 'Y' in both entries. (This way you can grant general privileges in the `db` table row and then selectively restrict them on a host-by-host basis using the `host` table entries.)

After determining the database-specific privileges granted by the `db` and `host` table entries, the server adds them to the global privileges granted by the `user` table. If the result allows the requested operation, access is granted. Otherwise, the server successively checks the user's table and column privileges in the `tables_priv` and `columns_priv` tables, adds those to the user's privileges, and allows or denies access based on the result.

Expressed in boolean terms, the preceding description of how a user's privileges are calculated may be summarized like this:

```
global privileges
OR (database privileges AND host privileges)
OR table privileges
OR column privileges
```

It may not be apparent why, if the global `user` row privileges are initially found to be insufficient for the requested operation, the server adds those privileges to the database, table, and column privileges later. The reason is that a request might require more than one type of privilege. For example, if you execute an `INSERT INTO ... SELECT` statement, you need both the `INSERT` and the `SELECT` privileges. Your privileges might be such that the `user` table row grants one privilege and the `db` table row grants the other. In this case, you have the necessary privileges to perform the request, but the server cannot tell that from either table by itself; the privileges granted by the entries in both tables must be combined.

The `host` table is not affected by the `GRANT` or `REVOKE` statements, so it is unused in most MySQL installations. If you modify it directly, you can use it for some specialized purposes, such as to maintain a list of secure servers. For example, at TcX, the `host` table contains a list of all machines on the local network. These are granted all privileges.

You can also use the `host` table to indicate hosts that are *not* secure. Suppose that you have a machine `public.your.domain` that is located in a public area that you do not consider secure. You can allow access to all hosts on your network except that machine by using `host` table entries like this:

```
+-----+-----+-----+
| Host | Db | ... |
+-----+-----+-----+
| public.your.domain | % | ... (all privileges set to 'N') |
| %.your.domain | % | ... (all privileges set to 'Y') |
+-----+-----+-----+
```

Naturally, you should always test your changes to the grant tables (for example, by using `SHOW GRANTS`) to make sure that your access privileges are actually set up the way you think they are.

### 5.7.7. When Privilege Changes Take Effect

When `mysqld` starts, it reads all grant table contents into memory. The in-memory tables become effective for access control at that point.

When the server reloads the grant tables, privileges for existing client connections are affected as follows:

- Table and column privilege changes take effect with the client's next request.
- Database privilege changes take effect at the next `USE db_name` statement.

**Note:** Client applications may cache the database name; thus, this effect may not be visible to them without actually changing to a different database or executing a `FLUSH PRIVILEGES` statement.

- Changes to global privileges and passwords take effect the next time the client connects.

If you modify the grant tables indirectly using statements such as `GRANT`, `REVOKE`, or `SET PASSWORD`, the server notices these changes and loads the grant tables into memory again immediately.

If you modify the grant tables directly using statements such as `INSERT`, `UPDATE`, or `DELETE`, your changes have no effect on privilege checking until you either restart the server or tell it to reload the tables. To reload the grant tables manually, issue a `FLUSH PRIVILEGES` statement or execute a `mysqladmin flush-privileges` or `mysqladmin reload` command.

If you change the grant tables directly but forget to reload them, your changes have *no effect* until you restart the server. This may leave you wondering why your changes do not seem to make any difference!

## 5.7.8. Causes of Access denied Errors

If you encounter problems when you try to connect to the MySQL server, the following items describe some courses of action you can take to correct the problem.

- Make sure that the server is running. If it is not running, you cannot connect to it. For example, if you attempt to connect to the server and see a message such as one of those following, one cause might be that the server is not running:

```
shell> mysql
ERROR 2003: Can't connect to MySQL server on 'host_name' (111)
shell> mysql
ERROR 2002: Can't connect to local MySQL server through socket
'/tmp/mysql.sock' (111)
```

It might also be that the server is running, but you are trying to connect using a TCP/IP port, named pipe, or Unix socket file different from the one on which the server is listening. To correct this when you invoke a client program, specify a `--port` option to indicate the proper port number, or a `--socket` option to indicate the proper named pipe or Unix socket file. To find out where the socket file is, you can use this command:

```
shell> netstat -ln | grep mysql
```

- The grant tables must be properly set up so that the server can use them for access control. For some distribution types (such as binary distributions on Windows, or RPM distributions on Linux), the installation process initializes the `mysql` database containing the grant tables. For distributions that do not do this, you must initialize the grant tables manually by running the `mysql_install_db` script. For details, see [Section 2.10.2, “Unix Post-Installation Procedures”](#).

One way to determine whether you need to initialize the grant tables is to look for a `mysql` direct-

ory under the data directory. (The data directory normally is named `data` or `var` and is located under your MySQL installation directory.) Make sure that you have a file named `user.MYD` in the `mysql` database directory. If you do not, execute the `mysql_install_db` script. After running this script and starting the server, test the initial privileges by executing this command:

```
shell> mysql -u root test
```

The server should let you connect without error.

- After a fresh installation, you should connect to the server and set up your users and their access permissions:

```
shell> mysql -u root mysql
```

The server should let you connect because the MySQL `root` user has no password initially. That is also a security risk, so setting the password for the `root` accounts is something you should do while you're setting up your other MySQL accounts. For instructions on setting the initial passwords, see [Section 2.10.3, “Securing the Initial MySQL Accounts”](#).

- If you have updated an existing MySQL installation to a newer version, did you run the `mysql_fix_privilege_tables` script? If not, do so. The structure of the grant tables changes occasionally when new capabilities are added, so after an upgrade you should always make sure that your tables have the current structure. For instructions, see [Section 5.5.1, “mysql\\_fix\\_privilege\\_tables — Upgrade MySQL System Tables”](#).
- If a client program receives the following error message when it tries to connect, it means that the server expects passwords in a newer format than the client is capable of generating:

```
shell> mysql
Client does not support authentication protocol requested
by server; consider upgrading MySQL client
```

For information on how to deal with this, see [Section 5.7.9, “Password Hashing as of MySQL 4.1”](#), and [Section A.2.3, “Client does not support authentication protocol”](#).

- If you try to connect as `root` and get the following error, it means that you do not have an row in the `user` table with a `User` column value of `'root'` and that `mysqld` cannot resolve the host-name for your client:

```
Access denied for user '@'unknown' to database mysql
```

In this case, you must restart the server with the `--skip-grant-tables` option and edit your `/etc/hosts` file or `\windows\hosts` file to add an entry for your host.

- Remember that client programs use connection parameters specified in option files or environment variables. If a client program seems to be sending incorrect default connection parameters when you have not specified them on the command line, check your environment and any applicable option files. For example, if you get `Access denied` when you run a client without any options, make sure that you have not specified an old password in any of your option files!

You can suppress the use of option files by a client program by invoking it with the `--no-defaults` option. For example:

```
shell> mysqladmin --no-defaults -u root version
```

The option files that clients use are listed in [Section 4.3.2, “Using Option Files”](#). Environment variables are listed in [Appendix F, \*Environment Variables\*](#).



- If you get the following error, it means that you are using an incorrect `root` password:

```
shell> mysqladmin -u root -pxxxx ver
Access denied for user 'root'@'localhost' (using password: YES)
```

If the preceding error occurs even when you have not specified a password, it means that you have an incorrect password listed in some option file. Try the `--no-defaults` option as described in the previous item.

For information on changing passwords, see [Section 5.8.5, “Assigning Account Passwords”](#).

If you have lost or forgotten the `root` password, you can restart `mysqld` with `--skip-grant-tables` to change the password. See [Section A.4.1, “How to Reset the Root Password”](#).

- If you change a password by using `SET PASSWORD`, `INSERT`, or `UPDATE`, you must encrypt the password using the `PASSWORD()` function. If you do not use `PASSWORD()` for these statements, the password will not work. For example, the following statement sets a password, but fails to encrypt it, so the user is not able to connect afterward:

```
SET PASSWORD FOR 'abe'@'host_name' = 'eagle';
```

Instead, set the password like this:

```
SET PASSWORD FOR 'abe'@'host_name' = PASSWORD('eagle');
```

The `PASSWORD()` function is unnecessary when you specify a password using the `GRANT` statement or the `mysqladmin password` command. Each of those automatically uses `PASSWORD()` to encrypt the password. See [Section 5.8.5, “Assigning Account Passwords”](#).

- `localhost` is a synonym for your local hostname, and is also the default host to which clients try to connect if you specify no host explicitly. However, connections to `localhost` on Unix systems do not work if you are using a MySQL version older than 3.23.27 that uses MIT-pthreads: `localhost` connections are made using Unix socket files, which were not supported by MIT-pthreads at that time.

To avoid this problem on such systems, you can use a `--host=127.0.0.1` option to name the server host explicitly. This will make a TCP/IP connection to the local `mysqld` server. You can also use TCP/IP by specifying a `--host` option that uses the actual hostname of the local host. In this case, the hostname must be specified in a `user` table row on the server host, even though you are running the client program on the same host as the server.

- If you get an `Access denied` error when trying to connect to the database with `mysql -u user_name`, you may have a problem with the `user` table. Check this by executing `mysql -u root mysql` and issuing this SQL statement:

```
SELECT * FROM user;
```

The result should include a row with the `Host` and `User` columns matching your computer's hostname and your MySQL username.

- The `Access denied` error message tells you who you are trying to log in as, the client host from which you are trying to connect, and whether you were using a password. Normally, you should have one row in the `user` table that exactly matches the hostname and username that were given in the error message. For example, if you get an error message that contains `using password: NO`, it means that you tried to log in without a password.

- If the following error occurs when you try to connect from a host other than the one on which the MySQL server is running, it means that there is no row in the `user` table with a `Host` value that matches the client host:

```
Host ... is not allowed to connect to this MySQL server
```

You can fix this by setting up an account for the combination of client hostname and username that you are using when trying to connect.

If you do not know the IP number or hostname of the machine from which you are connecting, you should put a row with `'%'` as the `Host` column value in the `user` table. After trying to connect from the client machine, use a `SELECT USER()` query to see how you really did connect. (Then change the `'%'` in the `user` table row to the actual hostname that shows up in the log. Otherwise, your system is left insecure because it allows connections from any host for the given username.)

(Note that if you are running a version of MySQL older than 3.23.11, the output from `USER()` does not include the hostname. In this case, you must restart the server with the `--log` option, then obtain the hostname from the log.)

On Linux, another reason that this error might occur is that you are using a binary MySQL version that is compiled with a different version of the `glibc` library than the one you are using. In this case, you should either upgrade your operating system or `glibc`, or download a source distribution of MySQL version and compile it yourself. A source RPM is normally trivial to compile and install, so this is not a big problem.

- If you specify a hostname when trying to connect, but get an error message where the hostname is not shown or is an IP number, it means that the MySQL server got an error when trying to resolve the IP number of the client host to a name:

```
shell> mysqladmin -u root -pxxxxx -h some_hostname ver
Access denied for user 'root'@'' (using password: YES)
```

This indicates a DNS problem. To fix it, execute `mysqladmin flush-hosts` to reset the internal DNS hostname cache. See [Section 7.5.5, “How MySQL Uses DNS”](#).

Some permanent solutions are:

- Determine what is wrong with your DNS server and fix it.
- Specify IP numbers rather than hostnames in the MySQL grant tables.
- Put an entry for the client machine name in `/etc/hosts` or `\windows\hosts`.
- Start `mysqld` with the `--skip-name-resolve` option.
- Start `mysqld` with the `--skip-host-cache` option.
- On Unix, if you are running the server and the client on the same machine, connect to `localhost`. Unix connections to `localhost` use a Unix socket file rather than TCP/IP.
- On Windows, if you are running the server and the client on the same machine and the server supports named pipe connections, connect to the hostname `.` (period). Connections to `.` use a named pipe rather than TCP/IP.
- If `mysql -u root test` works but `mysql -h your_hostname -u root test` results in `Access denied` (where `your_hostname` is the actual hostname of the local host), you may not have the correct name for your host in the `user` table. A common problem here is that the `Host` value in the `user` table row specifies an unqualified hostname, but your system's name resolution

routines return a fully qualified domain name (or vice versa). For example, if you have an entry with host 'tcx' in the `user` table, but your DNS tells MySQL that your hostname is 'tcx.subnet.se', the entry does not work. Try adding an entry to the `user` table that contains the IP number of your host as the `Host` column value. (Alternatively, you could add an entry to the `user` table with a `Host` value that contains a wildcard; for example, 'tcx.%'. However, use of hostnames ending with '%' is insecure and is *not* recommended.)

- If `mysql -u user_name test` works but `mysql -u user_name other_db_name` does not, you have not granted database access for `other_db_name` to the given user.
- If `mysql -u user_name` works when executed on the server host, but `mysql -h host_name -u user_name` does not work when executed on a remote client host, you have not enabled access to the server for the given username from the remote host.
- If you cannot figure out why you get `Access denied`, remove from the `user` table all entries that have `Host` values containing wildcards (entries that contain '%' or '\_'). A very common error is to insert a new entry with `Host='%'` and `User='some_user'`, thinking that this allows you to specify `localhost` to connect from the same machine. The reason that this does not work is that the default privileges include an entry with `Host='localhost'` and `User=''`. Because that entry has a `Host` value 'localhost' that is more specific than '%', it is used in preference to the new entry when connecting from `localhost`! The correct procedure is to insert a second entry with `Host='localhost'` and `User='some_user'`, or to delete the entry with `Host='localhost'` and `User=''`. After deleting the entry, remember to issue a `FLUSH PRIVILEGES` statement to reload the grant tables.
- If you get the following error, you may have a problem with the `db` or `host` table:

```
Access to database denied
```

If the entry selected from the `db` table has an empty value in the `Host` column, make sure that there are one or more corresponding entries in the `host` table specifying which hosts the `db` table entry applies to.

- If you are able to connect to the MySQL server, but get an `Access denied` message whenever you issue a `SELECT ... INTO OUTFILE` or `LOAD DATA INFILE` statement, your entry in the `user` table does not have the `FILE` privilege enabled.
- If you change the grant tables directly (for example, by using `INSERT`, `UPDATE`, or `DELETE` statements) and your changes seem to be ignored, remember that you must execute a `FLUSH PRIVILEGES` statement or a `mysqladmin flush-privileges` command to cause the server to re-read the privilege tables. Otherwise, your changes have no effect until the next time the server is restarted. Remember that after you change the `root` password with an `UPDATE` command, you will not need to specify the new password until after you flush the privileges, because the server will not know you've changed the password yet!
- If your privileges seem to have changed in the middle of a session, it may be that a MySQL administrator has changed them. Reloading the grant tables affects new client connections, but it also affects existing connections as indicated in [Section 5.7.7, “When Privilege Changes Take Effect”](#).
- If you have access problems with a Perl, PHP, Python, or ODBC program, try to connect to the server with `mysql -u user_name db_name` or `mysql -u user_name -p your_pass db_name`. If you are able to connect using the `mysql` client, the problem lies with your program, not with the access privileges. (There is no space between `-p` and the password; you can also use the `--password=your_pass` syntax to specify the password. If you use the `-p --password` option with no password value, MySQL prompts you for the password.)
- For testing, start the `mysqld` server with the `--skip-grant-tables` option. Then you can change the MySQL grant tables and use the `mysqlaccess` script to check whether your modifica-

tions have the desired effect. When you are satisfied with your changes, execute `mysqladmin flush-privileges` to tell the `mysqld` server to start using the new grant tables. (Reloading the grant tables overrides the `--skip-grant-tables` option. This enables you to tell the server to begin using the grant tables again without stopping and restarting it.)

- If everything else fails, start the `mysqld` server with a debugging option (for example, `-debug=d,general,query`). This prints host and user information about attempted connections, as well as information about each command issued. See [Section E.1.2, “Creating Trace Files”](#).
- If you have any other problems with the MySQL grant tables and feel you must post the problem to the mailing list, always provide a dump of the MySQL grant tables. You can dump the tables with the `mysqldump mysql` command. To file a bug report, see the instructions at [Section 1.8, “How to Report Bugs or Problems”](#). In some cases, you may need to restart `mysqld` with `-skip-grant-tables` to run `mysqldump`.

## 5.7.9. Password Hashing as of MySQL 4.1

MySQL user accounts are listed in the `user` table of the `mysql` database. Each MySQL account is assigned a password, although what is stored in the `Password` column of the `user` table is not the plain-text version of the password, but a hash value computed from it. Password hash values are computed by the `PASSWORD()` function.

MySQL uses passwords in two phases of client/server communication:

- When a client attempts to connect to the server, there is an initial authentication step in which the client must present a password that has a hash value matching the hash value stored in the `user` table for the account that the client wants to use.
- After the client connects, it can (if it has sufficient privileges) set or change the password hashes for accounts listed in the `user` table. The client can do this by using the `PASSWORD()` function to generate a password hash, or by using the `GRANT` or `SET PASSWORD` statements.

In other words, the server *uses* hash values during authentication when a client first attempts to connect. The server *generates* hash values if a connected client invokes the `PASSWORD()` function or uses a `GRANT` or `SET PASSWORD` statement to set or change a password.

The password hashing mechanism was updated in MySQL 4.1 to provide better security and to reduce the risk of passwords being intercepted. However, this new mechanism is understood only by the 4.1 server and 4.1 clients, which can result in some compatibility problems. A 4.1 client can connect to a pre-4.1 server, because the client understands both the old and new password hashing mechanisms. However, a pre-4.1 client that attempts to connect to a 4.1 server may run into difficulties. For example, a 4.0 `mysql` client that attempts to connect to a 4.1 server may fail with the following error message:

```
shell> mysql -h localhost -u root
Client does not support authentication protocol requested
by server; consider upgrading MySQL client
```

Another common example of this phenomenon occurs for attempts to use the older PHP `mysql` extension after upgrading to MySQL 4.1 or newer. (See [Section 17.3.1, “Common Problems with MySQL and PHP”](#).)

The following discussion describes the differences between the old and new password mechanisms, and what you should do if you upgrade your server to 4.1 but need to maintain backward compatibility with pre-4.1 clients. Additional information can be found in [Section A.2.3, “Client does not support authentication protocol”](#). This information is of particular importance to PHP pro-

grammers migrating MySQL databases from version 4.0 or lower to version 4.1 or higher.

**Note:** This discussion contrasts 4.1 behavior with pre-4.1 behavior, but the 4.1 behavior described here actually begins with 4.1.1. MySQL 4.1.0 is an “odd” release because it has a slightly different mechanism than that implemented in 4.1.1 and up. Differences between 4.1.0 and more recent versions are described further in [Section 5.7.9.2, “Password Hashing in MySQL 4.1.0”](#).

Prior to MySQL 4.1, password hashes computed by the `PASSWORD()` function are 16 bytes long. Such hashes look like this:

```
mysql> SELECT PASSWORD('mypass');
+-----+
| PASSWORD('mypass') |
+-----+
| 6f8c114b58f2ce9e   |
+-----+
```

The `Password` column of the `user` table (in which these hashes are stored) also is 16 bytes long before MySQL 4.1.

As of MySQL 4.1, the `PASSWORD()` function has been modified to produce a longer 41-byte hash value:

```
mysql> SELECT PASSWORD('mypass');
+-----+
| PASSWORD('mypass') |
+-----+
| *6C8989366EAF75BB670AD8EA7A7FC1176A95CEF4 |
+-----+
```

Accordingly, the `Password` column in the `user` table also must be 41 bytes long to store these values:

- If you perform a new installation of MySQL 4.1, the `Password` column is made 41 bytes long automatically.
- If you upgrade an older installation to 4.1, you should run the `mysql_fix_privilege_tables` script to increase the length of the `Password` column from 16 to 41 bytes. (The script does not change existing password values, which remain 16 bytes long.)

A widened `Password` column can store password hashes in both the old and new formats. The format of any given password hash value can be determined two ways:

- The obvious difference is the length (16 bytes versus 41 bytes).
- A second difference is that password hashes in the new format always begin with a ‘\*’ character, whereas passwords in the old format never do.

The longer password hash format has better cryptographic properties, and client authentication based on long hashes is more secure than that based on the older short hashes.

The differences between short and long password hashes are relevant both for how the server uses passwords during authentication and for how it generates password hashes for connected clients that perform password-changing operations.

The way in which the server uses password hashes during authentication is affected by the width of the `Password` column:

- If the column is short, only short-hash authentication is used.
- If the column is long, it can hold either short or long hashes, and the server can use either format:
  - Pre-4.1 clients can connect, although because they know only about the old hashing mechanism, they can authenticate only using accounts that have short hashes.
  - 4.1 clients can authenticate using accounts that have short or long hashes.

Even for short-hash accounts, the authentication process is actually a bit more secure for 4.1 and later clients than for older clients. In terms of security, the gradient from least to most secure is:

- Pre-4.1 client authenticating with short password hash
- 4.1 client authenticating with short password hash
- 4.1 client authenticating with long password hash

The way in which the server generates password hashes for connected clients is affected by the width of the `Password` column and by the `--old-passwords` option. A 4.1 server generates long hashes only if certain conditions are met: The `Password` column must be wide enough to hold long values and the `--old-passwords` option must not be given. These conditions apply as follows:

- The `Password` column must be wide enough to hold long hashes (41 bytes). If the column has not been updated and still has the pre-4.1 width of 16 bytes, the server notices that long hashes cannot fit into it and generates only short hashes when a client performs password-changing operations using `PASSWORD()`, `GRANT`, or `SET PASSWORD`. This is the behavior that occurs if you have upgraded to 4.1 but have not yet run the `mysql_fix_privilege_tables` script to widen the `Password` column.
- If the `Password` column is wide, it can store either short or long password hashes. In this case, `PASSWORD()`, `GRANT`, and `SET PASSWORD` generate long hashes unless the server was started with the `--old-passwords` option. That option forces the server to generate short password hashes instead.

The purpose of the `--old-passwords` option is to enable you to maintain backward compatibility with pre-4.1 clients under circumstances where the server would otherwise generate long password hashes. The option does not affect authentication (4.1 clients can still use accounts that have long password hashes), but it does prevent creation of a long password hash in the `user` table as the result of a password-changing operation. Were that to occur, the account no longer could be used by pre-4.1 clients. Without the `--old-passwords` option, the following undesirable scenario is possible:

- An old client connects to an account that has a short password hash.
- The client changes its own password. Without `--old-passwords`, this results in the account having a long password hash.
- The next time the old client attempts to connect to the account, it cannot, because the account has a long password hash that requires the new hashing mechanism during authentication. (Once an account has a long password hash in the user table, only 4.1 clients can authenticate for it, because pre-4.1 clients do not understand long hashes.)

This scenario illustrates that, if you must support older pre-4.1 clients, it is dangerous to run a 4.1 server without using the `--old-passwords` option. By running the server with `--old-passwords`, password-changing operations do not generate long password hashes and thus do not cause accounts to become inaccessible to older clients. (Those clients cannot inadvertently lock themselves out by changing their password and ending up with a long password hash.)

The downside of the `--old-passwords` option is that any passwords you create or change use short hashes, even for 4.1 clients. Thus, you lose the additional security provided by long password hashes. If you want to create an account that has a long hash (for example, for use by 4.1 clients), you must do so while running the server without `--old-passwords`.

The following scenarios are possible for running a 4.1 or later server:

**Scenario 1:** Short `Password` column in user table:

- Only short hashes can be stored in the `Password` column.
- The server uses only short hashes during client authentication.
- For connected clients, password hash-generating operations involving `PASSWORD()`, `GRANT`, or `SET PASSWORD` use short hashes exclusively. Any change to an account's password results in that account having a short password hash.
- The `--old-passwords` option can be used but is superfluous because with a short `Password` column, the server generates only short password hashes anyway.

**Scenario 2:** Long `Password` column; server not started with `--old-passwords` option:

- Short or long hashes can be stored in the `Password` column.
- 4.1 and later clients can authenticate using accounts that have short or long hashes.
- Pre-4.1 clients can authenticate only using accounts that have short hashes.
- For connected clients, password hash-generating operations involving `PASSWORD()`, `GRANT`, or `SET PASSWORD` use long hashes exclusively. A change to an account's password results in that account having a long password hash.

As indicated earlier, a danger in this scenario is that it is possible for accounts that have a short password hash to become inaccessible to pre-4.1 clients. A change to such an account's password made via `GRANT`, `PASSWORD()`, or `SET PASSWORD` results in the account being given a long password hash. From that point on, no pre-4.1 client can authenticate to that account until the client upgrades to 4.1.

To deal with this problem, you can change a password in a special way. For example, normally you use `SET PASSWORD` as follows to change an account password:

```
SET PASSWORD FOR 'some_user'@'some_host' = PASSWORD('mypass');
```

To change the password but create a short hash, use the `OLD_PASSWORD()` function instead:

```
SET PASSWORD FOR 'some_user'@'some_host' = OLD_PASSWORD('mypass');
```

`OLD_PASSWORD()` is useful for situations in which you explicitly want to generate a short hash.

**Scenario 3:** Long `Password` column; server started with `--old-passwords` option:

- Short or long hashes can be stored in the `Password` column.
- 4.1 clients can authenticate for accounts that have short or long hashes (but note that it is possible to create long hashes only when the server is started without `--old-passwords`).
- Pre-4.1 clients can authenticate only for accounts that have short hashes.
- For connected clients, password hash-generating operations involving `PASSWORD()`, `GRANT`, or `SET PASSWORD` use short hashes exclusively. Any change to an account's password results in that account having a short password hash.

In this scenario, you cannot create accounts that have long password hashes, because the `--old-passwords` option prevents generation of long hashes. Also, if you create an account with a long hash before using the `--old-passwords` option, changing the account's password while `--old-passwords` is in effect results in the account being given a short password, causing it to lose the security benefits of a longer hash.

The disadvantages for these scenarios may be summarized as follows:

In scenario 1, you cannot take advantage of longer hashes that provide more secure authentication.

In scenario 2, accounts with short hashes become inaccessible to pre-4.1 clients if you change their passwords without explicitly using `OLD_PASSWORD()`.

In scenario 3, `--old-passwords` prevents accounts with short hashes from becoming inaccessible, but password-changing operations cause accounts with long hashes to revert to short hashes, and you cannot change them back to long hashes while `--old-passwords` is in effect.

### 5.7.9.1. Implications of Password Hashing Changes for Application Programs

An upgrade to MySQL 4.1 can cause a compatibility issue for applications that use `PASSWORD()` to generate passwords for their own purposes. Applications really should not do this, because `PASSWORD()` should be used only to manage passwords for MySQL accounts. But some applications use `PASSWORD()` for their own purposes anyway.

If you upgrade to 4.1 and run the server under conditions where it generates long password hashes, an application that uses `PASSWORD()` for its own passwords breaks. The recommended course of action is to modify the application to use another function, such as `SHA1()` or `MD5()`, to produce hashed values. If that is not possible, you can use the `OLD_PASSWORD()` function, which is provided to generate short hashes in the old format. But note that `OLD_PASSWORD()` may one day no longer be supported.

If the server is running under circumstances where it generates short hashes, `OLD_PASSWORD()` is available but is equivalent to `PASSWORD()`.

PHP programmers migrating their MySQL databases from version 4.0 or lower to version 4.1 or higher should see [Section 17.3, “MySQL PHP API”](#).

### 5.7.9.2. Password Hashing in MySQL 4.1.0

Password hashing in MySQL 4.1.0 differs from hashing in 4.1.1 and up. The 4.1.0 differences are:

- Password hashes are 45 bytes long rather than 41 bytes.
- The `PASSWORD()` function is non-repeatable. That is, with a given argument `X`, successive calls to `PASSWORD(X)` generate different results.



These differences make authentication in 4.1.0 incompatible with that of releases that follow it. If you have upgraded to MySQL 4.1.0, it is recommended that you upgrade to a newer version as soon as possible. After you do, reassign any long passwords in the `user` table so that they are compatible with the 41-byte format.

## 5.8. MySQL User Account Management

This section describes how to set up accounts for clients of your MySQL server. It discusses the following topics:

- The meaning of account names and passwords as used in MySQL and how that compares to names and passwords used by your operating system
- How to set up new accounts and remove existing accounts
- How to change passwords
- Guidelines for using passwords securely
- How to use secure connections with SSL

### 5.8.1. MySQL Usernames and Passwords

A MySQL account is defined in terms of a username and the client host or hosts from which the user can connect to the server. The account also has a password. There are several distinctions between the way usernames and passwords are used by MySQL and the way they are used by your operating system:

- Usernames, as used by MySQL for authentication purposes, have nothing to do with usernames (login names) as used by Windows or Unix. On Unix, most MySQL clients by default try to log in using the current Unix username as the MySQL username, but that is for convenience only. The default can be overridden easily, because client programs allow any username to be specified with a `-u` or `--user` option. Because this means that anyone can attempt to connect to the server using any username, you cannot make a database secure in any way unless all MySQL accounts have passwords. Anyone who specifies a username for an account that has no password is able to connect successfully to the server.
- MySQL usernames can be up to 16 characters long. This limit is hard-coded in the MySQL servers and clients, and trying to circumvent it by modifying the definitions of the tables in the `mysql` database *does not work*.

**Note:** *You should never alter any of the tables in the `mysql` database in any manner whatsoever except by means of the procedure prescribed by MySQL AB that is described in [Section 5.5.1](#), “`mysql_fix_privilege_tables` — Upgrade MySQL System Tables”. Attempting to redefine MySQL’s system tables in any other fashion results in undefined (and unsupported!) behavior.*

Operating system usernames are completely unrelated to MySQL usernames and may even be of a different maximum length. For example, Unix usernames typically are limited to eight characters.

- MySQL usernames can be up to 16 characters long. *Changing the maximum length is not supported*. If you try to change it, for example by changing the length of the `User` column in the `mysql` database tables, this will result in unpredictable behavior. (Altering privilege tables is not supported in any case.) Operating system usernames might have a different maximum length. For example, Unix usernames typically are limited to eight characters.
- MySQL passwords have nothing to do with passwords for logging in to your operating system.

There is no necessary connection between the password you use to log in to a Windows or Unix machine and the password you use to access the MySQL server on that machine.

- MySQL encrypts passwords using its own algorithm. This encryption is different from that used during the Unix login process. MySQL password encryption is the same as that implemented by the `PASSWORD()` SQL function. Unix password encryption is the same as that implemented by the `ENCRYPT()` SQL function. See the descriptions of the `PASSWORD()` and `ENCRYPT()` functions in [Section 12.9.2, “Encryption and Compression Functions”](#). From version 4.1 on, MySQL employs a stronger authentication method that has better password protection during the connection process than in earlier versions. It is secure even if TCP/IP packets are sniffed or the `mysql` database is captured. (In earlier versions, even though passwords are stored in encrypted form in the `user` table, knowledge of the encrypted password value could be used to connect to the MySQL server.)

When you install MySQL, the grant tables are populated with an initial set of accounts. These accounts have names and access privileges that are described in [Section 2.10.3, “Securing the Initial MySQL Accounts”](#), which also discusses how to assign passwords to them. Thereafter, you normally set up, modify, and remove MySQL accounts using statements such as `GRANT` and `REVOKE`. See [Section 13.5.1, “Account Management Statements”](#).

When you connect to a MySQL server with a command-line client, you should specify the username and password for the account that you want to use:

```
shell> mysql --user=monty --password=guess db_name
```

If you prefer short options, the command looks like this:

```
shell> mysql -u monty -pguess db_name
```

There must be *no space* between the `-p` option and the following password value. See [Section 5.7.4, “Connecting to the MySQL Server”](#).

The preceding commands include the password value on the command line, which can be a security risk. See [Section 5.8.6, “Keeping Your Password Secure”](#). To avoid this problem, specify the `--password` or `-p` option without any following password value:

```
shell> mysql --user=monty --password db_name
shell> mysql -u monty -p db_name
```

When the password option has no password value, the client program prints a prompt and waits for you to enter the password. (In these examples, `db_name` is *not* interpreted as a password because it is separated from the preceding password option by a space.)

On some systems, the library routine that MySQL uses to prompt for a password automatically limits the password to eight characters. That is a problem with the system library, not with MySQL. Internally, MySQL does not have any limit for the length of the password. To work around the problem, change your MySQL password to a value that is eight or fewer characters long, or put your password in an option file.

## 5.8.2. Adding New User Accounts to MySQL

You can create MySQL accounts in two ways:

- By using `GRANT` statements
- By manipulating the MySQL grant tables directly with statements such as `INSERT`, `UPDATE`, or

## DELETE

The preferred method is to use `GRANT` statements, because they are more concise and less error-prone. `GRANT` is described in [Section 13.5.1.2, “GRANT Syntax”](#).

Another option for creating accounts is to use one of several available third-party programs that offer capabilities for MySQL account administration. `phpMyAdmin` is one such program.

The following examples show how to use the `mysql` client program to set up new users. These examples assume that privileges are set up according to the defaults described in [Section 2.10.3, “Securing the Initial MySQL Accounts”](#). This means that to make changes, you must connect to the MySQL server as the MySQL `root` user, and the `root` account must have the `INSERT` privilege for the `mysql` database and the `RELOAD` administrative privilege.

First, use the `mysql` program to connect to the server as the MySQL `root` user:

```
shell> mysql --user=root mysql
```

If you have assigned a password to the `root` account, you also need to supply a `--password` or `-p` option for this `mysql` command and also for those later in this section.

After connecting to the server as `root`, you can add new accounts. The following statements use `GRANT` to set up four new accounts:

```
mysql> GRANT ALL PRIVILEGES ON *.* TO 'monty'@'localhost'
-> IDENTIFIED BY 'some_pass' WITH GRANT OPTION;
mysql> GRANT ALL PRIVILEGES ON *.* TO 'monty'@'%'
-> IDENTIFIED BY 'some_pass' WITH GRANT OPTION;
mysql> GRANT RELOAD,PROCESS ON *.* TO 'admin'@'localhost';
mysql> GRANT USAGE ON *.* TO 'dummy'@'localhost';
```

The accounts created by these `GRANT` statements have the following properties:

- Two of the accounts have a username of `monty` and a password of `some_pass`. Both accounts are superuser accounts with full privileges to do anything. One account (`'monty'@'localhost'`) can be used only when connecting from the local host. The other (`'monty'@'%'`) can be used to connect from any other host. Note that it is necessary to have both accounts for `monty` to be able to connect from anywhere as `monty`. Without the `localhost` account, the anonymous-user account for `localhost` that is created by `mysql_install_db` would take precedence when `monty` connects from the local host. As a result, `monty` would be treated as an anonymous user. The reason for this is that the anonymous-user account has a more specific `Host` column value than the `'monty'@'%'` account and thus comes earlier in the `user` table sort order. (`user` table sorting is discussed in [Section 5.7.5, “Access Control, Stage 1: Connection Verification”](#).)
- One account has a username of `admin` and no password. This account can be used only by connecting from the local host. It is granted the `RELOAD` and `PROCESS` administrative privileges. These privileges allow the `admin` user to execute the `mysqladmin reload`, `mysqladmin refresh`, and `mysqladmin flush-xxx` commands, as well as `mysqladmin processlist`. No privileges are granted for accessing any databases. You could add such privileges later by issuing additional `GRANT` statements.
- One account has a username of `dummy` and no password. This account can be used only by connecting from the local host. No privileges are granted. The `USAGE` privilege in the `GRANT` statement enables you to create an account without giving it any privileges. It has the effect of setting all the global privileges to `'N'`. It is assumed that you will grant specific privileges to the account later.

As an alternative to `GRANT`, you can create the same accounts directly by issuing `INSERT` statements and then telling the server to reload the grant tables:

```
shell> mysql --user=root mysql
mysql> INSERT INTO user
-> VALUES('localhost','monty',PASSWORD('some_pass'),
-> 'Y','Y','Y','Y','Y','Y','Y','Y','Y','Y','Y','Y','Y');
mysql> INSERT INTO user
-> VALUES('%','monty',PASSWORD('some_pass'),
-> 'Y','Y','Y','Y','Y','Y','Y','Y','Y','Y','Y','Y','Y');
mysql> INSERT INTO user SET Host='localhost',User='admin',
-> Reload_priv='Y', Process_priv='Y';
mysql> INSERT INTO user (Host,User>Password)
-> VALUES('localhost','dummy','');
mysql> FLUSH PRIVILEGES;
```

The reason for using `FLUSH PRIVILEGES` when you create accounts with `INSERT` is to tell the server to re-read the grant tables. Otherwise, the changes go unnoticed until you restart the server. With `GRANT`, `FLUSH PRIVILEGES` is unnecessary.

The reason for using the `PASSWORD()` function with `INSERT` is to encrypt the password. The `GRANT` statement encrypts the password for you, so `PASSWORD()` is unnecessary.

The 'Y' values enable privileges for the accounts. Depending on your MySQL version, you may have to use a different number of 'Y' values in the first two `INSERT` statements. (Versions prior to 3.22.11 have fewer privilege columns, and versions from 4.0.2 on have more.) For the `admin` account, the more readable extended `INSERT` syntax using `SET` that is available starting with MySQL 3.22.11 is used.

In the `INSERT` statement for the `dummy` account, only the `Host`, `User`, and `Password` columns in the `user` table row are assigned values. None of the privilege columns are set explicitly, so MySQL assigns them all the default value of 'N'. This is equivalent to what `GRANT USAGE` does.

Note that to set up a superuser account, it is necessary only to create a `user` table entry with all of the privilege columns set to 'Y'. `user` table privileges are global, so no entries in any of the other grant tables are needed.

The next examples create three accounts and give them access to specific databases. Each of them has a username of `custom` and password of `obscure`.

To create the accounts with `GRANT`, use the following statements:

```
shell> mysql --user=root mysql
mysql> GRANT SELECT,INSERT,UPDATE,DELETE,CREATE,DROP
-> ON bankaccount.*
-> TO 'custom'@'localhost'
-> IDENTIFIED BY 'obscure';
mysql> GRANT SELECT,INSERT,UPDATE,DELETE,CREATE,DROP
-> ON expenses.*
-> TO 'custom'@'whitehouse.gov'
-> IDENTIFIED BY 'obscure';
mysql> GRANT SELECT,INSERT,UPDATE,DELETE,CREATE,DROP
-> ON customer.*
-> TO 'custom'@'server.domain'
-> IDENTIFIED BY 'obscure';
```

The three accounts can be used as follows:

- The first account can access the `bankaccount` database, but only from the local host.
- The second account can access the `expenses` database, but only from the host `whitehouse.gov`.
- The third account can access the `customer` database, but only from the host `server.domain`.

To set up the `custom` accounts without `GRANT`, use `INSERT` statements as follows to modify the grant tables directly:

```
shell> mysql --user=root mysql
mysql> INSERT INTO user (Host,User,Password)
-> VALUES ('localhost','custom',PASSWORD('obscure'));
mysql> INSERT INTO user (Host,User,Password)
-> VALUES ('whitehouse.gov','custom',PASSWORD('obscure'));
mysql> INSERT INTO user (Host,User,Password)
-> VALUES ('server.domain','custom',PASSWORD('obscure'));
mysql> INSERT INTO db
-> (Host,Db,User,Select_priv,Insert_priv,
-> Update_priv,Delete_priv,Create_priv,Drop_priv)
-> VALUES ('localhost','bankaccount','custom',
-> 'Y','Y','Y','Y','Y','Y');
mysql> INSERT INTO db
-> (Host,Db,User,Select_priv,Insert_priv,
-> Update_priv,Delete_priv,Create_priv,Drop_priv)
-> VALUES ('whitehouse.gov','expenses','custom',
-> 'Y','Y','Y','Y','Y','Y');
mysql> INSERT INTO db
-> (Host,Db,User,Select_priv,Insert_priv,
-> Update_priv,Delete_priv,Create_priv,Drop_priv)
-> VALUES ('server.domain','customer','custom',
-> 'Y','Y','Y','Y','Y','Y');
mysql> FLUSH PRIVILEGES;
```

The first three `INSERT` statements add `user` table entries that allow the user `custom` to connect from the various hosts with the given password, but grant no global privileges (all privileges are set to the default value of 'N'). The next three `INSERT` statements add `db` table entries that grant privileges to `custom` for the `bankaccount`, `expenses`, and `customer` databases, but only when accessed from the proper hosts. As usual when you modify the grant tables directly, you must tell the server to reload them with `FLUSH PRIVILEGES` so that the privilege changes take effect.

If you want to give a specific user access from all machines in a given domain (for example, `mydomain.com`), you can issue a `GRANT` statement that uses the '%' wildcard character in the host part of the account name:

```
mysql> GRANT ...
-> ON *.*
-> TO 'myname'@'%.mydomain.com'
-> IDENTIFIED BY 'mypass';
```

To do the same thing by modifying the grant tables directly, do this:

```
mysql> INSERT INTO user (Host,User,Password,...)
-> VALUES ('%.mydomain.com','myname',PASSWORD('mypass'),...);
mysql> FLUSH PRIVILEGES;
```

### 5.8.3. Removing User Accounts from MySQL

To remove an account, use the `DROP USER` statement, which was added in MySQL 4.1.1. For older versions of MySQL, use `DELETE` instead. The account removal procedure is described in [Section 13.5.1.1, “DROP USER Syntax”](#).

### 5.8.4. Limiting Account Resources

Before MySQL 4.0.2, the only available method for limiting use of MySQL server resources is to set the `max_user_connections` system variable to a non-zero value. But that method is strictly global. It does not allow for management of individual accounts. Also, it limits only the number of simultaneous connections made using a single account, not what a client can do once connected. Both types of control are interest to many MySQL administrators, particularly those for Internet Service Providers.

Starting from MySQL 4.0.2, you can limit the following server resources for individual accounts:

- The number of queries that an account can issue per hour
- The number of updates that an account can issue per hour
- The number of times an account can connect to the server per hour

Any statement that a client can issue counts against the query limit. Only statements that modify databases or tables count against the update limit.

An account in this context is a single row in the `user` table. Each account is uniquely identified by its `User` and `Host` column values.

As a prerequisite for using this feature, the `user` table in the `mysql` database must contain the resource-related columns. Resource limits are stored in the `max_questions`, `max_updates`, `max_connections`, and `max_user_connections` columns. If your `user` table does not have these columns, it must be upgraded; see [Section 5.5.1, “mysql\\_fix\\_privilege\\_tables — Upgrade MySQL System Tables”](#).

To set resource limits with a `GRANT` statement, use a `WITH` clause that names each resource to be limited and a per-hour count indicating the limit value. For example, to create a new account that can access the `customer` database, but only in a limited fashion, issue this statement:

```
mysql> GRANT ALL ON customer.* TO 'francis'@'localhost'  
-> IDENTIFIED BY 'frank'  
-> WITH MAX_QUERIES_PER_HOUR 20  
->      MAX_UPDATES_PER_HOUR 10  
->      MAX_CONNECTIONS_PER_HOUR 5;
```

The limit types need not all be named in the `WITH` clause, but those named can be present in any order. The value for each per-hour limit should be an integer representing a count per hour. If the `GRANT` statement has no `WITH` clause, the limits are each set to the default value of zero (that is, no limit).

To set or change limits for an existing account, use a `GRANT USAGE` statement at the global level (`ON *.*`). The following statement changes the query limit for `francis` to 100:

```
mysql> GRANT USAGE ON *.* TO 'francis'@'localhost'  
-> WITH MAX_QUERIES_PER_HOUR 100;
```

This statement leaves the account's existing privileges unchanged and modifies only the limit values specified.

To remove an existing limit, set its value to zero. For example, to remove the limit on how many times per hour `francis` can connect, use this statement:

```
mysql> GRANT USAGE ON *.* TO 'francis'@'localhost'  
-> WITH MAX_CONNECTIONS_PER_HOUR 0;
```

Resource-use counting takes place when any account has a non-zero limit placed on its use of any of the resources.

As the server runs, it counts the number of times each account uses resources. If an account reaches its limit on number of connections within the last hour, further connections for the account are rejected until that hour is up. Similarly, if the account reaches its limit on the number of queries or updates, further queries or updates are rejected until the hour is up. In all such cases, an appropriate error message is issued.

The current per-hour resource-use counts can be reset globally for all accounts, or individually for a given account:

- To reset the current counts to zero for all accounts, issue a `FLUSH USER_RESOURCES` statement. The counts also can be reset by reloading the grant tables (for example, with a `FLUSH PRIVILEGES` statement or a `mysqladmin reload` command).
- The counts for an individual account can be set to zero by re-granting it any of its limits. To do this, use `GRANT USAGE` as described earlier and specify a limit value equal to the value that the account currently has.

Counter resets do not affect the `MAX_USER_CONNECTIONS` limit.

All counts begin at zero when the server starts; counts are not carried over through a restart.

## 5.8.5. Assigning Account Passwords

Passwords may be assigned from the command line by using the `mysqladmin` command:

```
shell> mysqladmin -u user_name -h host_name password "newpwd"
```

The account for which this command resets the password is the one with a `user` table row that matches `user_name` in the `User` column and the client host *from which you connect* in the `Host` column.

Another way to assign a password to an account is to issue a `SET PASSWORD` statement:

```
mysql> SET PASSWORD FOR 'jeffrey'@'%' = PASSWORD('biscuit');
```

Only users such as `root` that have update access to the `mysql` database can change the password for other users. If you are not connected as an anonymous user, you can change your own password by omitting the `FOR` clause:

```
mysql> SET PASSWORD = PASSWORD('biscuit');
```

You can also use a `GRANT USAGE` statement at the global level (`ON *.*`) to assign a password to an account without affecting the account's current privileges:

```
mysql> GRANT USAGE ON *.* TO 'jeffrey'@'%' IDENTIFIED BY 'biscuit';
```

Although it is generally preferable to assign passwords using one of the preceding methods, you can also do so by modifying the `user` table directly:

- To establish a password when creating a new account, provide a value for the `Password` column:

```
shell> mysql -u root mysql
mysql> INSERT INTO user (Host,User>Password)
-> VALUES('%','jeffrey',PASSWORD('biscuit'));
mysql> FLUSH PRIVILEGES;
```

- To change the password for an existing account, use `UPDATE` to set the `Password` column value:

```
shell> mysql -u root mysql
mysql> UPDATE user SET Password = PASSWORD('bagel')
-> WHERE Host = '%' AND User = 'francis';
mysql> FLUSH PRIVILEGES;
```

When you assign an account a non-empty password using `SET PASSWORD`, `INSERT`, or `UPDATE`, you must use the `PASSWORD()` function to encrypt it. `PASSWORD()` is necessary because the `user` table stores passwords in encrypted form, not as plaintext. If you forget that fact, you are likely to set passwords like this:

```
shell> mysql -u root mysql
mysql> INSERT INTO user (Host,User,Password)
-> VALUES('%','jeffrey','biscuit');
mysql> FLUSH PRIVILEGES;
```

The result is that the literal value `'biscuit'` is stored as the password in the `user` table, not the encrypted value. When `jeffrey` attempts to connect to the server using this password, the value is encrypted and compared to the value stored in the `user` table. However, the stored value is the literal string `'biscuit'`, so the comparison fails and the server rejects the connection:

```
shell> mysql -u jeffrey -pbiscuit test
Access denied
```

If you assign passwords using the `GRANT ... IDENTIFIED BY` statement or the `mysqladmin password` command, they both take care of encrypting the password for you. The `PASSWORD()` function is unnecessary.

**Note:** `PASSWORD()` encryption is different from Unix password encryption. See [Section 5.8.1, “MySQL Usernames and Passwords”](#).

## 5.8.6. Keeping Your Password Secure

On an administrative level, you should never grant access to the `user` grant table to any non-administrative accounts. Passwords in the `user` table are stored in encrypted form, but in versions of MySQL earlier than 4.1, knowing the encrypted password for an account makes it possible to connect to the server using that account.

When you run a client program to connect to the MySQL server, it is inadvisable to specify your password in a way that exposes it to discovery by other users. The methods you can use to specify your password when you run client programs are listed here, along with an assessment of the risks of each method:

- Use a `-pyour_pass` or `--password=your_pass` option on the command line. For example:

```
shell> mysql -u francis -pfrank db_name
```

This is convenient *but insecure*, because your password becomes visible to system status programs such as `ps` that may be invoked by other users to display command lines. MySQL clients typically overwrite the command-line password argument with zeros during their initialization sequence. However, there is still a brief interval during which the value is visible. On some systems this strategy is ineffective, anyway, and the password remains visible to `ps`. (SystemV Unix systems and perhaps others are subject to this problem.)

- Use the `-p` or `--password` option with no password value specified. In this case, the client program solicits the password from the terminal:

```
shell> mysql -u francis -p db_name
Enter password: *****
```

The `*` characters indicate where you enter your password. The password is not displayed as you enter it.



It is more secure to enter your password this way than to specify it on the command line because it is not visible to other users. However, this method of entering a password is suitable only for programs that you run interactively. If you want to invoke a client from a script that runs non-interactively, there is no opportunity to enter the password from the terminal. On some systems, you may even find that the first line of your script is read and interpreted (incorrectly) as your password.

- Store your password in an option file. For example, on Unix you can list your password in the `[client]` section of the `.my.cnf` file in your home directory:

```
[client]
password=your_pass
```

If you store your password in `.my.cnf`, the file should not be accessible to anyone but yourself. To ensure this, set the file access mode to `400` or `600`. For example:

```
shell> chmod 600 .my.cnf
```

[Section 4.3.2, “Using Option Files”](#), discusses option files in more detail.

- Store your password in the `MYSQL_PWD` environment variable. This method of specifying your MySQL password must be considered *extremely insecure* and should not be used. Some versions of `ps` include an option to display the environment of running processes. If you set `MYSQL_PWD`, your password is exposed to any other user who runs `ps`. Even on systems without such a version of `ps`, it is unwise to assume that there are no other methods by which users can examine process environments. See [Appendix F, Environment Variables](#).

All in all, the safest methods are to have the client program prompt for the password or to specify the password in a properly protected option file.

## 5.8.7. Using Secure Connections

Beginning with version 4.0.0, MySQL has support for secure (encrypted) connections between MySQL clients and the server using the Secure Sockets Layer (SSL) protocol. This section discusses how to use SSL connections. For information on requiring users to use SSL connections, see [Section 13.5.1.2, “GRANT Syntax”](#).

The standard configuration of MySQL is intended to be as fast as possible, so encrypted connections are not used by default. Doing so would make the client/server protocol much slower. Encrypting data is a CPU-intensive operation that requires the computer to do additional work and can delay other MySQL tasks. For applications that require the security provided by encrypted connections, the extra computation is warranted.

MySQL allows encryption to be enabled on a per-connection basis. You can choose a normal unencrypted connection or a secure encrypted SSL connection according the requirements of individual applications.

Secure connections are based on the OpenSSL API and are available through the MySQL C API. Replication uses the C API, so secure connections can be used between master and slave servers.

### 5.8.7.1. Basic SSL Concepts

To understand how MySQL uses SSL, it is necessary to explain some basic SSL and X509 concepts. People who are familiar with these can skip this part of the discussion.

By default, MySQL uses unencrypted connections between the client and the server. This means that

someone with access to the network could watch all your traffic and look at the data being sent or received. They could even change the data while it is in transit between client and server. To improve security a little, you can compress client/server traffic by using the `--compress` option when invoking client programs. However, this does not foil a determined attacker.

When you need to move information over a network in a secure fashion, an unencrypted connection is unacceptable. Encryption is the way to make any kind of data unreadable. In fact, today's practice requires many additional security elements from encryption algorithms. They should resist many kind of known attacks such as changing the order of encrypted messages or replaying data twice.

SSL is a protocol that uses different encryption algorithms to ensure that data received over a public network can be trusted. It has mechanisms to detect any data change, loss, or replay. SSL also incorporates algorithms that provide identity verification using the X509 standard.

X509 makes it possible to identify someone on the Internet. It is most commonly used in e-commerce applications. In basic terms, there should be some company called a “Certificate Authority” (or CA) that assigns electronic certificates to anyone who needs them. Certificates rely on asymmetric encryption algorithms that have two encryption keys (a public key and a secret key). A certificate owner can show the certificate to another party as proof of identity. A certificate consists of its owner's public key. Any data encrypted with this public key can be decrypted only using the corresponding secret key, which is held by the owner of the certificate.

If you need more information about SSL, X509, or encryption, use your favorite Internet search engine to search for the keywords in which you are interested.

### 5.8.7.2. Using SSL Connections

To use SSL connections between the MySQL server and client programs, your system must support OpenSSL and your version of MySQL must be 4.0.0 or newer and built with SSL support.

To get secure connections to work with MySQL and SSL, you must do the following:

1. Install the OpenSSL library if it has not already been installed. We have tested MySQL with OpenSSL 0.9.6. To obtain OpenSSL, visit <http://www.openssl.org>.
2. If you are not using a binary (precompiled) version of MySQL that has been built with SSL support, configure a MySQL source distribution to use SSL. When you configure MySQL, invoke the `configure` script with the `--with-vio` and `--with-openssl` options:

```
shell> ./configure --with-vio --with-openssl
```

3. Make sure that you have upgraded your grant tables to include the SSL-related columns in the `mysql.user` table. This is necessary if your grant tables date from a version of MySQL older than 4.0. The upgrade procedure is described in [Section 5.5.1](#), “`mysql_fix_privilege_tables` — Upgrade MySQL System Tables”.
4. To check whether a server binary is compiled with SSL support, invoke it with the `--ssl` option. An error will occur if the server does not support SSL:

```
shell> mysqld --ssl --help
060525 14:18:52 [ERROR] mysqld: unknown option '--ssl'
```

To check whether a running `mysqld` server supports SSL, examine the value of the `have_openssl` system variable:

```
mysql> SHOW VARIABLES LIKE 'have_openssl';
+-----+-----+
| Variable_name | Value |
+-----+-----+
```

```
+-----+-----+
| have_openssl | YES |
+-----+-----+
```

If the value is `YES`, the server supports OpenSSL connections.

To start the MySQL server so that it allows clients to connect via SSL, use the options that identify the key and certificate files the server needs when establishing a secure connection:

```
shell> mysqld --ssl-ca=cacert.pem \
--ssl-cert=server-cert.pem \
--ssl-key=server-key.pem
```

- `--ssl-ca` identifies the Certificate Authority (CA) certificate.
- `--ssl-cert` identifies the server public key. This can be sent to the client and authenticated against the CA certificate that it has.
- `--ssl-key` identifies the server private key.

To establish a secure connection to a MySQL server with yaSSL support, start a client like this:

```
shell> mysql --ssl-ca=cacert.pem \
--ssl-cert=client-cert.pem \
--ssl-key=client-key.pem
```

In other words, the options are similar to those used for the server. Note that the Certificate Authority certificate has to be the same.

A client can determine whether the current connection with the server uses SSL by checking the value of the `Ssl_cipher` status variable. The value of `Ssl_cipher` is non-empty if SSL is used, and empty otherwise. For example:

```
mysql> SHOW STATUS LIKE 'Ssl_cipher';
+-----+-----+
| Variable_name | Value |
+-----+-----+
| Ssl_cipher    | DHE-RSA-AES256-SHA |
+-----+-----+
```

For the `mysql` client, you can use the `STATUS` or `\s` command and check the `SSL` line:

```
mysql> \s
...
SSL:                Not in use
...
```

Or:

```
mysql> \s
...
SSL:                Cipher in use is DHE-RSA-AES256-SHA
...
```

To establish a secure connection from within an application program, use the `mysql_ssl_set()` C API function to set the appropriate certificate options before calling `mysql_real_connect()`. See [Section 17.2.3.65](#), “`mysql_ssl_set()`”.

### 5.8.7.3. SSL Command Options

The following list describes options that are used for specifying the use of SSL, certificate files, and key files. These options are available beginning with MySQL 4.0. They can be given on the command line or in an option file.

These options are not available unless MySQL has been built with SSL support. See [Section 5.8.7.2, “Using SSL Connections”](#).

- `--ssl`

For the server, this option specifies that the server allows SSL connections. For a client program, it allows the client to connect to the server using SSL. This option is not sufficient in itself to cause an SSL connection to be used. You must also specify the `--ssl-ca`, `--ssl-cert`, and `--ssl-key` options.

This option is more often used in its opposite form to override any other SSL options and indicate that SSL should *not* be used. To do this, specify the option as `--skip-ssl` or `--ssl=0`.

Note that use of `--ssl` does not *require* an SSL connection. For example, if the server or client is compiled without SSL support, a normal unencrypted connection is used.

The secure way to ensure that an SSL connection is used is to create an account on the server that includes a `REQUIRE SSL` clause in the `GRANT` statement. Then use this account to connect to the server, with both a server and client that have SSL support enabled.

- `--ssl-ca=file_name`

The path to a file with a list of trusted SSL CAs.

- `--ssl-capath=directory_name`

The path to a directory that contains trusted SSL CA certificates in PEM format.

- `--ssl-cert=file_name`

The name of the SSL certificate file to use for establishing a secure connection.

- `--ssl-cipher=cipher_list`

A list of allowable ciphers to use for SSL encryption. `cipher_list` has the same format as the `openssl ciphers` command.

Example: `--ssl-cipher=ALL:-AES:-EXP`

- `--ssl-key=file_name`

The name of the SSL key file to use for establishing a secure connection.

### 5.8.7.4. Setting Up SSL Certificates for MySQL

Here is an example of setting up SSL certificates for MySQL using OpenSSL:

```
DIR=`pwd`/openssl
PRIV=$DIR/private

mkdir $DIR $PRIV $DIR/newcerts
cp /usr/share/ssl/openssl.cnf $DIR
```

```
replace ./demoCA $DIR -- $DIR/openssl.cnf

# Create necessary files: $database, $serial and $new_certs_dir
# directory (optional)

touch $DIR/index.txt
echo "01" > $DIR/serial

#
# Generation of Certificate Authority(CA)
#

openssl req -new -x509 -keyout $PRIV/cakey.pem -out $DIR/cacert.pem \
    -config $DIR/openssl.cnf

# Sample output:
# Using configuration from /home/monty/openssl/openssl.cnf
# Generating a 1024 bit RSA private key
# .....+++++
# .....+++++
# writing new private key to '/home/monty/openssl/private/cakey.pem'
# Enter PEM pass phrase:
# Verifying password - Enter PEM pass phrase:
# -----
# You are about to be asked to enter information that will be
# incorporated into your certificate request.
# What you are about to enter is what is called a Distinguished Name
# or a DN.
# There are quite a few fields but you can leave some blank
# For some fields there will be a default value,
# If you enter '.', the field will be left blank.
# -----
# Country Name (2 letter code) [AU]:FI
# State or Province Name (full name) [Some-State]:.
# Locality Name (eg, city) []:
# Organization Name (eg, company) [Internet Widgits Pty Ltd]:MySQL AB
# Organizational Unit Name (eg, section) []:
# Common Name (eg, YOUR name) []:MySQL admin
# Email Address []:

#
# Create server request and key
#
openssl req -new -keyout $DIR/server-key.pem -out \
    $DIR/server-req.pem -days 3600 -config $DIR/openssl.cnf

# Sample output:
# Using configuration from /home/monty/openssl/openssl.cnf
# Generating a 1024 bit RSA private key
# ..+++++
# .....+++++
# writing new private key to '/home/monty/openssl/server-key.pem'
# Enter PEM pass phrase:
# Verifying password - Enter PEM pass phrase:
# -----
# You are about to be asked to enter information that will be
# incorporated into your certificate request.
# What you are about to enter is what is called a Distinguished Name
# or a DN.
# There are quite a few fields but you can leave some blank
# For some fields there will be a default value,
# If you enter '.', the field will be left blank.
# -----
# Country Name (2 letter code) [AU]:FI
# State or Province Name (full name) [Some-State]:.
# Locality Name (eg, city) []:
# Organization Name (eg, company) [Internet Widgits Pty Ltd]:MySQL AB
# Organizational Unit Name (eg, section) []:
# Common Name (eg, YOUR name) []:MySQL server
# Email Address []:

#
# Please enter the following 'extra' attributes
# to be sent with your certificate request
# A challenge password []:
# An optional company name []:

#
# Remove the passphrase from the key (optional)
#
```

```
openssl rsa -in $DIR/server-key.pem -out $DIR/server-key.pem

#
# Sign server cert
#
openssl ca -policy policy_anything -out $DIR/server-cert.pem \
  -config $DIR/openssl.cnf -infile $DIR/server-req.pem

# Sample output:
# Using configuration from /home/monty/openssl/openssl.cnf
# Enter PEM pass phrase:
# Check that the request matches the signature
# Signature ok
# The Subjects Distinguished Name is as follows
# countryName          :PRINTABLE:'FI'
# organizationName     :PRINTABLE:'MySQL AB'
# commonName           :PRINTABLE:'MySQL admin'
# Certificate is to be certified until Sep 13 14:22:46 2003 GMT
# (365 days)
# Sign the certificate? [y/n]:y
#
#
# 1 out of 1 certificate requests certified, commit? [y/n]y
# Write out database with 1 new entries
# Data Base Updated

#
# Create client request and key
#
openssl req -new -keyout $DIR/client-key.pem -out \
  $DIR/client-req.pem -days 3600 -config $DIR/openssl.cnf

# Sample output:
# Using configuration from /home/monty/openssl/openssl.cnf
# Generating a 1024 bit RSA private key
# .....++++++
# .....++++++
# writing new private key to '/home/monty/openssl/client-key.pem'
# Enter PEM pass phrase:
# Verifying password - Enter PEM pass phrase:
# -----
# You are about to be asked to enter information that will be
# incorporated into your certificate request.
# What you are about to enter is what is called a Distinguished Name
# or a DN.
# There are quite a few fields but you can leave some blank
# For some fields there will be a default value,
# If you enter '.', the field will be left blank.
# -----
# Country Name (2 letter code) [AU]:FI
# State or Province Name (full name) [Some-State]:.
# Locality Name (eg, city) []:
# Organization Name (eg, company) [Internet Widgits Pty Ltd]:MySQL AB
# Organizational Unit Name (eg, section) []:
# Common Name (eg, YOUR name) []:MySQL user
# Email Address []:
#
# Please enter the following 'extra' attributes
# to be sent with your certificate request
# A challenge password []:
# An optional company name []:

#
# Remove a passphrase from the key (optional)
#
openssl rsa -in $DIR/client-key.pem -out $DIR/client-key.pem

#
# Sign client cert
#
openssl ca -policy policy_anything -out $DIR/client-cert.pem \
  -config $DIR/openssl.cnf -infile $DIR/client-req.pem

# Sample output:
# Using configuration from /home/monty/openssl/openssl.cnf
# Enter PEM pass phrase:
# Check that the request matches the signature
# Signature ok
# The Subjects Distinguished Name is as follows
```

```
# countryName          :PRINTABLE:'FI'
# organizationName    :PRINTABLE:'MySQL AB'
# commonName          :PRINTABLE:'MySQL user'
# Certificate is to be certified until Sep 13 16:45:17 2003 GMT
# (365 days)
# Sign the certificate? [y/n]:y
#
#
# 1 out of 1 certificate requests certified, commit? [y/n]y
# Write out database with 1 new entries
# Data Base Updated

#
# Create a my.cnf file that you can use to test the certificates
#

cnf=" "
cnf="$cnf [client]"
cnf="$cnf ssl-ca=$DIR/cacert.pem"
cnf="$cnf ssl-cert=$DIR/client-cert.pem"
cnf="$cnf ssl-key=$DIR/client-key.pem"
cnf="$cnf [mysqld]"
cnf="$cnf ssl-ca=$DIR/cacert.pem"
cnf="$cnf ssl-cert=$DIR/server-cert.pem"
cnf="$cnf ssl-key=$DIR/server-key.pem"
echo $cnf | replace " " '
' > $DIR/my.cnf
```

To test SSL connections, start the server as follows, where `$DIR` is the pathname to the directory where the sample `my.cnf` option file is located:

```
shell> mysqld --defaults-file=$DIR/my.cnf &
```

Then invoke a client program using the same option file:

```
shell> mysql --defaults-file=$DIR/my.cnf
```

If you have a MySQL source distribution, you can also test your setup by modifying the preceding `my.cnf` file to refer to the demonstration certificate and key files in the `SSL` directory of the distribution.

### 5.8.7.5. Connecting to MySQL Remotely from Windows with SSH

Here is a note that describes how to get a secure connection to a remote MySQL server with SSH (by David Carlson <[dcarlson@mplcomm.com](mailto:dcarlson@mplcomm.com)>):

1. Install an SSH client on your Windows machine. As a user, the best non-free one I have found is from `SecureCRT` from <http://www.vandyke.com/>. Another option is `f-secure` from <http://www.f-secure.com/>. You can also find some free ones on Google at [http://directory.google.com/Top/Computers/Security/Products\\_and\\_Tools/Cryptography/SSH/Clients/Windows/](http://directory.google.com/Top/Computers/Security/Products_and_Tools/Cryptography/SSH/Clients/Windows/).
2. Start your Windows SSH client. Set `Host_Name = yourmysqlserver_URL_or_IP`. Set `userid=your_userid` to log in to your server. This `userid` value might not be the same as the username of your MySQL account.
3. Set up port forwarding. Either do a remote forward (Set `local_port: 3306`, `remote_host: yourmysqlservername_or_ip`, `remote_port: 3306`) or a local forward (Set `port: 3306`, `host: localhost`, `remote port: 3306`).
4. Save everything, otherwise you will have to redo it the next time.
5. Log in to your server with the SSH session you just created.

6. On your Windows machine, start some ODBC application (such as Access).
7. Create a new file in Windows and link to MySQL using the ODBC driver the same way you normally do, except type in `localhost` for the MySQL host server, not `yourmysqlservername`.

At this point, you should have an ODBC connection to MySQL, encrypted using SSH.

## 5.9. Backup and Recovery

This section discusses how to make database backups (full and incremental) and how to perform table maintenance. The syntax of the SQL statements described here is given in [Chapter 13, \*SQL Statement Syntax\*](#). Much of the information here pertains primarily to [MyISAM](#) tables. Additional information about [InnoDB](#) backup procedures is given in [Section 14.2.9, “Backing Up and Recovering an InnoDB Database”](#).

### 5.9.1. Database Backups

Because MySQL tables are stored as files, it is easy to do a backup. To get a consistent backup, do a `LOCK TABLES` on the relevant tables, followed by `FLUSH TABLES` for the tables. See [Section 13.4.5, “LOCK TABLES and UNLOCK TABLES Syntax”](#), and [Section 13.5.5.2, “FLUSH Syntax”](#). You need only a read lock; this allows other clients to continue to query the tables while you are making a copy of the files in the database directory. The `FLUSH TABLES` statement is needed to ensure that the all active index pages are written to disk before you start the backup.

To make an SQL-level backup of a table, you can use `SELECT INTO ... OUTFILE`. For this statement, the output file cannot already exist because allowing files to be overwritten would constitute a security risk. See [Section 13.2.7, “SELECT Syntax”](#).

Another technique for backing up a database is to use the `mysqldump` program or the `mysqlhotcopy` script. See [Section 8.12, “mysqldump — A Database Backup Program”](#), and [Section 8.13, “mysqlhotcopy — A Database Backup Program”](#).

1. Create a full backup of your database:

```
shell> mysqldump --tab=/path/to/some/dir --opt db_name
```

Or:

```
shell> mysqlhotcopy db_name /path/to/some/dir
```

You can also create a binary backup simply by copying all table files (`*.frm`, `*.MYD`, and `*.MYI` files), as long as the server isn't updating anything. The `mysqlhotcopy` script uses this method. (But note that these methods do not work if your database contains [InnoDB](#) tables. [InnoDB](#) does not store table contents in database directories, and `mysqlhotcopy` works only for [MyISAM](#) and [ISAM](#) tables.)

2. Stop `mysqld` if it is running, then start it with the `--log-bin[=file_name]` option. See [Section 5.11.4, “The Binary Log”](#). The binary log files provide you with the information you need to replicate changes to the database that are made subsequent to the point at which you executed `mysqldump`.

For [InnoDB](#) tables, it is possible to perform an online backup that takes no locks on tables; see [Section 8.12, “mysqldump — A Database Backup Program”](#).



MySQL supports incremental backups: You need to start the server with the `--log-bin` option to enable binary logging; see [Section 5.11.4, “The Binary Log”](#). At the moment you want to make an incremental backup (containing all changes that happened since the last full or incremental backup), you should rotate the binary log by using `FLUSH LOGS`. This done, you need to copy to the backup location all binary logs which range from the one of the moment of the last full or incremental backup to the last but one. These binary logs are the incremental backup; at restore time, you apply them as explained further below. The next time you do a full backup, you should also rotate the binary log using `FLUSH LOGS`, `mysqldump --flush-logs`, or `mysqlhotcopy --flushlog`. See [Section 8.12, “mysqldump — A Database Backup Program”](#), and [Section 8.13, “mysqlhotcopy — A Database Backup Program”](#).

If your MySQL server is a slave replication server, then regardless of the backup method you choose, you should also back up the `master.info` and `relay-log.info` files when you back up your slave's data. These files are always needed to resume replication after you restore the slave's data. If your slave is subject to replicating `LOAD DATA INFILE` commands, you should also back up any `SQL_LOAD-*` files that may exist in the directory specified by the `--slave-load-tmpdir` option. (This location defaults to the value of the `tmpdir` variable if not specified.) The slave needs these files to resume replication of any interrupted `LOAD DATA INFILE` operations.

If you have to restore `MyISAM` tables, try to recover them using `REPAIR TABLE` or `myisamchk -r` first. That should work in 99.9% of all cases. If `myisamchk` fails, try the following procedure. Note that it works only if you have enabled binary logging by starting MySQL with the `--log-bin` option.

1. Restore the original `mysqldump` backup, or binary backup.
2. Execute the following command to re-run the updates in the binary logs:

```
shell> mysqlbinlog binlog.[0-9]* | mysql
```

In some cases, you may want to re-run only certain binary logs, from certain positions (usually you want to re-run all binary logs from the date of the restored backup, excepting possibly some incorrect statements). See [Section 8.10, “mysqlbinlog — Utility for Processing Binary Log Files”](#), for more information on the `mysqlbinlog` utility and how to use it.

If you are using the update logs instead, you can process their contents like this:

```
shell> ls -l -t -r hostname.[0-9]* | xargs cat | mysql
```

`ls` is used to sort the update log filenames into the right order.

You can also make selective backups of individual files:

- To dump the table, use `SELECT * INTO OUTFILE 'file_name' FROM tbl_name`.
- To reload the table, use `LOAD DATA INFILE 'file_name' REPLACE . . .`. To avoid duplicate rows, the table must have a `PRIMARY KEY` or a `UNIQUE` index. The `REPLACE` keyword causes old rows to be replaced with new ones when a new row duplicates an old row on a unique key value.

If you have performance problems with your server while making backups, one strategy that can help is to set up replication and perform backups on the slave rather than on the master. See [Section 6.1, “Introduction to Replication”](#).

If you are using a Veritas filesystem, you can make a backup like this:

1. From a client program, execute `FLUSH TABLES WITH READ LOCK`.
2. From another shell, execute `mount vxfs snapshot`.
3. From the first client, execute `UNLOCK TABLES`.
4. Copy files from the snapshot.
5. Unmount the snapshot.

## 5.9.2. Example Backup and Recovery Strategy

This section discusses a procedure for performing backups that allows you to recover data after several types of crashes:

- Operating system crash
- Power failure
- Filesystem crash
- Hardware problem (hard drive, motherboard, and so forth)

The following instructions assume a minimum version of MySQL 4.1.8, because some `mysqldump` options used here are not available in earlier versions.

The example commands do not include options such as `--user` and `--password` for the `mysqldump` and `mysql` programs. You should include such options as necessary so that the MySQL server allows you to connect to it.

We assume that data is stored in the `InnoDB` storage engine, which has support for transactions and automatic crash recovery. We also assume that the MySQL server is under load at the time of the crash. If it were not, no recovery would ever be needed.

For cases of operating system crashes or power failures, we can assume that MySQL's disk data is available after a restart. The `InnoDB` data files might not contain consistent data due to the crash, but `InnoDB` reads its logs and finds in them the list of pending committed and non-committed transactions that have not been flushed to the data files. `InnoDB` automatically rolls back those transactions that were not committed, and flushes to its data files those that were committed. Information about this recovery process is conveyed to the user through the MySQL error log. The following is an example log excerpt:

```
InnoDB: Database was not shut down normally.
InnoDB: Starting recovery from log files...
InnoDB: Starting log scan based on checkpoint at
InnoDB: log sequence number 0 13674004
InnoDB: Doing recovery: scanned up to log sequence number 0 13739520
InnoDB: Doing recovery: scanned up to log sequence number 0 13805056
InnoDB: Doing recovery: scanned up to log sequence number 0 13870592
InnoDB: Doing recovery: scanned up to log sequence number 0 13936128
...
InnoDB: Doing recovery: scanned up to log sequence number 0 20555264
InnoDB: Doing recovery: scanned up to log sequence number 0 20620800
InnoDB: Doing recovery: scanned up to log sequence number 0 20664692
InnoDB: 1 uncommitted transaction(s) which must be rolled back
InnoDB: Starting rollback of uncommitted transactions
InnoDB: Rolling back trx no 16745
InnoDB: Rolling back of trx no 16745 completed
InnoDB: Rollback of uncommitted transactions completed
InnoDB: Starting an apply batch of log records to the database...
InnoDB: Apply batch completed
InnoDB: Started
```

```
mysqld: ready for connections
```

For the cases of filesystem crashes or hardware problems, we can assume that the MySQL disk data is *not* available after a restart. This means that MySQL fails to start successfully because some blocks of disk data are no longer readable. In this case, it is necessary to reformat the disk, install a new one, or otherwise correct the underlying problem. Then it is necessary to recover our MySQL data from backups, which means that we must already have made backups. To make sure that is the case, we should design a backup policy.

### 5.9.2.1. Backup Policy

We all know that backups must be scheduled periodically. A full backups (a snapshot of the data at a point in time) can be done in MySQL with several tools. For example, [InnoDB Hot Backup](#) provides online non-blocking physical backup of the [InnoDB](#) data files, and [mysqldump](#) provides on-line logical backup. This discussion uses [mysqldump](#).

Assume that we make a backup on Sunday at 1 p.m., when load is low. The following command makes a full backup of all our [InnoDB](#) tables in all databases:

```
shell> mysqldump --single-transaction --all-databases > backup_sunday_1_PM.sql
```

This is an online, non-blocking backup that does not disturb the reads and writes on the tables. We assumed earlier that our tables are [InnoDB](#) tables, so `--single-transaction` uses a consistent read and guarantees that data seen by [mysqldump](#) does not change. (Changes made by other clients to [InnoDB](#) tables are not seen by the [mysqldump](#) process.) If we do also have other types of tables, we must assume that they are not changed during the backup. For example, for the [MyISAM](#) tables in the [mysql](#) database, we must assume that no administrative changes are being made to MySQL accounts during the backup.

The resulting `.sql` file produced by [mysqldump](#) contains a set of SQL `INSERT` statements that can be used to reload the dumped tables at a later time.

Full backups are necessary, but they are not always convenient. They produce large backup files and take time to generate. They are not optimal in the sense that each successive full backup includes all data, even that part that has not changed since the previous full backup. After we have made the initial full backup, it is more efficient to make incremental backups. They are smaller and take less time to produce. The tradeoff is that, at recovery time, you cannot restore your data just by reloading the full backup. You must also process the incremental backups to recover the incremental changes.

To make incremental backups, we need to save the incremental changes. The MySQL server should always be started with the `--log-bin` option so that it stores these changes in a file while it updates data. This option enables binary logging, so that the server writes each SQL statement that updates data into a file called a MySQL binary log. Looking at the data directory of a MySQL server that was started with the `--log-bin` option and that has been running for some days, we find these MySQL binary log files:

```
-rw-rw---- 1 guilhem guilhem 1277324 Nov 10 23:59 gbichot2-bin.000001
-rw-rw---- 1 guilhem guilhem      4 Nov 10 23:59 gbichot2-bin.000002
-rw-rw---- 1 guilhem guilhem    79 Nov 11 11:06 gbichot2-bin.000003
-rw-rw---- 1 guilhem guilhem   508 Nov 11 11:08 gbichot2-bin.000004
-rw-rw---- 1 guilhem guilhem 220047446 Nov 12 16:47 gbichot2-bin.000005
-rw-rw---- 1 guilhem guilhem  998412 Nov 14 10:08 gbichot2-bin.000006
-rw-rw---- 1 guilhem guilhem   361 Nov 14 10:07 gbichot2-bin.index
```

Each time it restarts, the MySQL server creates a new binary log file using the next number in the sequence. While the server is running, you can also tell it to close the current binary log file and begin a new one manually by issuing a `FLUSH LOGS` SQL statement or with a [mysqladmin flush-logs](#) command. [mysqldump](#) also has an option to flush the logs. The `.index` file in the data directory con-

tains the list of all MySQL binary logs in the directory. This file is used for replication.

The MySQL binary logs are important for recovery because they form the set of incremental backups. If you make sure to flush the logs when you make your full backup, then any binary log files created afterward contain all the data changes made since the backup. Let's modify the previous `mysqldump` command a bit so that it flushes the MySQL binary logs at the moment of the full backup, and so that the dump file contains the name of the new current binary log:

```
shell> mysqldump --single-transaction --flush-logs --master-data=2 \
--all-databases > backup_sunday_1_PM.sql
```

After executing this command, the data directory contains a new binary log file, `gbichot2-bin.000007`. The resulting `.sql` file includes these lines:

```
-- Position to start replication or point-in-time recovery from
-- CHANGE MASTER TO MASTER_LOG_FILE='gbichot2-bin.000007',MASTER_LOG_POS=4;
```

Because the `mysqldump` command made a full backup, those lines mean two things:

- The `.sql` file contains all changes made before any changes written to the `gbichot2-bin.000007` binary log file or newer.
- All data changes logged after the backup are not present in the `.sql`, but are present in the `gbichot2-bin.000007` binary log file or newer.

On Monday at 1 p.m., we can create an incremental backup by flushing the logs to begin a new binary log file. For example, executing a `mysqladmin flush-logs` command creates `gbichot2-bin.000008`. All changes between the Sunday 1 p.m. full backup and Monday 1 p.m. will be in the `gbichot2-bin.000007` file. This incremental backup is important, so it is a good idea to copy it to a safe place. (For example, back it up on tape or DVD, or copy it to another machine.) On Tuesday at 1 p.m., execute another `mysqladmin flush-logs` command. All changes between Monday 1 p.m. and Tuesday 1 p.m. will be in the `gbichot2-bin.000008` file (which also should be copied somewhere safe).

The MySQL binary logs take up disk space. To free up space, purge them from time to time. One way to do this is by deleting the binary logs that are no longer needed, such as when we make a full backup:

```
shell> mysqldump --single-transaction --flush-logs --master-data=2 \
--all-databases --delete-master-logs > backup_sunday_1_PM.sql
```

**Note:** Deleting the MySQL binary logs with `mysqldump --delete-master-logs` can be dangerous if your server is a replication master server, because slave servers might not yet fully have processed the contents of the binary log. The description for the `PURGE MASTER LOGS` statement explains what should be verified before deleting the MySQL binary logs. See [Section 13.6.1.1, “PURGE MASTER LOGS Syntax”](#).

## 5.9.2.2. Using Backups for Recovery

Now, suppose that we have a catastrophic crash on Wednesday at 8 a.m. that requires recovery from backups. To recover, first we restore the last full backup we have (the one from Sunday 1 p.m.). The full backup file is just a set of SQL statements, so restoring it is very easy:

```
shell> mysql < backup_sunday_1_PM.sql
```

At this point, the data is restored to its state as of Sunday 1 p.m.. To restore the changes made since then, we must use the incremental backups; that is, the `gbichot2-bin.000007` and `gbichot2-bin.000008` binary log files. Fetch the files if necessary from where they were backed up,

and then process their contents like this:

```
shell> mysqlbinlog gbichot2-bin.000007 gbichot2-bin.000008 | mysql
```

We now have recovered the data to its state as of Tuesday 1 p.m., but still are missing the changes from that date to the date of the crash. To not lose them, we would have needed to have the MySQL server store its MySQL binary logs into a safe location (RAID disks, SAN, ...) different from the place where it stores its data files, so that these logs were not on the destroyed disk. (That is, we can start the server with a `--log-bin` option that specifies a location on a different physical device from the one on which the data directory resides. That way, the logs are safe even if the device containing the directory is lost.) If we had done this, we would have the `gbichot2-bin.000009` file at hand, and we could apply it using `mysqlbinlog` and `mysql` to restore the most recent data changes with no loss up to the moment of the crash.

### 5.9.2.3. Backup Strategy Summary

In case of an operating system crash or power failure, `InnoDB` itself does all the job of recovering data. But to make sure that you can sleep well, observe the following guidelines:

- Always run the MySQL server with the `--log-bin` option, or even `--log-bin=log_name`, where the log file name is located on some safe media different from the drive on which the data directory is located. If you have such safe media, this technique can also be good for disk load balancing (which results in a performance improvement).
- Make periodic full backups, using the `mysqldump` command shown earlier in [Section 5.9.2.1](#), “Backup Policy”, that makes an online, non-blocking backup.
- Make periodic incremental backups by flushing the logs with `FLUSH LOGS` or `mysqladmin flush-logs`.

## 5.9.3. Point-in-Time Recovery

If a MySQL server was started with the `--log-bin` option to enable binary logging, you can use the `mysqlbinlog` utility to recover data from the binary log files, starting from a specified point in time (for example, since your last backup) until the present or another specified point in time. For information on enabling the binary log and using `mysqlbinlog`, see [Section 5.11.4](#), “The Binary Log”, and [Section 8.10](#), “`mysqlbinlog` — Utility for Processing Binary Log Files”.

To restore data from a binary log, you must know the location and name of the current binary log file. By default, the server creates binary log files in the data directory, but a pathname can be specified with the `--log-bin` option to place the files in a different location. Typically the option is given in an option file (that is, `my.cnf` or `my.ini`, depending on your system). It can also be given on the command line when the server is started. To determine the name of the current binary log file, issue the following statement:

```
mysql> SHOW BINLOG EVENTS\G
```

If you prefer, you can execute the following command from the command line instead:

```
shell> mysql -u root -p -E -e "SHOW BINLOG EVENTS"
```

Enter the `root` password for your server when `mysql` prompts you for it.

### 5.9.3.1. Specifying Times for Recovery

To indicate the start and end times for recovery, specify the `--start-date` and `--stop-date` options for `mysqlbinlog`, in `DATETIME` format. As an example, suppose that exactly at 10:00 a.m. on April 20, 2005 an SQL statement was executed that deleted a large table. To restore the table and data, you could restore the previous night's backup, and then execute the following command:

```
shell> mysqlbinlog --stop-date="2005-04-20 9:59:59" \  
/var/log/mysql/bin.123456 | mysql -u root -p
```

This command recovers all of the data up until the date and time given by the `--stop-date` option. If you did not detect the erroneous SQL statement that was entered until hours later, you will probably also want to recover the activity that occurred afterward. Based on this, you could run `mysqlbinlog` again with a start date and time, like so:

```
shell> mysqlbinlog --start-date="2005-04-20 10:01:00" \  
/var/log/mysql/bin.123456 | mysql -u root -p
```

In this command, the SQL statements logged from 10:01 a.m. on will be re-executed. The combination of restoring of the previous night's dump file and the two `mysqlbinlog` commands restores everything up until one second before 10:00 a.m. and everything from 10:01 a.m. on. You should examine the log to be sure of the exact times to specify for the commands. To display the log file contents without executing them, use this command:

```
shell> mysqlbinlog /var/log/mysql/bin.123456 > /tmp/mysql_restore.sql
```

Then open the file with a text editor to examine it.

### 5.9.3.2. Specifying Positions for Recovery

Instead of specifying dates and times, the `--start-position` and `--stop-position` options for `mysqlbinlog` can be used for specifying log positions. They work the same as the start and stop date options, except that you specify log position numbers rather than dates. Using positions may enable you to be more precise about which part of the log to recover, especially if many transactions occurred around the same time as a damaging SQL statement. To determine the position numbers, run `mysqlbinlog` for a range of times near the time when the unwanted transaction was executed, but redirect the results to a text file for examination. This can be done like so:

```
shell> mysqlbinlog --start-date="2005-04-20 9:55:00" \  
--stop-date="2005-04-20 10:05:00" \  
/var/log/mysql/bin.123456 > /tmp/mysql_restore.sql
```

This command creates a small text file in the `/tmp` directory that contains the SQL statements around the time that the deleterious SQL statement was executed. Open this file with a text editor and look for the statement that you don't want to repeat. Determine the positions in the binary log for stopping and resuming the recovery and make note of them. Positions are labeled as `log_pos` followed by a number. After restoring the previous backup file, use the position numbers to process the binary log file. For example, you would use commands something like these:

```
shell> mysqlbinlog --stop-position="368312" /var/log/mysql/bin.123456 \  
| mysql -u root -p  
  
shell> mysqlbinlog --start-position="368315" /var/log/mysql/bin.123456 \  
| mysql -u root -p
```

The first command recovers all the transactions up until the stop position given. The second command recovers all transactions from the starting position given until the end of the binary log. Because the output of `mysqlbinlog` includes `SET TIMESTAMP` statements before each SQL statement recorded, the recovered data and related MySQL logs will reflect the original times at which the transactions were executed.

## 5.9.4. Table Maintenance and Crash Recovery

This section discusses how to use `myisamchk` to check or repair `MyISAM` tables (tables that have `.MYD` and `.MYI` files for storing data and indexes). The same concepts apply to using `isamchk` to check or repair `ISAM` tables (tables that have `.ISD` and `.ISM` files for storing data and indexes). For general `myisamchk` or `isamchk` background, see [Section 8.3, “myisamchk — MyISAM Table-Maintenance Utility”](#).

You can use `myisamchk` to get information about your database tables or to check, repair, or optimize them. The following sections describe how to perform these operations and how to set up a table maintenance schedule.

Even though table repair with `myisamchk` is quite secure, it is always a good idea to make a backup *before* doing a repair or any maintenance operation that could make a lot of changes to a table.

`myisamchk` operations that affect indexes can cause `FULLTEXT` indexes to be rebuilt with full-text parameters that are incompatible with the values used by the MySQL server. To avoid this problem, follow the guidelines in [Section 8.3.1, “myisamchk General Options”](#).

In many cases, you may find it simpler to do `MyISAM` table maintenance using the SQL statements that perform operations that `myisamchk` can do:

- To check or repair `MyISAM` tables, use `CHECK TABLE` or `REPAIR TABLE`.
- To optimize `MyISAM` tables, use `OPTIMIZE TABLE`.
- To analyze `MyISAM` tables, use `ANALYZE TABLE`.

These statements were introduced in different versions, but all are available from MySQL 3.23.14 on. These statements can be used directly or by means of the `mysqlcheck` client program. One advantage of these statements over `myisamchk` is that the server does all the work. With `myisamchk`, you must make sure that the server does not use the tables at the same time so that there is no unwanted interaction between `myisamchk` and the server. See [Section 13.5.2.1, “ANALYZE TABLE Syntax”](#), [Section 13.5.2.3, “CHECK TABLE Syntax”](#), [Section 13.5.2.5, “OPTIMIZE TABLE Syntax”](#), and [Section 13.5.2.6, “REPAIR TABLE Syntax”](#).

### 5.9.4.1. Using `myisamchk` for Crash Recovery

This section describes how to check for and deal with data corruption in MySQL databases. If your tables become corrupted frequently, you should try to find the reason why. See [Section A.4.2, “What to Do If MySQL Keeps Crashing”](#).

For an explanation of how `MyISAM` tables can become corrupted, see [Section 14.1.4, “MyISAM Table Problems”](#).

If you run `mysqld` with external locking disabled (which is the default as of MySQL 4.0), you cannot reliably use `myisamchk` to check a table when `mysqld` is using the same table. If you can be certain that no one will access the tables through `mysqld` while you run `myisamchk`, you only have to execute `mysqladmin flush-tables` before you start checking the tables. If you cannot guarantee this, you must stop `mysqld` while you check the tables. If you run `myisamchk` to check tables that `mysqld` is updating at the same time, you may get a warning that a table is corrupt even when it is not.

If the server is run with external locking enabled, you can use `myisamchk` to check tables at any time. In this case, if the server tries to update a table that `myisamchk` is using, the server will wait for `myisamchk` to finish before it continues.

If you use `myisamchk` to repair or optimize tables, you *must* always ensure that the `mysqld` server is

not using the table (this also applies if external locking is disabled). If you don't stop `mysqld`, you should at least do a `mysqladmin flush-tables` before you run `myisamchk`. Your tables *may become corrupted* if the server and `myisamchk` access the tables simultaneously.

When performing crash recovery, it is important to understand that each `MyISAM` table `tbl_name` in a database corresponds to three files in the database directory:

File	Purpose
<code>tbl_name.frm</code>	Definition (format) file
<code>tbl_name.MYD</code>	Data file
<code>tbl_name.MYI</code>	Index file

Each of these three file types is subject to corruption in various ways, but problems occur most often in data files and index files.

`myisamchk` works by creating a copy of the `.MYD` data file row by row. It ends the repair stage by removing the old `.MYD` file and renaming the new file to the original file name. If you use `--quick`, `myisamchk` does not create a temporary `.MYD` file, but instead assumes that the `.MYD` file is correct and generates only a new index file without touching the `.MYD` file. This is safe, because `myisamchk` automatically detects whether the `.MYD` file is corrupt and aborts the repair if it is. You can also specify the `--quick` option twice to `myisamchk`. In this case, `myisamchk` does not abort on some errors (such as duplicate-key errors) but instead tries to resolve them by modifying the `.MYD` file. Normally the use of two `--quick` options is useful only if you have too little free disk space to perform a normal repair. In this case, you should at least make a backup of the table before running `myisamchk`.

### 5.9.4.2. How to Check `MyISAM` Tables for Errors

To check a `MyISAM` table, use the following commands:

- `myisamchk tbl_name`

This finds 99.99% of all errors. What it cannot find is corruption that involves *only* the data file (which is very unusual). If you want to check a table, you should normally run `myisamchk` without options or with the `-s` (silent) option.

- `myisamchk -m tbl_name`

This finds 99.999% of all errors. It first checks all index entries for errors and then reads through all rows. It calculates a checksum for all key values in the rows and verifies that the checksum matches the checksum for the keys in the index tree.

- `myisamchk -e tbl_name`

This does a complete and thorough check of all data (`-e` means “extended check”). It does a check-read of every key for each row to verify that they indeed point to the correct row. This may take a long time for a large table that has many indexes. Normally, `myisamchk` stops after the first error it finds. If you want to obtain more information, you can add the `-v` (verbose) option. This causes `my-isamchk` to keep going, up through a maximum of 20 errors.

- `myisamchk -e -i tbl_name`

This is like the previous command, but the `-i` option tells `myisamchk` to print additional statistical information.

In most cases, a simple `myisamchk` command with no arguments other than the table name is suffi-



cient to check a table.

### 5.9.4.3. How to Repair Tables

The discussion in this section describes how to use `myisamchk` on MyISAM tables (extensions `.MYI` and `.MYD`). If you are using ISAM tables (extensions `.ISM` and `.ISD`), you should use `isamchk` instead; the concepts are similar.

If you are using MySQL 3.23.16 and above, you can (and should) use the `CHECK TABLE` and `REPAIR TABLE` statements to check and repair MyISAM tables. See [Section 13.5.2.3, “CHECK TABLE Syntax”](#), and [Section 13.5.2.6, “REPAIR TABLE Syntax”](#).

Symptoms of corrupted tables include queries that abort unexpectedly and observable errors such as these:

- `tbl_name.frm` is locked against change
- Can't find file `tbl_name.MYI` (Errcode: `nnn`)
- Unexpected end of file
- Record file is crashed
- Got error `nnn` from table handler

To get more information about the error, run `pererror nnn`, where `nnn` is the error number. The following example shows how to use `pererror` to find the meanings for the most common error numbers that indicate a problem with a table:

```
shell> pererror 126 127 132 134 135 136 141 144 145
126 = Index file is crashed / Wrong file format
127 = Record-file is crashed
132 = Old database file
134 = Record was already deleted (or record file crashed)
135 = No more room in record file
136 = No more room in index file
141 = Duplicate unique key or constraint on write or update
144 = Table is crashed and last repair failed
145 = Table was marked as crashed and should be repaired
```

Note that error 135 (no more room in record file) and error 136 (no more room in index file) are not errors that can be fixed by a simple repair. In this case, you must use `ALTER TABLE` to increase the `MAX_ROWS` and `AVG_ROW_LENGTH` table option values:

```
ALTER TABLE tbl_name MAX_ROWS=xxx AVG_ROW_LENGTH=yyy;
```

If you do not know the current table option values, use `SHOW CREATE TABLE`.

For the other errors, you must repair your tables. `myisamchk` can usually detect and fix most problems that occur.

The repair process involves up to four stages, described here. Before you begin, you should change location to the database directory and check the permissions of the table files. On Unix, make sure that they are readable by the user that `mysqld` runs as (and to you, because you need to access the files you are checking). If it turns out you need to modify files, they must also be writable by you.

This section is for the cases where a table check fails (such as those described in [Section 5.9.4.2, “How to Check MyISAM Tables for Errors”](#)), or you want to use the extended features that `myisamchk` provides.

The options that you can use for table maintenance with `myisamchk` and `isamchk` are described in [Section 8.3, “myisamchk — MyISAM Table-Maintenance Utility”](#).

If you are going to repair a table from the command line, you must first stop the `mysqld` server. Note that when you do `mysqladmin shutdown` on a remote server, the `mysqld` server is still alive for a while after `mysqladmin` returns, until all statement-processing has stopped and all index changes have been flushed to disk.

### Stage 1: Checking your tables

Run `myisamchk *.MYI` or `myisamchk -e *.MYI` if you have more time. Use the `-s` (silent) option to suppress unnecessary information.

If the `mysqld` server is stopped, you should use the `--update-state` option to tell `myisamchk` to mark the table as “checked.”

You have to repair only those tables for which `myisamchk` announces an error. For such tables, proceed to Stage 2.

If you get unexpected errors when checking (such as `out of memory` errors), or if `myisamchk` crashes, go to Stage 3.

### Stage 2: Easy safe repair

First, try `myisamchk -r -q tbl_name` (`-r -q` means “quick recovery mode”). This attempts to repair the index file without touching the data file. If the data file contains everything that it should and the delete links point at the correct locations within the data file, this should work, and the table is fixed. Start repairing the next table. Otherwise, use the following procedure:

1. Make a backup of the data file before continuing.
2. Use `myisamchk -r tbl_name` (`-r` means “recovery mode”). This removes incorrect rows and deleted rows from the data file and reconstructs the index file.
3. If the preceding step fails, use `myisamchk --safe-recover tbl_name`. Safe recovery mode uses an old recovery method that handles a few cases that regular recovery mode does not (but is slower).

Note: If you want a repair operation to go much faster, you should set the values of the `sort_buffer_size` and `key_buffer_size` variables each to about 25% of your available memory when running `myisamchk` or `isamchk`.

If you get unexpected errors when repairing (such as `out of memory` errors), or if `myisamchk` crashes, go to Stage 3.

### Stage 3: Difficult repair

You should reach this stage only if the first 16KB block in the index file is destroyed or contains incorrect information, or if the index file is missing. In this case, it is necessary to create a new index file. Do so as follows:

1. Move the data file to a safe place.
2. Use the table description file to create new (empty) data and index files:

```
shell> mysql db_name
mysql> SET AUTOCOMMIT=1;
```

```
mysql> TRUNCATE TABLE tbl_name;  
mysql> quit
```

If your version of MySQL does not have `TRUNCATE TABLE`, use `DELETE FROM tbl_name` instead.

3. Copy the old data file back onto the newly created data file. (Do not just move the old file back onto the new file. You want to retain a copy in case something goes wrong.)

Go back to Stage 2. `myisamchk -r -q` should work. (This should not be an endless loop.)

As of MySQL 4.0.2, you can also use the `REPAIR TABLE tbl_name USE_FRM` SQL statement, which performs the whole procedure automatically. There is also no possibility of unwanted interaction between a utility and the server, because the server does all the work when you use `REPAIR TABLE`. See [Section 13.5.2.6, “REPAIR TABLE Syntax”](#).

#### Stage 4: Very difficult repair

You should reach this stage only if the `.frm` description file has also crashed. That should never happen, because the description file is not changed after the table is created:

1. Restore the description file from a backup and go back to Stage 3. You can also restore the index file and go back to Stage 2. In the latter case, you should start with `myisamchk -r`.
2. If you do not have a backup but know exactly how the table was created, create a copy of the table in another database. Remove the new data file, and then move the `.frm` description and `.MYI` index files from the other database to your crashed database. This gives you new description and index files, but leaves the `.MYD` data file alone. Go back to Stage 2 and attempt to reconstruct the index file.

### 5.9.4.4. Table Optimization

To coalesce fragmented rows and eliminate wasted space that results from deleting or updating rows, run `myisamchk` in recovery mode:

```
shell> myisamchk -r tbl_name
```

You can optimize a table in the same way by using the `OPTIMIZE TABLE` SQL statement. `OPTIMIZE TABLE` does a table repair and a key analysis, and also sorts the index tree so that key lookups are faster. There is also no possibility of unwanted interaction between a utility and the server, because the server does all the work when you use `OPTIMIZE TABLE`. See [Section 13.5.2.5, “OPTIMIZE TABLE Syntax”](#).

`myisamchk` has a number of other options that you can use to improve the performance of a table:

- `--analyze, -a`
- `--sort-index, -S`
- `--sort-records=index_num, -R index_num`

For a full description of all available options, see [Section 8.3, “myisamchk — MyISAM Table-Maintenance Utility”](#).

### 5.9.4.5. Getting Information About a Table

To obtain a description of a table or statistics about it, use the commands shown here. We explain some of the information in more detail later.

- `myisamchk -d tbl_name`

Runs `myisamchk` in “describe mode” to produce a description of your table. If you start the MySQL server with external locking disabled, `myisamchk` may report an error for a table that is updated while it runs. However, because `myisamchk` does not change the table in describe mode, there is no risk of destroying data.

- `myisamchk -d -v tbl_name`

Adding `-v` runs `myisamchk` in verbose mode so that it produces more information about what it is doing.

- `myisamchk -eis tbl_name`

Shows only the most important information from a table. This operation is slow because it must read the entire table.

- `myisamchk -eiv tbl_name`

This is like `-eis`, but tells you what is being done.

Sample output for some of these commands follows. They are based on a table with these data and index file sizes:

```
-rw-rw-r-- 1 monty   tcx      317235748 Jan 12 17:30 company.MYD
-rw-rw-r-- 1 davida  tcx      96482304 Jan 12 18:35 company.MYI
```

Example of `myisamchk -d` output:

```
MyISAM file:      company.MYI
Record format:   Fixed length
Data records:    1403698 Deleted blocks:      0
Recordlength:   226
```

```
table description:
Key Start Len Index  Type
1   2     8  unique double
2  15    10  multip. text packed stripped
3  219    8   multip. double
4   63   10  multip. text packed stripped
5  167    2   multip. unsigned short
6  177    4   multip. unsigned long
7  155    4   multip. text
8  138    4   multip. unsigned long
9  177    4   multip. unsigned long
   193    1   text
```

Example of `myisamchk -d -v` output:

```
MyISAM file:      company
Record format:   Fixed length
File-version:    1
Creation time:   1999-10-30 12:12:51
Recover time:   1999-10-31 19:13:01
Status:         checked
Data records:    1403698 Deleted blocks:      0
Datafile parts: 1403698 Deleted data:      0
Datafile pointer (bytes): 3 Keyfile pointer (bytes): 3
```

```
Max datafile length: 3791650815 Max keyfile length: 4294967294
Recordlength: 226
```

```
table description:
Key Start Len Index Type Rec/key Root Blocksize
1 2 8 unique double 1 15845376 1024
2 15 10 multip. text packed stripped 2 25062400 1024
3 219 8 multip. double 73 40907776 1024
4 63 10 multip. text packed stripped 5 48097280 1024
5 167 2 multip. unsigned short 4840 55200768 1024
6 177 4 multip. unsigned long 1346 65145856 1024
7 155 4 multip. text 4995 75090944 1024
8 138 4 multip. unsigned long 87 85036032 1024
9 177 4 multip. unsigned long 178 96481280 1024
193 1 text
```

Example of `myisamchk -eis` output:

```
Checking MyISAM file: company
Key: 1: Keyblocks used: 97% Packed: 0% Max levels: 4
Key: 2: Keyblocks used: 98% Packed: 50% Max levels: 4
Key: 3: Keyblocks used: 97% Packed: 0% Max levels: 4
Key: 4: Keyblocks used: 99% Packed: 60% Max levels: 3
Key: 5: Keyblocks used: 99% Packed: 0% Max levels: 3
Key: 6: Keyblocks used: 99% Packed: 0% Max levels: 3
Key: 7: Keyblocks used: 99% Packed: 0% Max levels: 3
Key: 8: Keyblocks used: 99% Packed: 0% Max levels: 3
Key: 9: Keyblocks used: 98% Packed: 0% Max levels: 4
Total: Keyblocks used: 98% Packed: 17%

Records: 1403698 M.recordlength: 226
Packed: 0%
Recordspace used: 100% Empty space: 0%
Blocks/Record: 1.00
Record blocks: 1403698 Delete blocks: 0
Recorddata: 317235748 Deleted data: 0
Lost space: 0 Linkdata: 0

User time 1626.51, System time 232.36
Maximum resident set size 0, Integral resident set size 0
Non physical pagefaults 0, Physical pagefaults 627, Swaps 0
Blocks in 0 out 0, Messages in 0 out 0, Signals 0
Voluntary context switches 639, Involuntary context switches 28966
```

Example of `myisamchk -eiv` output:

```
Checking MyISAM file: company
Data records: 1403698 Deleted blocks: 0
- check file-size
- check delete-chain
block_size 1024:
index 1:
index 2:
index 3:
index 4:
index 5:
index 6:
index 7:
index 8:
index 9:
No recordlinks
- check index reference
- check data record references index: 1
Key: 1: Keyblocks used: 97% Packed: 0% Max levels: 4
- check data record references index: 2
Key: 2: Keyblocks used: 98% Packed: 50% Max levels: 4
- check data record references index: 3
Key: 3: Keyblocks used: 97% Packed: 0% Max levels: 4
- check data record references index: 4
Key: 4: Keyblocks used: 99% Packed: 60% Max levels: 3
- check data record references index: 5
Key: 5: Keyblocks used: 99% Packed: 0% Max levels: 3
- check data record references index: 6
Key: 6: Keyblocks used: 99% Packed: 0% Max levels: 3
- check data record references index: 7
Key: 7: Keyblocks used: 99% Packed: 0% Max levels: 3
- check data record references index: 8
```

```

Key: 8: Keyblocks used: 99% Packed: 0% Max levels: 3
- check data record references index: 9
Key: 9: Keyblocks used: 98% Packed: 0% Max levels: 4
Total: Keyblocks used: 9% Packed: 17%

- check records and index references
*** LOTS OF ROW NUMBERS DELETED ***

Records: 1403698 M.recordlength: 226 Packed: 0%
Recordspace used: 100% Empty space: 0% Blocks/Record: 1.00
Record blocks: 1403698 Delete blocks: 0
Recorddata: 317235748 Deleted data: 0
Lost space: 0 Linkdata: 0

User time 1639.63, System time 251.61
Maximum resident set size 0, Integral resident set size 0
Non physical pagefaults 0, Physical pagefaults 10580, Swaps 0
Blocks in 4 out 0, Messages in 0 out 0, Signals 0
Voluntary context switches 10604, Involuntary context switches 122798

```

Explanations for the types of information `myisamchk` produces are given here. “Keyfile” refers to the index file. “Record” and “row” are synonymous.

These options are not available unless MySQL has been built with SSL support. See [Section 5.8.7](#), “Using Secure Connections”.

- `MyISAM file`  
Name of the `MyISAM` (index) file.
- `File-version`  
Version of `MyISAM` format. Currently always 2.
- `Creation time`  
When the data file was created.
- `Recover time`  
When the index/data file was last reconstructed.
- `Data records`  
How many rows are in the table.
- `Deleted blocks`  
How many deleted blocks still have reserved space. You can optimize your table to minimize this space. See [Section 5.9.4.4](#), “Table Optimization”.
- `Datafile parts`  
For dynamic-row format, this indicates how many data blocks there are. For an optimized table without fragmented rows, this is the same as `Data records`.
- `Deleted data`  
How many bytes of unreclaimed deleted data there are. You can optimize your table to minimize this space. See [Section 5.9.4.4](#), “Table Optimization”.
- `Datafile pointer`

The size of the data file pointer, in bytes. It is usually 2, 3, 4, or 5 bytes. Most tables manage with 2 bytes, but this cannot be controlled from MySQL yet. For fixed tables, this is a row address. For dynamic tables, this is a byte address.

- `Keyfile pointer`

The size of the index file pointer, in bytes. It is usually 1, 2, or 3 bytes. Most tables manage with 2 bytes, but this is calculated automatically by MySQL. It is always a block address.

- `Max datafile length`

How long the table data file can become, in bytes.

- `Max keyfile length`

How long the table index file can become, in bytes.

- `Recordlength`

How much space each row takes, in bytes.

- `Record format`

The format used to store table rows. The preceding examples use `Fixed length`. Other possible values are `Compressed` and `Packed`.

- `table description`

A list of all keys in the table. For each key, `myisamchk` displays some low-level information:

- `Key`

This key's number.

- `Start`

Where in the row this portion of the index starts.

- `Len`

How long this portion of the index is. For packed numbers, this should always be the full length of the column. For strings, it may be shorter than the full length of the indexed column, because you can index a prefix of a string column.

- `Index`

Whether a key value can exist multiple times in the index. Possible values are `unique` or `multiple`. (multiple).

- `Type`

What data type this portion of the index has. This is a `MyISAM` data type with the possible values `packed`, `stripped`, or `empty`.

- `Root`

Address of the root index block.

- `Blocksize`

The size of each index block. By default this is 1024, but the value may be changed at compile time when MySQL is built from source.

- `Rec/key`

This is a statistical value used by the optimizer. It tells how many rows there are per value for this index. A unique index always has a value of 1. This may be updated after a table is loaded (or greatly changed) with `myisamchk -a`. If this is not updated at all, a default value of 30 is given.

For the table shown in the examples, there are two `table description` lines for the ninth index. This indicates that it is a multiple-part index with two parts.

- `Keyblocks used`

What percentage of the keyblocks are used. When a table has just been reorganized with `myisamchk`, as for the table in the examples, the values are very high (very near the theoretical maximum).

- `Packed`

MySQL tries to pack key values that have a common suffix. This can only be used for indexes on `CHAR` and `VARCHAR` columns. For long indexed strings that have similar leftmost parts, this can significantly reduce the space used. In the third of the preceding examples, the fourth key is 10 characters long and a 60% reduction in space is achieved.

- `Max levels`

How deep the B-tree for this key is. Large tables with long key values get high values.

- `Records`

How many rows are in the table.

- `M.recordlength`

The average row length. This is the exact row length for tables with fixed-length rows, because all rows have the same length.

- `Packed`

MySQL strips spaces from the end of strings. The `Packed` value indicates the percentage of savings achieved by doing this.

- `Recordspace used`

What percentage of the data file is used.

- `Empty space`

What percentage of the data file is unused.

- `Blocks/Record`

Average number of blocks per row (that is, how many links a fragmented row is composed of). This is always 1.0 for fixed-format tables. This value should stay as close to 1.0 as possible. If it gets too large, you can reorganize the table. See [Section 5.9.4.4, “Table Optimization”](#).



- `Recordblocks`  
How many blocks (links) are used. For fixed-format tables, this is the same as the number of rows.
- `Deleteblocks`  
How many blocks (links) are deleted.
- `Recorddata`  
How many bytes in the data file are used.
- `Deleted data`  
How many bytes in the data file are deleted (unused).
- `Lost space`  
If a row is updated to a shorter length, some space is lost. This is the sum of all such losses, in bytes.
- `Linkdata`  
When the dynamic table format is used, row fragments are linked with pointers (4 to 7 bytes each). `Linkdata` is the sum of the amount of storage used by all such pointers.

If a table has been compressed with `myisampack`, `myisamchk -d` prints additional information about each table column. See [Section 8.5, “myisampack — Generate Compressed, Read-Only MyISAM Tables”](#), for an example of this information and a description of what it means.

### 5.9.4.6. Setting Up a Table Maintenance Schedule

It is a good idea to perform table checks on a regular basis rather than waiting for problems to occur. One way to check and repair `MyISAM` tables is with the `CHECK TABLE` and `REPAIR TABLE` statements. These are available starting with MySQL 3.23.16. See [Section 13.5.2.3, “CHECK TABLE Syntax”](#), and [Section 13.5.2.6, “REPAIR TABLE Syntax”](#).

Another way to check tables is to use `myisamchk`. For maintenance purposes, you can use `myisamchk -s`. The `-s` option (short for `--silent`) causes `myisamchk` to run in silent mode, printing messages only when errors occur.

It is also a good idea to check tables when the server starts. For example, whenever the machine has done a restart in the middle of an update, you usually need to check all the tables that could have been affected. (These are “expected” crashed tables.) To check `MyISAM` tables automatically, start the server with the `--myisam-recover` option, available as of MySQL 3.23.25. If your server is too old to support this option, you could add a test to `mysqld_safe` that runs `myisamchk` to check all tables that have been modified during the last 24 hours if there is an old `.pid` (process ID) file left after a restart. (The `.pid` file is created by `mysqld` when it starts and removed when it terminates normally. The presence of a `.pid` file at system startup time indicates that `mysqld` terminated abnormally.)

It is also a good idea to enable automatic `MyISAM` table checking. For example, whenever the machine has done a restart in the middle of an update, you usually need to check each table that could have been affected before it is used further. (These are “expected crashed tables.”) To check `MyISAM` tables automatically, start the server with the `--myisam-recover` option, available as of MySQL 3.23.25. See [Section 5.2.1, “Command Options”](#). If your server is too old to support this option, you could add a test to `mysqld_safe` that runs `myisamchk` to check all tables that have been modified during the last 24 hours if there is an old `.pid` (process ID) file left after a restart. (The `.pid` file is created by `mysqld` when it starts and removed when it terminates normally. The presence of a `.pid` file at system startup time indicates that `mysqld` terminated abnormally.)

You should also check your tables regularly during normal system operation. At MySQL AB, we run a `cron` job to check all our important tables once a week, using a line like this in a `crontab` file:

```
35 0 * * 0 /path/to/myisamchk --fast --silent /path/to/datadir/*/*.MYI
```

This prints out information about crashed tables so that we can examine and repair them when needed.

Because we have not had any unexpectedly crashed tables (tables that become corrupted for reasons other than hardware trouble) for several years, once a week is more than sufficient for us.

We recommend that to start with, you execute `myisamchk -s` each night on all tables that have been updated during the last 24 hours, until you come to trust MySQL as much as we do.

Normally, MySQL tables need little maintenance. If you are performing many updates to `MyISAM` tables with dynamic-sized rows (tables with `VARCHAR`, `BLOB`, or `TEXT` columns) or have tables with many deleted rows you may want to defragment/reclaim space from the tables from time to time. You can do this by using `OPTIMIZE TABLE` on the tables in question. Alternatively, if you can stop the `mysqld` server for a while, change location into the data directory and use this command while the server is stopped:

```
shell> myisamchk -r -s --sort-index --sort_buffer_size=16M */*.MYI
```

For `ISAM` tables, the command is similar:

```
shell> isamchk -r -s --sort-index -O sort_buffer_size=16M */*.ISM
```

## 5.10. MySQL Localization and International Usage

This section describes how to configure the server to use different character sets. It also discusses how to set the server's time zone and enable per-connection time zone support.

### 5.10.1. The Character Set Used for Data and Sorting

By default, MySQL uses the `latin1` (cp1252 West European) character set and the `latin1_swedish_ci` collation that sorts according to Swedish/Finnish rules. These defaults are suitable for the United States and most of Western Europe.

All MySQL binary distributions are compiled with `--with-extra-charsets=complex`. This adds code to all standard programs that enables them to handle `latin1` and all multi-byte character sets within the binary. Other character sets are loaded from a character-set definition file when needed.

The character set determines what characters are allowed in identifiers. The collation determines how strings are sorted by the `ORDER BY` and `GROUP BY` clauses of the `SELECT` statement.

You can change the default server character set and collation with the `--character-set-server` and `--collation-server` options when you start the server. The collation must be a legal collation for the default character set. (Use the `SHOW COLLATION` statement to determine which collations are available for each character set.) See [Section 5.2.1, “Command Options”](#). Before MySQL 4.1.3, use `--default-character-set` to change the character set; there is no option for changing the collation.

The character sets available depend on the `--with-charset=charset_name` and `--with-extra-charsets=list-of-charsets | complex | all | none` options to `configure`, and the character set configuration files listed in `SHAREDIR/charsets/Index`. See [Section 2.9.2, “Typical configure Options”](#).

If you change the character set when running MySQL, that may also change the sort order. Consequently, you must run `myisamchk -r -q --set-collation=collation_name` on all `MyISAM` tables, or your indexes may not be ordered correctly. (Use `--set-character-set` before MySQL 4.1.1.)

When a client connects to a MySQL server, the server indicates to the client what the server's default character set is. The client switches to this character set for this connection.

You should use `mysql_real_escape_string()` when escaping strings for an SQL query. `mysql_real_escape_string()` is identical to the old `mysql_escape_string()` function, except that it takes the `MYSQL` connection handle as the first parameter so that the appropriate character set can be taken into account when escaping characters.

If the client is compiled with paths that differ from where the server is installed and the user who configured MySQL didn't include all character sets in the MySQL binary, you must tell the client where it can find the additional character sets it needs if the server runs with a different character set from the client. You can do this by specifying a `--character-sets-dir` option to indicate the path to the directory in which the dynamic MySQL character sets are stored. For example, you can put the following in an option file:

```
[client]
character-sets-dir=/usr/local/mysql/share/mysql/charsets
```

You can force the client to use specific character set as follows:

```
[client]
default-character-set=charset_name
```

This is normally unnecessary, however.

### 5.10.1.1. Using the German Character Set

In MySQL 4.0, to get German sorting order, you should start `mysqld` with a `-default-character-set=latin1_de` option. This affects server behavior in several ways:

- When sorting and comparing strings, the following mapping is performed on the strings before doing the comparison:

```
ä -> ae
ö -> oe
Û -> ue
ß -> ss
```

- All accented characters are converted to their unaccented uppercase counterpart. All letters are converted to uppercase.
- When comparing strings with `LIKE`, the one-character to two-character mapping is not done. All letters are converted to uppercase. Accents are removed from all letters except `Û`, `ü`, `Ö`, `ö`, `Ä`, and `ä`.

In MySQL 4.1 and up, character set and collation are specified separately. You should select the `latin1` character set and either the `latin1_german1_ci` or `latin1_german2_ci` collation. For example, to start the server with the `latin1_german1_ci` collation, use the `-character-set-server=latin1` and `--collation-server=latin1_german1_ci` options.

For information on the differences between these two collations, see [Section 10.10.2, “West European Character Sets”](#).

## 5.10.2. Setting the Error Message Language

By default, `mysqld` produces error messages in English, but they can also be displayed in any of these other languages: Czech, Danish, Dutch, Estonian, French, German, Greek, Hungarian, Italian, Japanese, Korean, Norwegian, Norwegian-ny, Polish, Portuguese, Romanian, Russian, Slovak, Spanish, or Swedish.

To start `mysqld` with a particular language for error messages, use the `--language` or `-L` option. The option value can be a language name or the full path to the error message file. For example:

```
shell> mysqld --language=swedish
```

Or:

```
shell> mysqld --language=/usr/local/share/swedish
```

The language name should be specified in lowercase.

By default, the language files are located in the `share/LANGUAGE` directory under the MySQL base directory.

You can also change the content of the error messages produced by the server. Details can be found in the MySQL Internals manual, available at <http://dev.mysql.com/doc/>. If you upgrade to a newer version of MySQL after changing the error messages, remember to repeat your changes after the upgrade.

## 5.10.3. Adding a New Character Set

This section discusses the procedure for adding a new character set to MySQL. You must have a MySQL source distribution to use these instructions. To choose the proper procedure, determine whether the character set is simple or complex:

- If the character set does not need to use special string collating routines for sorting and does not need multi-byte character support, it is simple.
- If it needs either of those features, it is complex.

For example, `latin1` and `danish` are simple character sets, whereas `big5` and `czech` are complex character sets.

In the following instructions, the name of the character set is represented by `MYSET`.

For a simple character set, do the following:

1. Add `MYSET` to the end of the `sql/share/charsets/Index` file. Assign a unique number to it.
2. Create the file `sql/share/charsets/MYSET.conf`. (You can use a copy of `sql/share/charsets/latin1.conf` as the basis for this file.)

The syntax for the file is very simple:

- Comments start with a ‘#’ character and continue to the end of the line.
- Words are separated by arbitrary amounts of whitespace.

- When defining the character set, every word must be a number in hexadecimal format.
- The `ctype` array takes up the first 257 words. The `to_lower[ ]`, `to_upper[ ]` and `sort_order[ ]` arrays take up 256 words each after that.

See [Section 5.10.4, “The Character Definition Arrays”](#).

3. Add the character set name to the `CHARSETS_AVAILABLE` and `COMPILED_CHARSETS` lists in `configure.in`.
4. Reconfigure, recompile, and test.

For a complex character set, do the following:

1. Create the file `strings/ctype-MYSET.c` in the MySQL source distribution.
2. Add `MYSET` to the end of the `sql/share/charsets/Index` file. Assign a unique number to it.
3. Look at one of the existing `ctype-*.c` files (such as `strings/ctype-big5.c`) to see what needs to be defined. Note that the arrays in your file must have names like `ctype_MYSET`, `to_lower_MYSET`, and so on. These correspond to the arrays for a simple character set. See [Section 5.10.4, “The Character Definition Arrays”](#).
4. Near the top of the file, place a special comment like this:

```
/*
 * This comment is parsed by configure to create ctype.c,
 * so don't change it unless you know what you are doing.
 *
 * .configure. number_MYSET=MYNUMBER
 * .configure. strxfrm_multiply_MYSET=N
 * .configure. mbmaxlen_MYSET=N
 */
```

The `configure` program uses this comment to include the character set into the MySQL library automatically.

The `strxfrm_multiply` and `mbmaxlen` lines are explained in the following sections. You need include them only if you need the string collating functions or the multi-byte character set functions, respectively.

5. You should then create some of the following functions:
  - `my_strncoll_MYSET()`
  - `my_strcoll_MYSET()`
  - `my_strxfrm_MYSET()`
  - `my_like_range_MYSET()`

See [Section 5.10.5, “String Collating Support”](#).

6. Add the character set name to the `CHARSETS_AVAILABLE` and `COMPILED_CHARSETS` lists in `configure.in`.
7. Reconfigure, recompile, and test.

The `sql/share/charsets/README` file includes additional instructions.

If you want to have the character set included in the MySQL distribution, mail a patch to the MySQL `internals` mailing list. See [Section 1.7.1, “MySQL Mailing Lists”](#).

## 5.10.4. The Character Definition Arrays

`to_lower[ ]` and `to_upper[ ]` are simple arrays that hold the lowercase and uppercase characters corresponding to each member of the character set. For example:

```
to_lower['A'] should contain 'a'
to_upper['a'] should contain 'A'
```

`sort_order[ ]` is a map indicating how characters should be ordered for comparison and sorting purposes. Quite often (but not for all character sets) this is the same as `to_upper[ ]`, which means that sorting is case-insensitive. MySQL sorts characters based on the values of `sort_order[ ]` elements. For more complicated sorting rules, see the discussion of string collating in [Section 5.10.5, “String Collating Support”](#).

`ctype[ ]` is an array of bit values, with one element for one character. (Note that `to_lower[ ]`, `to_upper[ ]`, and `sort_order[ ]` are indexed by character value, but `ctype[ ]` is indexed by character value + 1. This is an old legacy convention for handling EOF.)

You can find the following bitmask definitions in `m_ctype.h`:

```
#define _U      01      /* Uppercase */
#define _L      02      /* Lowercase */
#define _N      04      /* Numeral (digit) */
#define _S      010     /* Spacing character */
#define _P      020     /* Punctuation */
#define _C      040     /* Control character */
#define _B      0100    /* Blank */
#define _X      0200    /* hexadecimal digit */
```

The `ctype[ ]` entry for each character should be the union of the applicable bitmask values that describe the character. For example, 'A' is an uppercase character (`_U`) as well as a hexadecimal digit (`_X`), so `ctype['A'+1]` should contain the value:

```
_U + _X = 01 + 0200 = 0201
```

## 5.10.5. String Collating Support

If the sorting rules for your language are too complex to be handled with the simple `sort_order[ ]` table, you need to use the string collating functions.

The best documentation for this is the existing character sets. Look at the `big5`, `czech`, `gbk`, `sjis`, and `tis160` character sets for examples.

You must specify the `strxfrm_multiply_MYSET=N` value in the special comment at the top of the file. `N` should be set to the maximum ratio the strings may grow during `my_strxfrm_MYSET` (it must be a positive integer).

## 5.10.6. Multi-Byte Character Support

If you want to add support for a new character set that includes multi-byte characters, you need to use the multi-byte character functions.

The best documentation for this is the existing character sets. Look at the `euc_kr`, `gb2312`, `gbk`,

`sjis`, and `ujis` character sets for examples. These are implemented in the `ctype-charset_name.c` files in the `strings` directory.

You must specify the `mbmaxlen_MYSET=N` value in the special comment at the top of the source file. `N` should be set to the size in bytes of the largest character in the set.

## 5.10.7. Problems With Character Sets

If you try to use a character set that is not compiled into your binary, you might run into the following problems:

- Your program uses an incorrect path to determine where the character sets are stored. (Default `/usr/local/mysql/share/mysql/charsets`). This can be fixed by using the `-character-sets-dir` option when you run the program in question.
- The character set is a multi-byte character set that cannot be loaded dynamically. In this case, you must recompile the program with support for the character set.
- The character set is a dynamic character set, but you do not have a configure file for it. In this case, you should install the configure file for the character set from a new MySQL distribution.
- If your `Index` file does not contain the name for the character set, your program displays the following error message:

```
ERROR 1105: File '/usr/local/share/mysql/charsets/? .conf'
not found (Errcode: 2)
```

In this case, you should either get a new `Index` file or manually add the name of any missing character sets to the current file.

For `MyISAM` tables, you can check the character set name and number for a table with `myisamchk -dvv tbl_name`.

## 5.10.8. MySQL Server Time Zone Support

Before MySQL 4.1.3, you can set the time zone for the server with the `--timezone=timezone_name` option to `mysqld_safe`. You can also set it by setting the `TZ` environment variable before you start `mysqld`.

The allowable values for `--timezone` or `TZ` are system-dependent. Consult your operating system documentation to see what values are acceptable.

Beginning with MySQL 4.1.3, the server maintains several time zone settings:

- The system time zone. When the server starts, it attempts to determine the time zone of the host machine and uses it to set the `system_time_zone` system variable. The value does not change thereafter.
- The server's current time zone. The global `time_zone` system variable indicates the time zone the server currently is operating in. The initial value for `time_zone` is `'SYSTEM'`, which indicates that the server time zone is the same as the system time zone. The initial value can be specified explicitly with the `--default-time-zone=timezone` option. If you have the `SUPER` privilege, you can set the global value at runtime with this statement:

```
mysql> SET GLOBAL time_zone = timezone;
```

- Per-connection time zones. Each client that connects has its own time zone setting, given by the session `time_zone` variable. Initially, the session variable takes its value from the global `time_zone` variable, but the client can change its own time zone with this statement:

```
mysql> SET time_zone = timezone;
```

The current values of the global and client-specific time zones can be retrieved like this:

```
mysql> SELECT @@global.time_zone, @@session.time_zone;
```

`timezone` values can be given as strings indicating an offset from UTC, such as `'+10:00'` or `'-6:00'`. If the time zone information tables in the `mysql` database have been created and populated, you can also use named time zones, such as `'Europe/Helsinki'`, `'US/Eastern'`, or `'MET'`. The value `'SYSTEM'` can be used to indicate that the time zone should be the same as the system time zone. Time zone names are not case sensitive.

The MySQL installation procedure creates the time zone tables in the `mysql` database, but does not load them. You must do so manually. (If you are upgrading to MySQL 4.1.3 or later from an earlier version, you should create the tables by upgrading your `mysql` database. Use the instructions in [Section 5.5.1](#), “[mysql\\_fix\\_privilege\\_tables — Upgrade MySQL System Tables](#)”.)

If your system has its own *zoneinfo database* (the set of files describing time zones), you should use the `mysql_tzinfo_to_sql` program for filling the time zone tables. Examples of such systems are Linux, FreeBSD, Sun Solaris, and Mac OS X. One likely location for these files is the `/usr/share/zoneinfo` directory. If your system does not have a zoneinfo database, you can use the downloadable package described later in this section.

The `mysql_tzinfo_to_sql` program is used to load the time zone tables. On the command line, pass the zoneinfo directory pathname to `mysql_tzinfo_to_sql` and send the output into the `mysql` program. For example:

```
shell> mysql_tzinfo_to_sql /usr/share/zoneinfo | mysql -u root mysql
```

`mysql_tzinfo_to_sql` reads your system's time zone files and generates SQL statements from them. `mysql` processes those statements to load the time zone tables.

`mysql_tzinfo_to_sql` also can be used to load a single time zone file, and to generate leap second information:

- To load a single time zone file `tz_file` that corresponds to a time zone name `tz_name`, invoke `mysql_tzinfo_to_sql` like this:

```
shell> mysql_tzinfo_to_sql tz_file tz_name | mysql -u root mysql
```

- If your time zone needs to account for leap seconds, initialize the leap second information like this, where `tz_file` is the name of your time zone file:

```
shell> mysql_tzinfo_to_sql --leap tz_file | mysql -u root mysql
```

If your system doesn't have a zoneinfo database (for example, Windows or HP-UX), you can use the package of pre-built time zone tables that is available for download at <http://dev.mysql.com/downloads/timezones.html>. This package contains `.frm`, `.MYD`, and `.MYI` files for the `MyISAM` time zone tables. These tables should be part of the `mysql` database, so you should place the files in the `mysql` subdirectory of your MySQL server's data directory. The server should be



stopped while you do this.

**Warning:** Please don't use the downloadable package if your system has a zoneinfo database. Use the `mysql_tzinfo_to_sql` utility instead. Otherwise, you may cause a difference in datetime handling between MySQL and other applications on your system.

For information about time zone settings in replication setup, please see [Section 6.7, “Replication Features and Known Problems”](#).

## 5.11. MySQL Server Logs

MySQL has several different log files that can help you find out what is going on inside `mysqld`:

Log Type	Information Written to Log
The error log	Problems encountered starting, running, or stopping <code>mysqld</code>
The isam log	All changes to the <code>ISAM</code> tables (used only for debugging the <code>ISAM</code> code)
The general query log	Established client connections and statements received from clients
The update log	All statements that change data (this log is deprecated)
The binary log	All statements that change data (also used for replication)
The slow log	All queries that took more than <code>long_query_time</code> seconds to execute or didn't use indexes

By default, all log files are created in the `mysqld` data directory. You can force `mysqld` to close and reopen the log files (or in some cases switch to a new log) by flushing the logs. Log flushing occurs when you issue a `FLUSH LOGS` statement or execute `mysqladmin flush-logs` or `mysqladmin refresh`. See [Section 13.5.5.2, “FLUSH Syntax”](#).

If you are using MySQL replication capabilities, slave replication servers maintain additional log files called relay logs. These are discussed in [Chapter 6, Replication](#).

### 5.11.1. The Error Log

The error log file contains information indicating when `mysqld` was started and stopped and also any critical errors that occur while the server is running. If `mysqld` notices a table that needs to be automatically checked or repaired, it writes a message to the error log.

On some operating systems, the error log contains a stack trace if `mysqld` dies. The trace can be used to determine where `mysqld` died. See [Section E.1.4, “Using a Stack Trace”](#).

If `mysqld` dies unexpectedly and `mysqld_safe` needs to restart it, `mysqld_safe` writes a `restarted mysqld` message to the error log.

Beginning with MySQL 4.0.10, you can specify where `mysqld` stores the error log file with the `--log-error[=file_name]` option. If no `file_name` value is given, `mysqld` uses the name `host_name.err` and writes the file in the data directory. (Prior to MySQL 4.0.10, the Windows error log name is `mysql.err`.) If you execute `FLUSH LOGS`, the error log is renamed with a suffix of `-old` and `mysqld` creates a new empty log file. (No renaming occurs if the `--log-error` option was not given.)

In older MySQL versions on Unix, error log handling was done by `mysqld_safe` which redirected the error file to `host_name.err`. You could change this filename by specifying a `--err-log=file_name` option to `mysqld_safe`.

If you do not specify `--log-error`, or (on Windows) if you use the `--console` option, errors are written to `stderr`, the standard error output. Usually this is your terminal.

On Windows, error output is always written to the `.err` file if `--console` is not given.

## 5.11.2. The General Query Log

The general query log is a general record of what `mysqld` is doing. The server writes information to this log when clients connect or disconnect, and it logs each SQL statement received from clients. The general query log can be very useful when you suspect an error in a client and want to know exactly what the client sent to `mysqld`.

Older versions of the `mysql.server` script (from MySQL 3.23.4 to 3.23.8) pass a `--log` option to `safe_mysqld` to enable the general query log. If you need better performance when you start using MySQL in a production environment, you can remove the `--log` option from `mysql.server` or change it to `--log-bin`. See [Section 5.11.4, “The Binary Log”](#).

`mysqld` writes statements to the query log in the order that it receives them. This may be different from the order in which they are executed. This is in contrast to the update log and the binary log, which are written after the query is executed, but before any locks are released. (Also, the query log contains all statements, whereas the update and binary logs do not contain statements that only select data.)

To enable the general query log, start `mysqld` with the `--log[=file_name]` or `-l [file_name]` option. If no `file_name` value is given, the default name is `host_name.log` in the data directory.

Server restarts and log flushing do not cause a new general query log file to be generated (although flushing closes and reopens it). On Unix, you can rename the file and create a new one by using the following commands:

```
shell> mv host_name.log host_name-old.log
shell> mysqladmin flush-logs
shell> cp host_name-old.log backup-directory
shell> rm host_name-old.log
```

On Windows, you cannot rename the log file while the server has it open. You must stop the server and rename the file, and then restart the server to create a new log file.

## 5.11.3. The Update Log

**Note:** The update log has been deprecated and replaced by the more useful, informative, and efficient binary log. See [Section 5.11.4, “The Binary Log”](#).

When started with the `--log-update[=file_name]` option, `mysqld` writes a log file containing all SQL statements that update data. If no `file_name` value is given, the default name is name of the host machine. If a filename is given, but it does not contain a leading path, the file is written in the data directory. If `file_name` does not have an extension, `mysqld` creates log files with names of the form `file_name.nnnnnn`, where `nnnnnn` is a number that is incremented each time you start the server or flush the logs.

**Note:** For this naming scheme to work, you must not create your own files with the same names as those that might be used in the log file sequence.

Update logging is “smart” in that *it logs only statements that actually update data*. Thus, an `UPDATE` or `DELETE` with a `WHERE` clause that finds no rows is not written to the log. Update logging also skips `UPDATE` statements that merely set a column to its existing value.

The update logging is done immediately after a query completes but before any locks are released or any

commit is done. This ensures that statements are logged in execution order.

If you want to update a database from update log files, you could do the following (assuming that your update logs have names of the form `file_name.nnnnnn`):

```
shell> ls -l -t -r file_name.[0-9]* | xargs cat | mysql
```

`ls` is used to sort the update log filenames into the right order.

This can be useful if you have to revert to backup files after a crash and you want to redo the updates that occurred between the time of the backup and the crash.

## 5.11.4. The Binary Log

The binary log contains all statements that update data or (starting from MySQL 4.1.3) potentially could have updated it (for example, a `DELETE` which matched no rows). Statements are stored in the form of “events” that describe the modifications. The binary log also contains information about how long each statement took that updated data.

**Note:** The binary log has replaced the old update log, which is being phased out of future MySQL release series after 4.1. The binary log contains all information that is available in the update log in a more efficient format and in a manner that is transaction-safe. If you are using transactions, you must use the MySQL binary log for backups instead of the old update log.

The binary log does not contain statements that do not modify any data. If you want to log all statements (for example, to identify a problem query), use the general query log. See [Section 5.11.2, “The General Query Log”](#).

The primary purpose of the binary log is to be able to update databases during a restore operation as fully as possible, because the binary log contains all updates done after a backup was made. The binary log is also used on master replication servers as a record of the statements to be sent to slave servers. See [Chapter 6, Replication](#).

Running the server with the binary log enabled makes performance about 1% slower. However, the benefits of the binary log for restore operations and in allowing you to set up replication generally outweigh this minor performance decrement.

When started with the `--log-bin[=base_name]` option, `mysqld` writes a log file containing all SQL commands that update data. If no *base\_name* value is given, the default name is the name of the host machine followed by `-bin`. If the basename is given, but not as an absolute pathname, the server writes the file in the data directory. It is recommended that you specify a basename; see [Section A.8.4, “Open Issues in MySQL”](#), for the reason.

If you supply an extension in the log name (for example, `--log-bin=base_name.extension`), the extension is silently removed and ignored.

`mysqld` appends a numeric extension to the binary log basename. The number increases each time the server creates a new log file, thus creating an ordered series of files. The server creates a new binary log file each time it starts or flushes the logs. The server also creates a new binary log file automatically when the current log's size reaches `max_binlog_size`. A binary log file may become larger than `max_binlog_size` if you are using large transactions because a transaction is written to the file in one piece, never split between files.

To keep track of which binary log files have been used, `mysqld` also creates a binary log index file that contains the names of all used binary log files. By default this has the same basename as the binary log file, with the extension `.index`. You can change the name of the binary log index file with the `--log-bin-index[=file_name]` option. You should not manually edit this file while `mysqld` is running; doing so would confuse `mysqld`.

Before MySQL 4.1.9, a write to a binary log file or binary log index file that failed due to a full disk or an exceeded quota resulted in corruption of the file. Starting from MySQL 4.1.9, writes to the binary log file and binary log index file are handled the same way as writes to `MyISAM` tables. See [Section A.4.3, “How MySQL Handles a Full Disk”](#).

You can delete all binary log files with the `RESET MASTER` statement, or a subset of them with `PURGE MASTER LOGS`. See [Section 13.5.5.5, “RESET Syntax”](#), and [Section 13.6.1, “SQL Statements for Controlling Master Servers”](#).

The binary log format has some known limitations that can affect recovery from backups, especially in old versions. These caveats, which also affect replication, are listed at [Section 6.7, “Replication Features and Known Problems”](#). One caveat which does not affect replication but only recovery with `mysqlbinlog`: before MySQL 4.1, `mysqlbinlog` could not prepare output suitable for `mysql` if the binary log contained interlaced statements originating from different clients that used temporary tables of the same name. This is fixed in MySQL 4.1. However, the problem still existed for `LOAD DATA INFILE` statements until it was fixed in MySQL 4.1.8.

You can use the following options to `mysqld` to affect what is logged to the binary log. See also the discussion that follows this option list.

If you are using replication, the options described here affect which statements are sent by a master server to its slaves. There are also options for slave servers that control which statements received from the master to execute or ignore. For details, see [Section 6.8, “Replication Startup Options”](#).

- `--binlog-do-db=db_name`

Tell the server to restrict binary logging to updates for which the default database is `db_name` (that is, the database selected by `USE`). All other databases that are not explicitly mentioned are ignored. If you use this option, you should ensure that you do updates only in the default database.

There is an exception to this for `CREATE DATABASE`, `ALTER DATABASE`, and `DROP DATABASE` statements. The server uses the database named in the statement (not the default database) to decide whether it should log the statement.

An example of what does not work as you might expect: If the server is started with `binlog-do-db=sales`, and you run `USE prices; UPDATE sales.january SET amount=amount+1000;`, this statement is *not* written into the binary log.

To log multiple databases, use multiple options, specifying the option once for each database.

- `--binlog-ignore-db=db_name`

Tell the server to suppress binary logging of updates for which the default database is `db_name` (that is, the database selected by `USE`). If you use this option, you should ensure that you do updates only in the default database.

As with the `--binlog-do-db` option, there is an exception for the `CREATE DATABASE`, `ALTER DATABASE`, and `DROP DATABASE` statements. The server uses the database named in the statement (not the default database) to decide whether it should log the statement.

An example of what does not work as you might expect: If the server is started with `binlog-ignore-db=sales`, and you run `USE prices; UPDATE sales.january SET amount=amount+1000;`, this statement *is* written into the binary log.

To ignore multiple databases, use multiple options, specifying the option once for each database.

The server evaluates the options for logging or ignoring updates to the binary log according to the fol-

lowing rules. As described previously, there is an exception for the `CREATE DATABASE`, `ALTER DATABASE`, and `DROP DATABASE` statements. In those cases, the database being *created*, *altered*, or *dropped* replaces the default database in the following rules:

1. Are there `--binlog-do-db` or `--binlog-ignore-db` rules?
  - No: Write the statement to the binary log and exit.
  - Yes: Go to the next step.
2. There are some rules (`--binlog-do-db`, `--binlog-ignore-db`, or both). Is there a default database (has any database been selected by `USE`)?
  - No: Do *not* write the statement, and exit.
  - Yes: Go to the next step.
3. There is a default database. Are there some `--binlog-do-db` rules?
  - Yes: Does the default database match any of the `--binlog-do-db` rules?
    - Yes: Write the statement and exit.
    - No: Do *not* write the statement, and exit.
  - No: Go to the next step.
4. There are some `--binlog-ignore-db` rules. Does the default database match any of the `--binlog-ignore-db` rules?
  - Yes: Do not write the statement, and exit.
  - No: Write the query and exit.

For example, a slave running with only `--binlog-do-db=sales` does not write to the binary log any statement for which the default database is different from `sales` (in other words, `--binlog-do-db` can sometimes mean “ignore other databases”).

If you are using replication, you should not delete old binary log files until you are sure that no slave still needs to use them. For example, if your slaves never run more than three days behind, once a day you can execute `mysqladmin flush-logs` on the master and then remove any logs that are more than three days old. You can remove the files manually, but it is preferable to use `PURGE MASTER LOGS`, which also safely updates the binary log index file for you (and which can take a date argument as of MySQL 4.1). See [Section 13.6.1, “SQL Statements for Controlling Master Servers”](#).

A client that has the `SUPER` privilege can disable binary logging of its own statements by using a `SET SQL_LOG_BIN=0` statement. See [Section 13.5.3, “SET Syntax”](#).

You can display the contents of binary log files with the `mysqlbinlog` utility. This can be useful when you want to reprocess statements in the log. For example, you can update a MySQL server from the binary log as follows:

```
shell> mysqlbinlog log_file | mysql -h server_name
```

See [Section 8.10, “mysqlbinlog — Utility for Processing Binary Log Files”](#), for more information on the `mysqlbinlog` utility and how to use it. `mysqlbinlog` also can be used with relay log files because they are written using the same format as binary log files.

Binary logging is done immediately after a statement completes but before any locks are released or any commit is done. This ensures that the log is logged in execution order.

Updates to non-transactional tables are stored in the binary log immediately after execution. Within an uncommitted transaction, all updates (`UPDATE`, `DELETE`, or `INSERT`) that change transactional tables such as `BDB` or `InnoDB` tables are cached until a `COMMIT` statement is received by the server. At that point, `mysqld` writes the entire transaction to the binary log before the `COMMIT` is executed. When the thread that handles the transaction starts, it allocates a buffer of `binlog_cache_size` to buffer statements. If a statement is bigger than this, the thread opens a temporary file to store the transaction. The temporary file is deleted when the thread ends.

Modifications to non-transactional tables cannot be rolled back. If a transaction that is rolled back includes modifications to non-transactional tables, the entire transaction is logged with a `ROLLBACK` statement at the end to ensure that the modifications to those tables are replicated. This is true as of MySQL 4.0.15.

The `Binlog_cache_use` status variable shows the number of transactions that used this buffer (and possibly a temporary file) for storing statements. The `Binlog_cache_disk_use` status variable shows how many of those transactions actually had to use a temporary file. These two variables can be used for tuning `binlog_cache_size` to a large enough value that avoids the use of temporary files.

The `max_binlog_cache_size` system variable (default 4GB) can be used to restrict the total size used to cache a multiple-statement transaction. If a transaction is larger than this, it fails and rolls back.

If you are using the update log or binary log, concurrent inserts are converted to normal inserts for `CREATE ... SELECT` or `INSERT ... SELECT` statements. This is done to ensure that you can re-create an exact copy of your tables by applying the log during a backup operation.

The binary log format differs between versions 3.23 and 4.0. (These format changes were required to implement enhancements to replication.) However, MySQL 4.1 has the same binary log format as 4.0. See [Section 6.5, “Replication Compatibility Between MySQL Versions”](#).

By default, the binary log is not synchronized to disk at each write. So if the operating system or machine (not only the MySQL server) crashes, there is a chance that the last statements of the binary log are lost. To prevent this, you can make the binary log be synchronized to disk after every *N* writes to the binary log, with the `sync_binlog` system variable. See [Section 5.2.2, “System Variables”](#). 1 is the safest value for `sync_binlog`, but also the slowest. Even with `sync_binlog` set to 1, there is still the chance of an inconsistency between the table content and binary log content in case of a crash. For example, if you are using `InnoDB` tables and the MySQL server processes a `COMMIT` statement, it writes the whole transaction to the binary log and then commits this transaction into `InnoDB`. If the server crashes between those two operations, the transaction is rolled back by `InnoDB` at restart but still exists in the binary log. This problem can be solved with the `--innodb-safe-binlog` option (available starting from MySQL 4.1.3), which adds consistency between the content of `InnoDB` tables and the binary log.

For this option to provide a greater degree of safety, the MySQL server should also be configured to synchronize the binary log and the `InnoDB` logs to disk at every transaction. The `InnoDB` logs are synchronized by default, and `sync_binlog=1` can be used to synchronize the binary log. The effect of this option is that at restart after a crash, after doing a rollback of transactions, the MySQL server cuts rolled back `InnoDB` transactions from the binary log. This ensures that the binary log reflects the exact data of `InnoDB` tables, and so, that the slave remains in synchrony with the master (not receiving a statement which has been rolled back).

Note that `--innodb-safe-binlog` can be used even if the MySQL server updates other storage engines than `InnoDB`. Only statements and transactions that affect `InnoDB` tables are subject to removal from the binary log at `InnoDB`'s crash recovery. If the MySQL server discovers at crash recovery that the binary log is shorter than it should have been, it lacks at least one successfully committed `InnoDB` transaction. This should not happen if `sync_binlog=1` and the disk/filesystem do an actual sync

when they are requested to (some don't), so the server prints an error message `The binary log <name> is shorter than its expected size`. In this case, this binary log is not correct and replication should be restarted from a fresh snapshot of the master's data.

### 5.11.5. The Slow Query Log

The slow query log consists of all SQL statements that took more than `long_query_time` seconds to execute. The time to acquire the initial table locks is not counted as execution time. The minimum and default values of `long_query_time` are 1 and 10, respectively.

`mysqld` writes a statement to the slow query log after it has been executed and after all locks have been released. Log order may be different from execution order.

To enable the slow query log, start `mysqld` with the `--log-slow-queries[=file_name]` option.

If no `file_name` value is given, the default is the name of the host machine with a suffix of `-slow.log`. If a filename is given, but not as an absolute pathname, the server writes the file in the data directory.

The slow query log can be used to find queries that take a long time to execute and are therefore candidates for optimization. However, examining a long slow query log can become a difficult task. To make this easier, you can process the slow query log using the `mysqldumpslow` command to summarize the queries that appear in the log. Use `mysqldumpslow --help` to see the options that this command supports.

Before MySQL 4.1, if you also use `--log-long-format` when logging slow queries, queries that are not using indexes are logged as well. Starting with MySQL 4.1 logging of queries not using indexes is enabled using the `--log-queries-not-using-indexes` option instead. The `--log-long-format` is deprecated as of MySQL 4.1, when `--log-short-format` was introduced. (The long log format is the default setting since version 4.1.). See [Section 5.2.1, “Command Options”](#).

In MySQL 4.0, slow administrative statements such as `OPTIMIZE TABLE`, `ANALYZE TABLE`, and `ALTER TABLE` were written to the slow query log. This logging was disabled in MySQL 4.1 until 4.1.13, when the `--log-slow-admin-statements` server option was added to specify logging of slow administrative statements.

Queries handled by the query cache are not added to the slow query log, nor are queries that would not benefit from the presence of an index because the table has zero rows or one row.

### 5.11.6. Server Log Maintenance

MySQL Server can create a number of different log files that make it easy to see what is going on. See [Section 5.11, “MySQL Server Logs”](#). However, you must clean up these files regularly to ensure that the logs do not take up too much disk space.

When using MySQL with logging enabled, you may want to back up and remove old log files from time to time and tell MySQL to start logging to new files. See [Section 5.9.1, “Database Backups”](#).

On a Linux (Red Hat) installation, you can use the `mysql-log-rotate` script for this. If you installed MySQL from an RPM distribution, this script should have been installed automatically. You should be careful with this script if you are using the binary log for replication. You should not remove binary logs until you are certain that their contents have been processed by all slaves.

On other systems, you must install a short script yourself that you start from `cron` (or its equivalent) for handling log files.

You can force MySQL to start using new log files by using `mysqladmin flush-logs` or by using the SQL statement `FLUSH LOGS`. If you are using MySQL 3.21, you must use `mysqladmin re-fresh`.

A log flushing operation does the following:

- If general query logging (`--log`) or slow query logging (`--log-slow-queries`) is used, the server closes and reopens the general query log file or slow query log file.
- If update logging (`--log-update`) or binary logging (`--log-bin`) is used, closes the log and opens a new log file with a higher sequence number.

The server creates a new binary log file when you flush the logs. However, it just closes and reopens the general and slow query log files. To cause new files to be created on Unix, rename the current logs before flushing them. At flush time, the server will open new logs with the original names. For example, if the general and slow query logs are named `mysql.log` and `mysql-slow.log`, you can use a series of commands like this:

```
shell> cd mysql-data-directory
shell> mv mysql.log mysql.old
shell> mv mysql-slow.log mysql-slow.old
shell> mysqladmin flush-logs
```

At this point, you can make a backup of `mysql.old` and `mysql-slow.log` and then remove them from disk.

On Windows, you cannot rename log files while the server has them open. You must stop the server and rename them, and then restart the server to create new logs.

## 5.12. Running Multiple MySQL Servers on the Same Machine

In some cases, you might want to run multiple `mysqld` servers on the same machine. You might want to test a new MySQL release while leaving your existing production setup undisturbed. Or you might want to give different users access to different `mysqld` servers that they manage themselves. (For example, you might be an Internet Service Provider that wants to provide independent MySQL installations for different customers.)

To run multiple servers on a single machine, each server must have unique values for several operating parameters. These can be set on the command line or in option files. See [Section 4.3, “Specifying Program Options”](#).

At least the following options must be different for each server:

- `--port=port_num`  
`--port` controls the port number for TCP/IP connections.
- `--socket=path`  
`--socket` controls the Unix socket file path on Unix and the name of the named pipe on Windows. On Windows, it is necessary to specify distinct pipe names only for those servers that support named-pipe connections.
- `--shared-memory-base-name=name`



The name of shared memory to use for shared-memory connections. This option is available only on Windows. The default name is `MYSQL`. The name is case sensitive. This option was added in MySQL 4.1.

- `--pid-file=file_name`

This option is used only on Unix. It indicates the pathname of the file in which the server writes its process ID.

If you use the following log file options, they must be different for each server:

- `--log=file_name`
- `--log-bin=file_name`
- `--log-update=file_name`
- `--log-error=file_name`
- `--log-isam=file_name`
- `--bdb-logdir=file_name`

[Section 5.11.6, “Server Log Maintenance”](#), discusses the log file options further.

For better performance, you can specify the following options differently for each server, to spread the load between several physical disks:

- `--tmpdir=path`
- `--bdb-tmpdir=path`

Having different temporary directories is also recommended to make it easier to determine which MySQL server created any given temporary file.

With very limited exceptions, each server should use a different data directory, which is specified using the `--datadir=path` option.

**Warning:** Normally, you should never have two servers that update data in the same databases. This may lead to unpleasant surprises if your operating system does not support fault-free system locking. If (despite this warning) you run multiple servers using the same data directory and they have logging enabled, you must use the appropriate options to specify log filenames that are unique to each server. Otherwise, the servers try to log to the same files. Please note that this kind of setup only works with `ISAM`, `MyISAM` and `MERGE` tables, and not with any of the other storage engines.

The warning against sharing a data directory among servers also applies in an NFS environment. Allowing multiple MySQL servers to access a common data directory over NFS is a *very bad idea*.

- The primary problem is that NFS is the speed bottleneck. It is not meant for such use.
- Another risk with NFS is that you must devise a way to ensure that two or more servers do not interfere with each other. Usually NFS file locking is handled by the `lockd` daemon, but at the moment there is no platform that performs locking 100% reliably in every situation.

Make it easy for yourself: Forget about sharing a data directory among servers over NFS. A better solution is to have one computer that contains several CPUs and use an operating system that handles threads efficiently.

If you have multiple MySQL installations in different locations, you can specify the base installation directory for each server with the `--basedir=path` option to cause each server to use a different data directory, log files, and PID file. (The defaults for all these values are determined relative to the base directory). In that case, the only other options you need to specify are the `--socket` and `--port` options. For example, suppose that you install different versions of MySQL using tarfile binary distributions. These install in different locations, so you can start the server for each installation using `bin/mysqld_safe` under its own corresponding base directory. `mysqld_safe` determines the proper `--basedir` option to pass to `mysqld`, and you need specify only the `--socket` and `--port` options to `mysqld_safe`. (For versions of MySQL older than 4.0, use `safe_mysqld` rather than `mysqld_safe`.)

As discussed in the following sections, it is possible to start additional servers by setting environment variables or by specifying appropriate command-line options. However, if you need to run multiple servers on a more permanent basis, it is more convenient to use option files to specify for each server those option values that must be unique to it. The `--defaults-file` option is useful for this purpose.

## 5.12.1. Running Multiple Servers on Windows

You can run multiple servers on Windows by starting them manually from the command line, each with appropriate operating parameters. On Windows NT-based systems, you also have the option of installing several servers as Windows services and running them that way. General instructions for running MySQL servers from the command line or as services are given in [Section 2.3, “Installing MySQL on Windows”](#). This section describes how to make sure that you start each server with different values for those startup options that must be unique per server, such as the data directory. These options are described in [Section 5.12, “Running Multiple MySQL Servers on the Same Machine”](#).

### 5.12.1.1. Starting Multiple Windows Servers at the Command Line

To start multiple servers manually from the command line, you can specify the appropriate options on the command line or in an option file. It is more convenient to place the options in an option file, but it is necessary to make sure that each server gets its own set of options. To do this, create an option file for each server and tell the server the filename with a `--defaults-file` option when you run it.

Suppose that you want to run `mysqld` on port 3307 with a data directory of `C:\mydata1`, and `mysqld-max` on port 3308 with a data directory of `C:\mydata2`. (To do this, make sure that before you start the servers, each data directory exists and has its own copy of the `mysql` database that contains the grant tables.) Then create two option files. For example, create one file named `C:\my-opts1.cnf` that looks like this:

```
[mysqld]
datadir = C:/mydata1
port = 3307
```

Create a second file named `C:\my-opts2.cnf` that looks like this:

```
[mysqld]
datadir = C:/mydata2
port = 3308
```

Then start each server with its own option file:

```
C:\> C:\mysql\bin\mysqld --defaults-file=C:\my-opts1.cnf
C:\> C:\mysql\bin\mysqld-max --defaults-file=C:\my-opts2.cnf
```

On NT, each server starts in the foreground (no new prompt appears until the server exits later), so you will need to issue those two commands in separate console windows.

To shut down the servers, you must connect to each using the appropriate port number:

```
C:\> C:\mysql\bin\mysqladmin --port=3307 shutdown
C:\> C:\mysql\bin\mysqladmin --port=3308 shutdown
```

Servers configured as just described allow clients to connect over TCP/IP. If your version of Windows supports named pipes and you also want to allow named-pipe connections, use the `mysqld-nt` or `mysqld-max-nt` servers and specify options that enable the named pipe and specify its name. Each server that supports named-pipe connections must use a unique pipe name. For example, the `C:\my-opts1.cnf` file might be written like this:

```
[mysqld]
datadir = C:/mydata1
port = 3307
enable-named-pipe
socket = mypipe1
```

Then start the server this way:

```
C:\> C:\mysql\bin\mysqld-nt --defaults-file=C:\my-opts1.cnf
```

Modify `C:\my-opts2.cnf` similarly for use by the second server.

A similar procedure applies for servers that you want to support shared-memory connections. Enable such connections with the `--shared-memory` option and specify a unique shared-memory name for each server with the `--shared-memory-base-name` option.

### 5.12.1.2. Starting Multiple Windows Servers as Services

On NT-based systems, a MySQL server can run as a Windows service. The procedures for installing, controlling, and removing a single MySQL service are described in [Section 2.3.11, “Starting MySQL as a Windows Service”](#).

As of MySQL 4.0.2, you can install multiple servers as services. In this case, you must make sure that each server uses a different service name in addition to all the other parameters that must be unique per server.

For the following instructions, assume that you want to run the `mysqld-nt` server from two different versions of MySQL that are installed at `C:\mysql-4.0.8` and `C:\mysql-4.0.17`, respectively. (This might be the case if you are running 4.0.8 as your production server, but want to test 4.0.17 before upgrading to it.)

The following principles apply when installing a MySQL service with the `--install` or `--install-manual` option:

- If you specify no service name, the server uses the default service name of `MySQL` and the server reads options from the `[mysqld]` group in the standard option files.
- If you specify a service name after the `--install` option, the server ignores the `[mysqld]` option group and instead reads options from the group that has the same name as the service. The server reads options from the standard option files.
- If you specify a `--defaults-file` option after the service name, the server ignores the standard option files and reads options only from the `[mysqld]` group of the named file.

**Note:** Before MySQL 4.0.17, only a server installed using the default service name (MySQL) or one installed explicitly with a service name of `mysqld` will read the `[mysqld]` group in the standard option files. As of 4.0.17, all servers read the `[mysqld]` group if they read the standard option files, even if they are installed using another service name. This allows you to use the `[mysqld]` group for options that should be used by all MySQL services, and an option group named after each service for use by the server installed with that service name.

Based on the preceding information, you have several ways to set up multiple services. The following instructions describe some examples. Before trying any of them, be sure that you shut down and remove any existing MySQL services first.

- **Approach 1:** Specify the options for all services in one of the standard option files. To do this, use a different service name for each server. Suppose that you want to run the 4.0.8 `mysqld-nt` using the service name of `mysqld1` and the 4.0.17 `mysqld-nt` using the service name `mysqld2`. In this case, you can use the `[mysqld1]` group for 4.0.8 and the `[mysqld2]` group for 4.0.17. For example, you can set up `C:\my.cnf` like this:

```
# options for mysqld1 service
[mysqld1]
basedir = C:/mysql-4.0.8
port = 3307
enable-named-pipe
socket = mypipe1

# options for mysqld2 service
[mysqld2]
basedir = C:/mysql-4.0.17
port = 3308
enable-named-pipe
socket = mypipe2
```

Install the services as follows, using the full server pathnames to ensure that Windows registers the correct executable program for each service:

```
C:\> C:\mysql-4.0.8\bin\mysqld-nt --install mysqld1
C:\> C:\mysql-4.0.17\bin\mysqld-nt --install mysqld2
```

To start the services, use the services manager, or use `NET START` with the appropriate service names:

```
C:\> NET START mysqld1
C:\> NET START mysqld2
```

To stop the services, use the services manager, or use `NET STOP` with the appropriate service names:

```
C:\> NET STOP mysqld1
C:\> NET STOP mysqld2
```

- **Approach 2:** Specify options for each server in separate files and use `--defaults-file` when you install the services to tell each server what file to use. In this case, each file should list options using a `[mysqld]` group.

With this approach, to specify options for the 4.0.8 `mysqld-nt`, create a file `C:\my-opts1.cnf` that looks like this:

```
[mysqld]
basedir = C:/mysql-4.0.8
port = 3307
enable-named-pipe
socket = mypipe1
```

For the 4.0.17 `mysqld-nt`, create a file `C:\my-opts2.cnf` that looks like this:

```
[mysqld]
basedir = C:/mysql-4.0.17
port = 3308
enable-named-pipe
socket = mypipe2
```

Install the services as follows (enter each command on a single line):

```
C:\> C:\mysql-4.0.8\bin\mysqld-nt --install mysqld1
      --defaults-file=C:\my-opts1.cnf
C:\> C:\mysql-4.0.17\bin\mysqld-nt --install mysqld2
      --defaults-file=C:\my-opts2.cnf
```

To use a `--defaults-file` option when you install a MySQL server as a service, you must precede the option with the service name.

After installing the services, start and stop them the same way as in the preceding example.

To remove multiple services, use `mysqld --remove` for each one, specifying a service name following the `--remove` option. If the service name is the default (MySQL), you can omit it.

## 5.12.2. Running Multiple Servers on Unix

The easiest way to run multiple servers on Unix is to compile them with different TCP/IP ports and Unix socket files so that each one is listening on different network interfaces. Compiling in different base directories for each installation also results automatically in a separate, compiled-in data directory, log file, and PID file location for each server.

Assume that an existing server is configured for the default TCP/IP port number (3306) and Unix socket file (`/tmp/mysql.sock`). To configure a new server to have different operating parameters, use a `configure` command something like this:

```
shell> ./configure --with-tcp-port=port_number \
                --with-unix-socket-path=file_name \
                --prefix=/usr/local/mysql-4.0.17
```

Here, `port_number` and `file_name` must be different from the default TCP/IP port number and Unix socket file pathname, and the `--prefix` value should specify an installation directory different from the one under which the existing MySQL installation is located.

If you have a MySQL server listening on a given port number, you can use the following command to find out what operating parameters it is using for several important configurable variables, including the base directory and Unix socket filename:

```
shell> mysqladmin --host=host_name --port=port_number variables
```

With the information displayed by that command, you can tell what option values *not* to use when configuring an additional server.

Note that if you specify `localhost` as a hostname, `mysqladmin` defaults to using a Unix socket file connection rather than TCP/IP. In MySQL 4.1, you can explicitly specify the connection protocol to use by using the `--protocol={TCP|SOCKET|PIPE|MEMORY}` option.

You don't have to compile a new MySQL server just to start with a different Unix socket file and TCP/IP port number. It is also possible to use the same server binary and start each invocation of it with dif-

ferent parameter values at runtime. One way to do so is by using command-line options:

```
shell> mysqld_safe --socket=file_name --port=port_number
```

To start a second server, provide different `--socket` and `--port` option values, and pass a `--datadir=path` option to `mysqld_safe` so that the server uses a different data directory.

Another way to achieve a similar effect is to use environment variables to set the Unix socket filename and TCP/IP port number:

```
shell> MYSQL_UNIX_PORT=/tmp/mysqld-new.sock
shell> MYSQL_TCP_PORT=3307
shell> export MYSQL_UNIX_PORT MYSQL_TCP_PORT
shell> mysql_install_db --user=mysql
shell> mysqld_safe --datadir=/path/to/datadir &
```

This is a quick way of starting a second server to use for testing. The nice thing about this method is that the environment variable settings apply to any client programs that you invoke from the same shell. Thus, connections for those clients are automatically directed to the second server.

[Appendix F, \*Environment Variables\*](#), includes a list of other environment variables you can use to affect `mysqld`.

For automatic server execution, the startup script that is executed at boot time should run the following command once for each server with an appropriate option file path for each command:

```
shell> mysqld_safe --defaults-file=file_name
```

Each option file should contain option values specific to a given server.

On Unix, the `mysqld_multi` script is another way to start multiple servers. See [Section 5.4.3](#), “`mysqld_multi` — Manage Multiple MySQL Servers”.

### 5.12.3. Using Client Programs in a Multiple-Server Environment

To connect with a client program to a MySQL server that is listening to different network interfaces from those compiled into your client, you can use one of the following methods:

- Start the client with `--host=host_name --port=port_number` to connect via TCP/IP to a remote server, with `--host=127.0.0.1 --port=port_number` to connect via TCP/IP to a local server, or with `--host=localhost --socket=file_name` to connect to a local server via a Unix socket file or a Windows named pipe.
- As of MySQL 4.1, start the client with `--protocol=tcp` to connect via TCP/IP, `--protocol=socket` to connect via a Unix socket file, `--protocol=pipe` to connect via a named pipe, or `--protocol=memory` to connect via shared memory. For TCP/IP connections, you may also need to specify `--host` and `--port` options. For the other types of connections, you may need to specify a `--socket` option to specify a Unix socket file or Windows named-pipe name, or a `--shared-memory-base-name` option to specify the shared-memory name. Shared-memory connections are supported only on Windows.
- On Unix, set the `MYSQL_UNIX_PORT` and `MYSQL_TCP_PORT` environment variables to point to the Unix socket file and TCP/IP port number before you start your clients. If you normally use a specific socket file or port number, you can place commands to set these environment variables in your `.login` file so that they apply each time you log in. See [Appendix F, \*Environment Variables\*](#).

- Specify the default Unix socket file and TCP/IP port number in the `[client]` group of an option file. For example, you can use `C:\my.cnf` on Windows, or the `.my.cnf` file in your home directory on Unix. See [Section 4.3.2, “Using Option Files”](#).
- In a C program, you can specify the socket file or port number arguments in the `mysql_real_connect()` call. You can also have the program read option files by calling `mysql_options()`. See [Section 17.2.3, “C API Function Descriptions”](#).
- If you are using the Perl `DBD::mysql` module, you can read options from MySQL option files. For example:

```
$dsn = "DBI:mysql:test;mysql_read_default_group=client;"  
      . "mysql_read_default_file=/usr/local/mysql/data/my.cnf";  
$dbh = DBI->connect($dsn, $user, $password);
```

See [Section 17.4, “MySQL Perl API”](#).

Other programming interfaces may provide similar capabilities for reading option files.

## 5.13. The MySQL Query Cache

From version 4.0.1 on, MySQL Server features a query cache. When in use, the query cache stores the text of a `SELECT` statement together with the corresponding result that was sent to the client. If an identical statement is received later, the server retrieves the results from the query cache rather than parsing and executing the statement again.

The query cache is extremely useful in an environment where you have tables that do not change very often and for which the server receives many identical queries. This is a typical situation for many Web servers that generate many dynamic pages based on database content.

**Note:** The query cache does not return stale data. When tables are modified, any relevant entries in the query cache are flushed.

**Note:** The query cache does not work in an environment where you have multiple `mysqld` servers updating the same `MyISAM` tables.

**Note:** The query cache is not used for server-side prepared statements. If you are using server-side prepared statements consider that these statements will not be satisfied by the query cache. See [Section 17.2.4, “C API Prepared Statements”](#).

Some performance data for the query cache follows. These results were generated by running the MySQL benchmark suite on a Linux Alpha 2×500MHz system with 2GB RAM and a 64MB query cache.

- If all the queries you are performing are simple (such as selecting a row from a table with one row), but still differ so that the queries cannot be cached, the overhead for having the query cache active is 13%. This could be regarded as the worst case scenario. In real life, queries tend to be much more complicated, so the overhead normally is significantly lower.
- Searches for a single row in a single-row table are 238% faster with the query cache than without it. This can be regarded as close to the minimum speedup to be expected for a query that is cached.

To disable the query cache at server startup, set the `query_cache_size` system variable to 0. By disabling the query cache code, there is no noticeable overhead. If you build MySQL from source, query cache capabilities can be excluded from the server entirely by invoking `configure` with the `-`

`-without-query-cache` option.

## 5.13.1. How the Query Cache Operates

This section describes how the query cache works when it is operational. [Section 5.13.3, “Query Cache Configuration”](#), describes how to control whether it is operational.

Incoming queries are compared to those in the query cache before parsing, so the following two queries are regarded as different by the query cache:

```
SELECT * FROM tbl_name
Select * from tbl_name
```

Queries must be *exactly* the same (byte for byte) to be seen as identical. In addition, query strings that are identical may be treated as different for other reasons. Queries that use different databases, different protocol versions, or different default character sets are considered different queries and are cached separately.

Before a query result is fetched from the query cache, MySQL checks that the user has `SELECT` privilege for all databases and tables involved. If this is not the case, the cached result is not used.

If a query result is returned from query cache, the server increments the `Qcache_hits` status variable, not `Com_select`. See [Section 5.13.4, “Query Cache Status and Maintenance”](#).

If a table changes, all cached queries that use the table become invalid and are removed from the cache. This includes queries that use `MERGE` tables that map to the changed table. A table can be changed by many types of statements, such as `INSERT`, `UPDATE`, `DELETE`, `TRUNCATE`, `ALTER TABLE`, `DROP TABLE`, or `DROP DATABASE`.

Cache entries for transactional `InnoDB` tables that have been changed are invalidated when a `COMMIT` is performed.

In MySQL 4.0, the query cache is disabled within transactions (it does not return results). Beginning with MySQL 4.1.1, the query cache also works within transactions when using `InnoDB` tables (it uses the table version number to detect whether its contents are still current).

A query that begins with a leading comment may be cached, but cannot be fetched from the cache.

The query cache works for `SELECT SQL_CALC_FOUND_ROWS . . .` and `SELECT FOUND_ROWS ( )` type queries. `FOUND_ROWS ( )` returns the correct value even if the preceding query was fetched from the cache because the number of found rows is also stored in the cache.

A query cannot be cached if it contains any of the functions shown in the following table:

<code>BENCHMARK ( )</code>	<code>CONNECTION_ID ( )</code>	<code>CURDATE ( )</code>
<code>CURRENT_DATE ( )</code>	<code>CURRENT_TIME ( )</code>	<code>CURRENT_TIMESTAMP ( )</code>
<code>CURTIME ( )</code>	<code>DATABASE ( )</code>	<code>ENCRYPT ( )</code> with one parameter
<code>FOUND_ROWS ( )</code>		<code>GET_LOCK ( )</code>
<code>LAST_INSERT_ID ( )</code>	<code>LOAD_FILE ( )</code>	<code>MASTER_POS_WAIT ( )</code>
<code>NOW ( )</code>	<code>RAND ( )</code>	<code>RELEASE_LOCK ( )</code>
<code>SYSDATE ( )</code>	<code>UNIX_TIMESTAMP ( )</code> with no parameters	<code>USER ( )</code>

A query also is not cached under these conditions:



- It refers to user-defined functions (UDFs).
- It refers to user variables.
- It refers to tables in the `mysql` system database.
- It is of any of the following forms:

```
SELECT ... IN SHARE MODE
SELECT ... FOR UPDATE
SELECT ... INTO OUTFILE ...
SELECT ... INTO DUMPFILE ...
SELECT * FROM ... WHERE autoincrement_col IS NULL
```

The last form is not cached because it is used as the ODBC workaround for obtaining the last insert ID value. See the MyODBC section of [Chapter 18, Connectors](#).

- It was issued as a prepared statement, even if no placeholders were employed. For example, the query used here is not cached:

```
char *my_sql_stmt = "SELECT a, b FROM table_c";
/* ... */
mysql_stmt_prepare(stmt, my_sql_stmt, strlen(my_sql_stmt));
```

See [Section 17.2.4, “C API Prepared Statements”](#).

- It uses `TEMPORARY` tables.
- It does not use any tables.
- The user has a column-level privilege for any of the involved tables.

### 5.13.2. Query Cache `SELECT` Options

Two query cache-related options may be specified in `SELECT` statements:

- `SQL_CACHE`

The query result is cached if the value of the `query_cache_type` system variable is `ON` or `DEMAND`.

- `SQL_NO_CACHE`

The query result is not cached.

Examples:

```
SELECT SQL_CACHE id, name FROM customer;
SELECT SQL_NO_CACHE id, name FROM customer;
```

### 5.13.3. Query Cache Configuration

The `have_query_cache` server system variable indicates whether the query cache is available:

```
mysql> SHOW VARIABLES LIKE 'have_query_cache';
+-----+-----+
| Variable_name | Value |
+-----+-----+
```

```
+-----+
| have_query_cache | YES |
+-----+
```

Several other system variables control query cache operation. These can be set in an option file or on the command line when starting `mysqld`. The query cache system variables all have names that begin with `query_cache_`. They are described briefly in [Section 5.2.2, “System Variables”](#), with additional configuration information given here.

To set the size of the query cache, set the `query_cache_size` system variable. Setting it to 0 disables the query cache. The default size is 0, so the query cache is disabled by default.

When you set `query_cache_size` to a non-zero value, keep in mind that the query cache needs a minimum size of about 40KB to allocate its structures. (The exact size depends on system architecture.) If you set the value too small, you'll get a warning, as in this example:

```
mysql> SET GLOBAL query_cache_size = 40000;
Query OK, 0 rows affected, 1 warning (0.00 sec)

mysql> SHOW WARNINGS\G
***** 1. row *****
Level: Warning
Code: 1282
Message: Query cache failed to set size 39936; new query cache size is 0

mysql> SET GLOBAL query_cache_size = 41984;
Query OK, 0 rows affected (0.00 sec)

mysql> SHOW VARIABLES LIKE 'query_cache_size';
+-----+
| Variable_name | Value |
+-----+
| query_cache_size | 41984 |
+-----+
```

If the query cache size is greater than 0, the `query_cache_type` variable influences how it works. This variable can be set to the following values:

- A value of 0 or `OFF` prevents caching or retrieval of cached results.
- A value of 1 or `ON` allows caching except of those statements that begin with `SELECT SQL_NO_CACHE`.
- A value of 2 or `DEMAND` causes caching of only those statements that begin with `SELECT SQL_CACHE`.

Setting the `GLOBAL query_cache_type` value determines query cache behavior for all clients that connect after the change is made. Individual clients can control cache behavior for their own connection by setting the `SESSION query_cache_type` value. For example, a client can disable use of the query cache for its own queries like this:

```
mysql> SET SESSION query_cache_type = OFF;
```

To control the maximum size of individual query results that can be cached, set the `query_cache_limit` system variable. The default value is 1MB.

When a query is to be cached, its result (the data sent to the client) is stored in the query cache during result retrieval. Therefore the data usually is not handled in one big chunk. The query cache allocates blocks for storing this data on demand, so when one block is filled, a new block is allocated. Because memory allocation operation is costly (timewise), the query cache allocates blocks with a minimum size given by the `query_cache_min_res_unit` system variable. When a query is executed, the last

result block is trimmed to the actual data size so that unused memory is freed. Depending on the types of queries your server executes, you might find it helpful to tune the value of `query_cache_min_res_unit`:

- The default value of `query_cache_min_res_unit` is 4KB. This should be adequate for most cases.
- If you have a lot of queries with small results, the default block size may lead to memory fragmentation, as indicated by a large number of free blocks. Fragmentation can force the query cache to prune (delete) queries from the cache due to lack of memory. In this case, you should decrease the value of `query_cache_min_res_unit`. The number of free blocks and queries removed due to pruning are given by the values of the `Qcache_free_blocks` and `Qcache_lowmem_prunes` status variables.
- If most of your queries have large results (check the `Qcache_total_blocks` and `Qcache_queries_in_cache` status variables), you can increase performance by increasing `query_cache_min_res_unit`. However, be careful to not make it too large (see the previous item).

`query_cache_min_res_unit` is present from MySQL 4.1.

### 5.13.4. Query Cache Status and Maintenance

You can check whether the query cache is present in your MySQL server using the following statement:

```
mysql> SHOW VARIABLES LIKE 'have_query_cache';
+-----+-----+
| Variable_name | Value |
+-----+-----+
| have_query_cache | YES  |
+-----+-----+
```

You can defragment the query cache to better utilize its memory with the `FLUSH QUERY CACHE` statement. The statement does not remove any queries from the cache.

The `RESET QUERY CACHE` statement removes all query results from the query cache. The `FLUSH TABLES` statement also does this.

To monitor query cache performance, use `SHOW STATUS` to view the cache status variables:

```
mysql> SHOW STATUS LIKE 'Qcache%';
+-----+-----+
| Variable_name | Value |
+-----+-----+
| Qcache_free_blocks | 36    |
| Qcache_free_memory | 138488 |
| Qcache_hits       | 79570 |
| Qcache_inserts    | 27087 |
| Qcache_lowmem_prunes | 3114  |
| Qcache_not_cached  | 22989 |
| Qcache_queries_in_cache | 415  |
| Qcache_total_blocks | 912   |
+-----+-----+
```

Descriptions of each of these variables are given in [Section 5.2.4, “Status Variables”](#). Some uses for them are described here.

The total number of `SELECT` queries is given by this formula:

```
Com_select
+ Qcache_hits
+ queries with errors found by parser
```

The `Com_select` value is given by this formula:

```
Qcache_inserts
+ Qcache_not_cached
+ queries with errors found during the column-privileges check
```

The query cache uses variable-length blocks, so `Qcache_total_blocks` and `Qcache_free_blocks` may indicate query cache memory fragmentation. After `FLUSH QUERY CACHE`, only a single free block remains.

Every cached query requires a minimum of two blocks (one for the query text and one or more for the query results). Also, every table that is used by a query requires one block. However, if two or more queries use the same table, only one table block needs to be allocated.

The information provided by the `Qcache_lowmem_prunes` status variable can help you tune the query cache size. It counts the number of queries that have been removed from the cache to free up memory for caching new queries. The query cache uses a least recently used (LRU) strategy to decide which queries to remove from the cache. Tuning information is given in [Section 5.13.3, “Query Cache Configuration”](#).

---

# Chapter 6. Replication

Replication capabilities allowing the databases on one MySQL server to be duplicated on another were introduced in MySQL 3.23.15. This chapter describes the various replication features provided by MySQL. It introduces replication concepts, shows how to set up replication servers, and serves as a reference to the available replication options. It also provides a list of frequently asked questions (with answers), and troubleshooting advice for solving problems.

For a description of the syntax of replication-related SQL statements, see [Section 13.6, “Replication Statements”](#).

## 6.1. Introduction to Replication

MySQL 3.23.15 and up features support for one-way, asynchronous replication, in which one server acts as the master, while one or more other servers act as slaves. This is in contrast to the *synchronous* replication which is a characteristic of MySQL Cluster (see [Chapter 15, MySQL Cluster](#)).

In single-master replication, the master server writes updates to its binary log files and maintains an index of those files to keep track of log rotation. The binary log files serve as a record of updates to be sent to any slave servers. When a slave connects to its master, it informs the master of the position up to which the slave read the logs at its last successful update. The slave receives any updates that have taken place since that time, and then blocks and waits for the master to notify it of new updates.

A slave server can itself serve as a master if you want to set up chained replication servers.

When you are using replication, all updates to the tables that are replicated should be performed on the master server. Otherwise, you must always be careful to avoid conflicts between updates that users make to tables on the master and updates that they make to tables on the slave.

Replication offers benefits for robustness, speed, and system administration:

- Robustness is increased with a master/slave setup. In the event of problems with the master, you can switch to the slave as a backup.
- Better response time for clients can be achieved by splitting the load for processing client queries between the master and slave servers. `SELECT` queries may be sent to the slave to reduce the query processing load of the master. Statements that modify data should still be sent to the master so that the master and slave do not get out of synchrony. This load-balancing strategy is effective if non-updating queries dominate, but that is the normal case.
- Another benefit of using replication is that you can perform database backups using a slave server without disturbing the master. The master continues to process updates while the backup is being made. See [Section 5.9.1, “Database Backups”](#).

## 6.2. Replication Implementation Overview

MySQL replication is based on the master server keeping track of all changes to your databases (updates, deletes, and so on) in its binary logs. Therefore, to use replication, you must enable binary logging on the master server. See [Section 5.11.4, “The Binary Log”](#).

Each slave server receives from the master the saved updates that the master has recorded in its binary log, so that the slave can execute the same updates on its copy of the data.

It is *extremely* important to realize that the binary log is simply a record starting from the fixed point in

time at which you enable binary logging. Any slaves that you set up need copies of the databases on your master *as they existed at the moment you enabled binary logging on the master*. If you start your slaves with databases that are not in the same state as those on the master when the binary log was started, your slaves are quite likely to fail.

After the slave has been set up with a copy of the master's data, it connects to the master and waits for updates to process. If the master fails, or the slave loses connectivity with your master, the slave keeps trying to connect periodically until it is able to resume listening for updates. The `-master-connect-retry` option controls the retry interval. The default is 60 seconds.

Each slave keeps track of where it left off when it last read from its master server. The master has no knowledge of how many slaves it has or which ones are up to date at any given time.

## 6.3. Replication Implementation Details

MySQL replication capabilities are implemented using three threads (one on the master server and two on the slave). When a `START SLAVE` statement is issued on a slave server, the slave creates an I/O thread, which connects to the master and asks it to send the updates recorded in its binary logs. The master creates a thread to send the binary log contents to the slave. This thread can be identified as the `Binlog Dump` thread in the output of `SHOW PROCESSLIST` on the master. The slave I/O thread reads the updates that the master `Binlog Dump` thread sends and copies them to local files, known as *relay logs*, in the slave's data directory. The third thread is the SQL thread, which the slave creates to read the relay logs and to execute the updates they contain.

In the preceding description, there are three threads per master/slave connection. A master that has multiple slaves creates one thread for each currently-connected slave, and each slave has its own I/O and SQL threads.

For versions of MySQL before 4.0.2, replication involves only two threads (one on the master and one on the slave). The slave I/O and SQL threads are combined as a single thread, and no relay log files are used.

The slave uses two threads so that reading updates from the master and executing them can be separated into two independent tasks. Thus, the task of reading statements is not slowed down if statement execution is slow. For example, if the slave server has not been running for a while, its I/O thread can quickly fetch all the binary log contents from the master when the slave starts, even if the SQL thread lags far behind. If the slave stops before the SQL thread has executed all the fetched statements, the I/O thread has at least fetched everything so that a safe copy of the statements is stored locally in the slave's relay logs, ready for execution the next time that the slave starts. This enables the master server to purge its binary logs sooner because it no longer needs to wait for the slave to fetch their contents.

The `SHOW PROCESSLIST` statement provides information that tells you what is happening on the master and on the slave regarding replication. The following example illustrates how the three threads show up in the output from `SHOW PROCESSLIST`. The output format is that used by `SHOW PROCESSLIST` as of MySQL version 4.0.15, when the content of the `State` column was changed to be more meaningful compared to earlier versions.

On the master server, the output from `SHOW PROCESSLIST` looks like this:

```
mysql> SHOW PROCESSLIST\G
***** 1. row *****
  Id: 2
  User: root
  Host: localhost:32931
  db: NULL
Command: Binlog Dump
  Time: 94
  State: Has sent all binlog to slave; waiting for binlog to
        be updated
  Info: NULL
```

Here, thread 2 is a [Binlog Dump](#) replication thread for a connected slave. The [State](#) information indicates that all outstanding updates have been sent to the slave and that the master is waiting for more updates to occur. If you see no [Binlog Dump](#) threads on a master server, this means that replication is not running — that is, that no slaves are currently connected.

On the slave server, the output from `SHOW PROCESSLIST` looks like this:

```
mysql> SHOW PROCESSLIST\G
***** 1. row *****
  Id: 10
  User: system user
  Host:
  db: NULL
Command: Connect
  Time: 11
  State: Waiting for master to send event
  Info: NULL
***** 2. row *****
  Id: 11
  User: system user
  Host:
  db: NULL
Command: Connect
  Time: 11
  State: Has read all relay log; waiting for the slave I/O
        thread to update it
  Info: NULL
```

This information indicates that thread 10 is the I/O thread that is communicating with the master server, and thread 11 is the SQL thread that is processing the updates stored in the relay logs. At the time that the `SHOW PROCESSLIST` was run, both threads were idle, waiting for further updates.

The value in the [Time](#) column can show how late the slave is compared to the master. See [Section 6.10](#), “[Replication FAQ](#)”.

### 6.3.1. Replication Master Thread States

The following list shows the most common states you may see in the [State](#) column for the master's [Binlog Dump](#) thread. If you see no [Binlog Dump](#) threads on a master server, this means that replication is not running — that is, that no slaves are currently connected.

- [Sending binlog event to slave](#)

Binary logs consist of *events*, where an event is usually an update plus some other information. The thread has read an event from the binary log and is now sending it to the slave.

- [Finished reading one binlog; switching to next binlog](#)

The thread has finished reading a binary log file and is opening the next one to send to the slave.

- [Has sent all binlog to slave; waiting for binlog to be updated](#)

The thread has read all outstanding updates from the binary logs and sent them to the slave. The thread is now idle, waiting for new events to appear in the binary log resulting from new updates occurring on the master.

- [Waiting to finalize termination](#)

A very brief state that occurs as the thread is stopping.

### 6.3.2. Replication Slave I/O Thread States

The following list shows the most common states you see in the `State` column for a slave server I/O thread. Beginning with MySQL 4.1.1, this state also appears in the `Slave_IO_State` column displayed by `SHOW SLAVE STATUS`, so you can get a good view of what is happening by using that statement.

- `Connecting to master`

The thread is attempting to connect to the master.

- `Checking master version`

A state that occurs very briefly, after the connection to the master is established.

- `Registering slave on master`

A state that occurs very briefly after the connection to the master is established.

- `Requesting binlog dump`

A state that occurs very briefly, after the connection to the master is established. The thread sends to the master a request for the contents of its binary logs, starting from the requested binary log file-name and position.

- `Waiting to reconnect after a failed binlog dump request`

If the binary log dump request failed (due to disconnection), the thread goes into this state while it sleeps, then tries to reconnect periodically. The interval between retries can be specified using the `-master-connect-retry` option.

- `Reconnecting after a failed binlog dump request`

The thread is trying to reconnect to the master.

- `Waiting for master to send event`

The thread has connected to the master and is waiting for binary log events to arrive. This can last for a long time if the master is idle. If the wait lasts for `slave_read_timeout` seconds, a timeout occurs. At that point, the thread considers the connection to be broken and makes an attempt to reconnect.

- `Queueing master event to the relay log`

The thread has read an event and is copying it to the relay log so that the SQL thread can process it.

- `Waiting to reconnect after a failed master event read`

An error occurred while reading (due to disconnection). The thread is sleeping for `master-connect-retry` seconds before attempting to reconnect.

- `Reconnecting after a failed master event read`

The thread is trying to reconnect to the master. When connection is established again, the state becomes `Waiting for master to send event`.

- `Waiting for the slave SQL thread to free enough relay log space`

You are using a non-zero `relay_log_space_limit` value, and the relay logs have grown large enough that their combined size exceeds this value. The I/O thread is waiting until the SQL thread frees enough space by processing relay log contents so that it can delete some relay log files.



- `Waiting for slave mutex on exit`

A state that occurs briefly as the thread is stopping.

### 6.3.3. Replication Slave SQL Thread States

The following list shows the most common states you may see in the `State` column for a slave server SQL thread:

- `Reading event from the relay log`

The thread has read an event from the relay log so that the event can be processed.

- `Has read all relay log; waiting for the slave I/O thread to update it`

The thread has processed all events in the relay log files, and is now waiting for the I/O thread to write new events to the relay log.

- `Waiting for slave mutex on exit`

A very brief state that occurs as the thread is stopping.

The `State` column for the I/O thread may also show the text of a statement. This indicates that the thread has read an event from the relay log, extracted the statement from it, and is executing it.

### 6.3.4. Replication Relay and Status Files

By default, relay logs filenames have the form `host_name-relay-bin.nnnnnn`, where `host_name` is the name of the slave server host and `nnnnnn` is a sequence number. Successive relay log files are created using successive sequence numbers, beginning with `000001` (`001` in MySQL 4.0 or older). The slave uses an index file to track the relay log files currently in use. The default relay log index filename is `host_name-relay-bin.index`. By default, the slave server creates relay log files in its data directory. The default filenames can be overridden with the `--relay-log` and `--relay-log-index` server options. See [Section 6.8, “Replication Startup Options”](#).

Relay logs have the same format as binary logs and can be read using `mysqlbinlog`. The SQL thread automatically deletes each relay log file as soon as it has executed all events in the file and no longer needs it. There is no explicit mechanism for deleting relay logs because the SQL thread takes care of doing so. However, as of MySQL 4.0.14, `FLUSH LOGS` rotates relay logs, which influences when the SQL thread deletes them.

A slave server creates a new relay log file under the following conditions:

- Each time the I/O thread starts.
- When the logs are flushed; for example, with `FLUSH LOGS` or `mysqladmin flush-logs`. (This creates a new relay log only as of MySQL 4.0.14.)
- When the size of the current relay log file becomes too large. The meaning of “too large” is determined as follows:
  - If the value of `max_relay_log_size` is greater than 0, that is the maximum relay log file size.

- If the value of `max_relay_log_size` is 0, `max_binlog_size` determines the maximum relay log file size. `max_binlog_size` always determines the relay log size before MySQL 4.0.14, the first version in which `max_relay_log_size` appears.

A slave replication server creates two additional small files in the data directory. These *status files* are named `master.info` and `relay-log.info` by default. Their names can be changed by using the `--master-info-file` and `--relay-log-info-file` options. See Section 6.8, “Replication Startup Options”.

The two status files contain information like that shown in the output of the `SHOW SLAVE STATUS` statement, which is discussed in Section 13.6.2, “SQL Statements for Controlling Slave Servers”. Because the status files are stored on disk, they survive a slave server's shutdown. The next time the slave starts up, it reads the two files to determine how far it has proceeded in reading binary logs from the master and in processing its own relay logs.

The I/O thread updates the `master.info` file. Before MySQL 4.1, the following table shows the correspondence between the lines in the file and the columns displayed by `SHOW SLAVE STATUS`.

Line	Description
1	<code>Master_Log_File</code>
2	<code>Read_Master_Log_Pos</code>
3	<code>Master_Host</code>
4	<code>Master_User</code>
5	Password (not shown by <code>SHOW SLAVE STATUS</code> )
6	<code>Master_Port</code>
7	<code>Connect_Retry</code>

As of MySQL 4.1, the file includes a line count and information about SSL options:

Line	Description
1	Number of lines in the file
2	<code>Master_Log_File</code>
3	<code>Read_Master_Log_Pos</code>
4	<code>Master_Host</code>
5	<code>Master_User</code>
6	Password (not shown by <code>SHOW SLAVE STATUS</code> )
7	<code>Master_Port</code>
8	<code>Connect_Retry</code>
9	<code>Master_SSL_Allowed</code>
10	<code>Master_SSL_CA_File</code>
11	<code>Master_SSL_CA_Path</code>
12	<code>Master_SSL_Cert</code>
13	<code>Master_SSL_Cipher</code>
14	<code>Master_SSL_Key</code>

The SQL thread updates the `relay-log.info` file. The following table shows the correspondence between the lines in the file and the columns displayed by `SHOW SLAVE STATUS`.

Line	Description
1	<code>Relay_Log_File</code>
2	<code>Relay_Log_Pos</code>
3	<code>Relay_Master_Log_File</code>
4	<code>Exec_Master_Log_Pos</code>

The contents of the `relay-log.info` file and the states shown by the `SHOW SLAVE STATES` command may not match if the `relay-log.info` file has not been flushed to disk. Ideally, you should only view `relay-log.info` on a slave that is offline (i.e. `mysqld` is not running). For a running system, `SHOW SLAVE STATUS` should be used.

When you back up the slave's data, you should back up these two status files as well, along with the relay log files. They are needed to resume replication after you restore the slave's data. If you lose the relay logs but still have the `relay-log.info` file, you can check it to determine how far the SQL thread has executed in the master binary logs. Then you can use `CHANGE MASTER TO` with the `MASTER_LOG_FILE` and `MASTER_LOG_POS` options to tell the slave to re-read the binary logs from that point. Of course, this requires that the binary logs still exist on the master server.

If your slave is subject to replicating `LOAD DATA INFILE` statements, you should also back up any `SQL_LOAD-*` files that exist in the directory that the slave uses for this purpose. The slave needs these files to resume replication of any interrupted `LOAD DATA INFILE` operations. The directory location is specified using the `--slave-load-tmpdir` option. If this option is not specified, the directory location is the value of the `tmpdir` system variable.

## 6.4. How to Set Up Replication

This section briefly describes how to set up complete replication of a MySQL server. It assumes that you want to replicate all databases on the master and have not previously configured replication. You must shut down your master server briefly to complete the steps outlined here.

This procedure is written in terms of setting up a single slave, but you can repeat it to set up multiple slaves.

Although this method is the most straightforward way to set up a slave, it is not the only one. For example, if you have a snapshot of the master's data, and the master already has its server ID set and binary logging enabled, you can set up a slave without shutting down the master or even blocking updates to it. For more details, please see [Section 6.10, "Replication FAQ"](#).

If you want to administer a MySQL replication setup, we suggest that you read this entire chapter through and try all statements mentioned in [Section 13.6.1, "SQL Statements for Controlling Master Servers"](#), and [Section 13.6.2, "SQL Statements for Controlling Slave Servers"](#). You should also familiarize yourself with the replication startup options described in [Section 6.8, "Replication Startup Options"](#).

**Note:** This procedure and some of the replication SQL statements shown in later sections refer to the `SUPER` privilege. Prior to MySQL 4.0.2, use the `PROCESS` privilege instead.

1. Make sure that you have a recent version of MySQL installed on the master and slaves, and that these versions are compatible according to the table shown in [Section 6.5, "Replication Compatibility Between MySQL Versions"](#).

If you encounter a problem, please do not report it as a bug until you have verified that the problem is present in the latest MySQL release.

2. Set up an account on the master server that the slave server can use to connect. This account must be given the `REPLICATION SLAVE` privilege. If the account is used only for replication (which is recommended), you don't need to grant any additional privileges.

Suppose that your domain is `mydomain.com` and that you want to create an account with a username of `repl` such that slave servers can use the account to access the master server from any host in your domain using a password of `slavepass`. To create the account, use this `GRANT` statement:

```
mysql> GRANT REPLICATION SLAVE ON *.*
-> TO 'repl'@'%.mydomain.com' IDENTIFIED BY 'slavepass';
```

For MySQL versions older than 4.0.2, the `REPLICATION SLAVE` privilege does not exist. Grant the `FILE` privilege instead:

```
mysql> GRANT FILE ON *.*
-> TO 'repl'@'%.mydomain.com' IDENTIFIED BY 'slavepass';
```

For additional information about setting up user accounts and privileges, see [Section 5.8, “MySQL User Account Management”](#).

3. Flush all the tables and block write statements by executing a `FLUSH TABLES WITH READ LOCK` statement:

```
mysql> FLUSH TABLES WITH READ LOCK;
```

For `InnoDB` tables, note that `FLUSH TABLES WITH READ LOCK` blocks `COMMIT` operations, too. (This is true as of MySQL version 4.0.20.) When you have acquired your global read lock, you can start a filesystem snapshot of your `InnoDB` tables. Internally (inside the `InnoDB` storage engine) the snapshot won't be consistent (because the `InnoDB` caches are not flushed), but there's no need to worry at all, because `InnoDB` will resolve this at startup, and consequently deliver a consistent result. This means that `InnoDB` will perform a crash recovery when started on this snapshot, but it will not be corrupted. If you want to have a consistent snapshot of your `InnoDB` tables, there's no way around taking down the MySQL server, though.

Leave running the client from which you issue the `FLUSH TABLES` statement so that the read lock remains in effect. (If you exit the client, the lock is released.) Then take a snapshot of the data on your master server.

The easiest way to create a snapshot is to use an archiving program to make a binary backup of the databases in your master's data directory. For example, use `tar` on Unix, or `PowerArchiver`, `WinRAR`, `WinZip`, or any similar software on Windows. To use `tar` to create an archive that includes all databases, change location into the master server's data directory, then execute this command:

```
shell> tar -cvf /tmp/mysql-snapshot.tar .
```

If you want the archive to include only a database called `this_db`, use this command instead:

```
shell> tar -cvf /tmp/mysql-snapshot.tar ./this_db
```

Then copy the archive file to the `/tmp` directory on the slave server host. On that machine, change location into the slave's data directory, and unpack the archive file using this command:

```
shell> tar -xvf /tmp/mysql-snapshot.tar
```

You may not want to replicate the `mysql` database if the slave server has a different set of user accounts from those that exist on the master. In this case, you should exclude it from the archive. You also need not include any log files in the archive, or the `master.info` or `relay-log.info` files.

While the read lock placed by `FLUSH TABLES WITH READ LOCK` is in effect, read the value of the current binary log name and offset on the master:

```
mysql > SHOW MASTER STATUS;
+-----+-----+-----+-----+
| File           | Position | Binlog_Do_DB | Binlog_Ignore_DB |
+-----+-----+-----+-----+
| mysql-bin.003 | 73       | test         | manual,mysql      |
+-----+-----+-----+-----+
```

The `File` column shows the name of the log and `Position` shows the offset within the file. In this example, the binary log file is `mysql-bin.003` and the offset is 73. Record these values. You need them later when you are setting up the slave. They represent the replication coordinates at which the slave should begin processing new updates from the master.

If the master has been running previously without binary logging enabled, the log name and position values displayed by `SHOW MASTER STATUS` or `mysqldump --master-data` will be empty. In that case, the values that you need to use later when specifying the slave's log file and position are the empty string ( ' ' ) and 4.

After you have taken the snapshot and recorded the log name and offset, you can re-enable write activity on the master:

```
mysql> UNLOCK TABLES;
```

If you are using `InnoDB` tables, ideally you should use the `InnoDB Hot Backup` tool, which takes a consistent snapshot without acquiring any locks on the master server, and records the log name and offset corresponding to the snapshot to be later used on the slave. `Hot Backup` is an additional non-free (commercial) tool that is not included in the standard MySQL distribution. See the `InnoDB Hot Backup` home page at <http://www.innodb.com/manual.php> for detailed information.

Without the `Hot Backup` tool, the quickest way to take a binary snapshot of `InnoDB` tables is to shut down the master server and copy the `InnoDB` data files, log files, and table format files (`.frm` files). To record the current log file name and offset, you should issue the following statements before you shut down the server:

```
mysql> FLUSH TABLES WITH READ LOCK;
mysql> SHOW MASTER STATUS;
```

Then record the log name and the offset from the output of `SHOW MASTER STATUS` as was shown earlier. After recording the log name and the offset, shut down the server *without* unlocking the tables to make sure that the server goes down with the snapshot corresponding to the current log file and offset:

```
shell> mysqladmin -u root shutdown
```

An alternative that works for both `MyISAM` and `InnoDB` tables is to take an SQL dump of the master instead of a binary copy as described in the preceding discussion. For this, you can use `mysqldump --master-data` on your master and later load the SQL dump file into your slave.

However, this is slower than doing a binary copy.

4. Make sure that the `[mysqld]` section of the `my.cnf` file on the master host includes a `log-bin` option. The section should also have a `server-id=master_id` option, where `master_id` must be a positive integer value from 1 to  $2^{32} - 1$ . For example:

```
[mysqld]
log-bin=mysql-bin
server-id=1
```

If those options are not present, add them and restart the server. The server cannot act as a replication master unless binary logging is enabled.

Note: For the greatest possible durability and consistency in a replication setup using InnoDB with transactions, you should use `innodb_flush_log_at_trx_commit=1`, `sync_binlog=1`, and `innodb_safe_binlog` in your master `my.cnf` file.

5. Stop the server that is to be used as a slave and add the following lines to its `my.cnf` file:

```
[mysqld]
server-id=slave_id
```

The `slave_id` value, like the `master_id` value, must be a positive integer value from 1 to  $2^{32} - 1$ . In addition, it is necessary that the ID of the slave be different from the ID of the master. For example:

```
[mysqld]
server-id=2
```

If you are setting up multiple slaves, each one must have a unique `server-id` value that differs from that of the master and from each of the other slaves. Think of `server-id` values as something similar to IP addresses: These IDs uniquely identify each server instance in the community of replication partners.

If you do not specify a `server-id` value, it is set to 1 if you have not defined `master-host`; otherwise it is set to 2. Note that in the case of `server-id` omission, a master refuses connections from all slaves, and a slave refuses to connect to a master. Thus, omitting `server-id` is good only for backup with a binary log.

6. If you made a binary backup of the master server's data, copy it to the slave server's data directory before starting the slave. Make sure that the privileges on the files and directories are correct. The system account that you use to run the slave server must be able to read and write the files, just as on the master.

If you made a backup using `mysqldump`, start the slave first. The dump file is loaded in a later step.

7. Start the slave server. If it has been replicating previously, start the slave server with the `-skip-slave-start` option so that it doesn't immediately try to connect to its master. You also may want to start the slave server with the `--log-warnings` option to get more messages in the error log about problems (for example, network or connection problems). The option is enabled by default as of MySQL 4.0.19 and 4.1.2, but as of MySQL 4.0.21 and 4.1.3, aborted connections are not logged to the error log unless the value is greater than 1.
8. If you made a backup of the master server's data using `mysqldump`, load the dump file into the slave server:

```
shell> mysql -u root -p < dump_file.sql
```

9. Execute the following statement on the slave, replacing the option values with the actual values relevant to your system:

```
mysql> CHANGE MASTER TO
->     MASTER_HOST='master_host_name',
->     MASTER_USER='replication_user_name',
->     MASTER_PASSWORD='replication_password',
->     MASTER_LOG_FILE='recorded_log_file_name',
->     MASTER_LOG_POS=recorded_log_position;
```

The following table shows the maximum allowable length for the string-valued options:

MASTER_HOST	60
MASTER_USER	16
MASTER_PASSWORD	32
MASTER_LOG_FILE	255

10. Start the slave threads:

```
mysql> START SLAVE;
```

After you have performed this procedure, the slave should connect to the master and catch up on any updates that have occurred since the snapshot was taken.

If you have forgotten to set the `server-id` option for the master, slaves cannot connect to it.

If you have forgotten to set the `server-id` option for the slave, you get the following error in the slave's error log:

```
Warning: You should set server-id to a non-0 value if master_host
is set; we will force server id to 2, but this MySQL server will
not act as a slave.
```

You also find error messages in the slave's error log if it is not able to replicate for any other reason.

Once a slave is replicating, you can find in its data directory one file named `master.info` and another named `relay-log.info`. The slave uses these two files to keep track of how much of the master's binary log it has processed. Do *not* remove or edit these files unless you know exactly what you are doing and fully understand the implications. Even in that case, it is preferred that you use the `CHANGE MASTER TO` statement to change replication parameters. The slave will use the values specified in the statement to update the status files automatically.

**Note:** The content of `master.info` overrides some of the server options specified on the command line or in `my.cnf`. See [Section 6.8, “Replication Startup Options”](#), for more details.

Once you have a snapshot of the master, you can use it to set up other slaves by following the slave portion of the procedure just described. You do not need to take another snapshot of the master; you can use the same one for each slave.

## 6.5. Replication Compatibility Between MySQL Versions

The original binary log format was developed in MySQL 3.23. It was changed in MySQL 4.0.

**Note:** You *cannot* replicate from a master that uses a newer binary log format to a slave that uses an older format (for example, from MySQL 4.1 to MySQL 3.23.) This has significant consequences for upgrading servers in a replication setup, as described in [Section 6.6, “Upgrading a Replication Setup”](#).

As far as replication is concerned, the binary log format for any MySQL 4.1.x version and any 4.0.x version is identical. However, replication from a 4.1 master to a 4.0 slave is unsupported, has not been tested thoroughly, and no further development or bug fixing is planned for this master/slave combination. Although the binary log format is the same for 4.0 and 4.1, there are other constraints, such as SQL-level compatibility issues. For example, a 4.1 master cannot replicate to a 4.0 slave if the replicated statements use SQL features available in 4.1 but not 4.0. These and other issues are discussed in [Section 6.7, “Replication Features and Known Problems”](#).

As a general rule, we recommended using recent MySQL versions, because replication capabilities are continually being improved. We also recommend using the same version for both the master and the slave. We recommend upgrading masters and slaves running alpha or beta versions to new versions.

## 6.6. Upgrading a Replication Setup

When you upgrade servers that participate in a replication setup, the procedure for upgrading depends on the current server versions and the version to which you are upgrading.

### 6.6.1. Upgrading Replication to 4.0 or 4.1

This section applies to upgrading replication from MySQL 3.23 to 4.0 or 4.1. A 4.0 server should be 4.0.3 or newer, as mentioned in [Section 6.5, “Replication Compatibility Between MySQL Versions”](#).

When you upgrade a master from MySQL 3.23 to MySQL 4.0 or 4.1, you should first ensure that all the slaves of this master are at 4.0 or 4.1. If that is not the case, you should first upgrade your slaves: Shut down each one, upgrade it, restart it, and restart replication.

The upgrade can safely be done using the following procedure, assuming that you have a 3.23 master to upgrade and the slaves are 4.0 or 4.1. Note that after the master has been upgraded, you should not restart replication using any old 3.23 binary logs, because this unfortunately confuses the 4.0 or 4.1 slaves.

1. Block all updates on the master by issuing a `FLUSH TABLES WITH READ LOCK` statement.
2. Wait until all the slaves have caught up with all changes from the master server. Use `SHOW MASTER STATUS` on the master to obtain its current binary log file and position. Then, for each slave, use those values with a `SELECT MASTER_POS_WAIT()` statement. The statement blocks on the slave and returns when the slave has caught up. Then run `STOP SLAVE` on the slave.
3. Stop the master server and upgrade it to MySQL 4.0 or 4.1.
4. Restart the master server and record the name of its newly created binary log. You can obtain the name of the file by issuing a `SHOW MASTER STATUS` statement on the master. Then issue these statements on each slave:

```
mysql> CHANGE MASTER TO MASTER_LOG_FILE='binary_log_name',
->     MASTER_LOG_POS=4;
mysql> START SLAVE;
```

## 6.7. Replication Features and Known Problems

In general, replication compatibility at the SQL level requires that any features used be supported by both the master and the slave servers. If you use a feature on a master server that is available only as of a



given version of MySQL, you cannot replicate to a slave that is older than that version. Such incompatibilities are likely to occur between series, so that, for example, you cannot replicate from MySQL 4.1 to 4.0. However, these incompatibilities also can occur for within-series replication. For example, the `CONVERT_TZ()` function is available in MySQL 4.1.3 and up. If you use this function on the master server, you cannot replicate to a slave server that is older than MySQL 4.1.3.

The following list provides details about what is supported and what is not. Additional InnoDB-specific information about replication is given in [Section 14.2.7.5, “InnoDB and MySQL Replication”](#).

- Replication of `AUTO_INCREMENT`, `LAST_INSERT_ID()`, and `TIMESTAMP` values is done correctly.

However, adding an `AUTO_INCREMENT` column to a table with `ALTER TABLE` might not produce the same ordering of the rows on the slave and the master. This occurs because the order in which the rows are numbered depends on the specific storage engine used for the table and the order in which the rows were inserted. If it is important to have the same order on the master and slave, the rows must be ordered before assigning an `AUTO_INCREMENT` number. Assuming that you want to add an `AUTO_INCREMENT` column to the table `t1`, the following statements produce a new table `t2` identical to `t1` but with an `AUTO_INCREMENT` column:

```
CREATE TABLE t2 LIKE t1;
ALTER TABLE t2 ADD id INT AUTO_INCREMENT PRIMARY KEY;
INSERT INTO t2 SELECT * FROM t1 ORDER BY col1, col2;
```

This assumes that the table `t1` has columns `col1` and `col2`.

**Important:** To guarantee the same ordering on both master and slave, *all* columns of `t1` must be referenced in the `ORDER BY` clause.

The instructions just given are subject to the limitations of `CREATE TABLE ... LIKE`: Foreign key definitions are ignored, as are the `DATA DIRECTORY` and `INDEX DIRECTORY` table options. If a table definition includes any of those characteristics, create `t2` using a `CREATE TABLE` statement that is identical to the one used to create `t1`, but with the addition of the `AUTO_INCREMENT` column.

Regardless of the method used to create and populate the copy having the `AUTO_INCREMENT` column, the final step is to drop the original table and then rename the copy:

```
DROP t1;
ALTER TABLE t2 RENAME t1;
```

See also [Section A.7.1, “Problems with ALTER TABLE”](#).

- The `USER()`, `UUID()`, and `LOAD_FILE()` functions are replicated without change and thus do not work reliably on the slave. This is also true for `CONNECTION_ID()` in slave versions older than 4.1.1. The **new** `PASSWORD()` function in MySQL 4.1 is well replicated in masters from 4.1.1 and up; your slaves also must be 4.1.1 or above to replicate it. If you have older slaves and need to replicate `PASSWORD()` from your 4.1.x master, you must start your master with the `-old-password` option, so that it uses the old implementation of `PASSWORD()`. (Note that the `PASSWORD()` implementation in MySQL 4.1.0 differs from every other version of MySQL. It is best to avoid 4.1.0 in a replication scenario.)
- User privileges are replicated only if the `mysql` database is replicated. That is, the `GRANT`, `REVOKE`, `SET PASSWORD`, and `DROP USER` (available as of MySQL 4.1.1) statements take effect on the slave only if the replication setup includes the `mysql` database.

If you're replicating all databases, but don't want statements that affect user privileges to be replicated, set up the slave to not replicate the `mysql` database, using the `-`

---

`--replicate-wild-ignore-table=mysql.%` option. That option is available as of MySQL 4.0.13. The slave will recognize that issuing privilege-related SQL statements won't have an effect, and thus not execute those statements.

- The `GET_LOCK()`, `RELEASE_LOCK()`, `IS_FREE_LOCK()`, and `IS_USED_LOCK()` functions that handle user-level locks are replicated without the slave knowing the concurrency context on master. Therefore, these functions should not be used to insert into a master's table because the content on the slave would differ. (For example, do not issue a statement such as `INSERT INTO mytable VALUES(GET_LOCK(...))`.)
- The `FOREIGN_KEY_CHECKS` variable is replicated as of MySQL 4.0.14. The `SQL_MODE`, `UNIQUE_CHECKS`, `SQL_AUTO_IS_NULL`, and `storage_engine` (also known as `table_type`) variables are not replicated in MySQL 4.1 or earlier versions.
- The following applies to replication between MySQL servers that use different character sets:
  1. You must *always* use the same *global* character set and collation on the master and the slave. (These are controlled by the `--character-set-server` and `--collation-server` options.) Otherwise, you may get duplicate-key errors on the slave, because a key that is unique in the master character set might not be unique in the slave character set.
  2. If the master is older than MySQL 4.1.3, the character set of any client should never be made different from its global value because this character set change is not known to the slave. In other words, clients should not use `SET NAMES`, `SET CHARACTER SET`, and so forth. If both the master and the slave are 4.1.3 or newer, clients can freely set session values for character set variables because these settings are written to the binary log and so are known to the slave. That is, clients can use `SET NAMES` or `SET CHARACTER SET` or can set variables such as `collation_client` or `collation_server`. However, clients are prevented from changing the *global* value of these variables; as stated previously, the master and slave must always have identical global character set values.
  3. If on the master you have databases with different character sets from the global `collation_server` value, you should design your `CREATE TABLE` statements so that they do not implicitly rely on the default database's character set, because there currently is a bug ([Bug#2326](http://bugs.mysql.com/2326) [http://bugs.mysql.com/2326]); a good workaround is to state the character set and collation explicitly in `CREATE TABLE`.
- The same system time zone should be set for both master and slave. Otherwise some statements will not be replicated properly, such as statements that use the `NOW()` or `FROM_UNIXTIME()` functions. You can set the time zone in which MySQL server runs by using the `--timezone=timezone_name` option of the `mysqld_safe` script or by setting the `TZ` environment variable. Also starting from version 4.1.3 both master and slave should have the same default connection time zone set, that is the `--default-time-zone` parameter should have the same value for both master and slave.
- `CONVERT_TZ(..., ..., @@global.time_zone)` is not properly replicated.
- Session variables are not replicated properly when used in statements which update tables; for example: `SET MAX_JOIN_SIZE=1000; INSERT INTO mytable VALUES(@@MAX_JOIN_SIZE);` will not insert the same data on the master and on the slave.
- It is possible to replicate transactional tables on the master using non-transactional tables on the slave. For example, you can replicate an `InnoDB` master table as a `MyISAM` slave table. However, if you do this, there are problems if the slave is stopped in the middle of a `BEGIN/COMMIT` block because the slave restarts at the beginning of the `BEGIN` block.
- Update statements that refer to user-defined variables (that is, variables of the form `@var_name`) are badly replicated in 3.23 and 4.0. This is fixed in 4.1.

- The slave can connect to the master using SSL if both are 4.1.1 or newer.
- Starting from MySQL 4.1.11, there is a global system variable `slave_transaction_retries`: If the replication slave SQL thread fails to execute a transaction because of an InnoDB deadlock or because it exceeded the InnoDB `innodb_lock_wait_timeout` or the NDBCluster `TransactionDeadlockDetectionTimeout` or `TransactionInactiveTimeout` value, the transaction automatically retries `slave_transaction_retries` times before stopping with an error. The default value is 0 in MySQL 4.1. Starting from MySQL 4.1.11, the total retry count can be seen in `SHOW STATUS`; see [Section 5.2.4, “Status Variables”](#).
- If a `DATA DIRECTORY` or `INDEX DIRECTORY` table option is used in a `CREATE TABLE` statement on the master server, the table option is also used on the slave. This can cause problems if no corresponding directory exists in the slave host filesystem or if it exists but is not accessible to the slave server. As of MySQL 4.0.15, there is an `sql_mode` option called `NO_DIR_IN_CREATE`. If the slave server is run with this SQL mode enabled, it ignores the `DATA DIRECTORY` and `INDEX DIRECTORY` table options when replicating `CREATE TABLE` statements. The result is that `MyISAM` data and index files are created in the table's database directory.
- It is possible for the data on the master and slave to become different if a statement is designed in such a way that the data modification is non-deterministic; that is, left to the will of the query optimizer. (This is in general not a good practice, even outside of replication.) For a detailed explanation of this issue, see [Section A.8.4, “Open Issues in MySQL”](#).
- If on the master a `LOAD DATA INFILE` is interrupted (for example, by a integrity constraint violation or a killed connection), the slave skips this `LOAD DATA INFILE` entirely. This means that if this command permanently inserted or updated table records before being interrupted, these modifications are *not* replicated to the slave.
- `LOAD DATA INFILE` does not replicate correctly when `--binlog-do-db` is used. ([Bug#19662](#) [<http://bugs.mysql.com/19662>])
- Before MySQL 4.1.1, the `FLUSH`, `ANALYZE TABLE`, `OPTIMIZE TABLE`, and `REPAIR TABLE` statements are not written to the binary log and thus are not replicated to the slaves. This is not normally a problem because these statements do not modify table data. However, it can cause difficulties under certain circumstances. If you replicate the privilege tables in the `mysql` database and update those tables directly without using the `GRANT` statement, you must issue a `FLUSH PRIVILEGES` statement on your slaves to put the new privileges into effect. Also if you use `FLUSH TABLES` when renaming a `MyISAM` table that is part of a `MERGE` table, you have to issue `FLUSH TABLES` manually on the slaves. As of MySQL 4.1.1, these statements are written to the binary log (unless you specify `NO_WRITE_TO_BINLOG`, or its alias `LOCAL`). Exceptions are that `FLUSH LOGS`, `FLUSH MASTER`, `FLUSH SLAVE`, and `FLUSH TABLES WITH READ LOCK` are not logged in any case. (Any of these may cause problems if replicated to a slave.) For a syntax example, see [Section 13.5.5.2, “FLUSH Syntax”](#).
- MySQL 4.1 and earlier support only replication scenarios involving one master and many slaves.
- When a server shuts down and restarts, its `MEMORY (HEAP)` tables become empty. As of MySQL 4.0.18, the master replicates this effect to slaves as follows: The first time that the master uses each `MEMORY` table after startup, it logs an event that notifies the slaves that the table needs to be emptied by writing a `DELETE` statement for that table to the binary log. See [Section 14.4, “The MEMORY \(HEAP\) Storage Engine”](#), for more information.
- Temporary tables are replicated except in the case where you shut down the slave server (not just the slave threads) and you have replicated temporary tables that are used in updates that have not yet been executed on the slave. If you shut down the slave server, the temporary tables needed by those updates are no longer available when the slave is restarted. To avoid this problem, do not shut down the slave while it has temporary tables open. Instead, use the following procedure:

1. Issue a `STOP SLAVE` statement.
  2. Use `SHOW STATUS` to check the value of the `Slave_open_temp_tables` variable.
  3. If the value is 0, issue a `mysqladmin shutdown` command to stop the slave.
  4. If the value is not 0, restart the slave threads with `START SLAVE`.
  5. Repeat the procedure later until the `Slave_open_temp_tables` variable is 0 and you can stop the slave.
- The syntax for multiple-table `DELETE` statements that use table aliases changed between MySQL 4.0 and 4.1. In MySQL 4.0, you should use the true table name to refer to any table from which rows should be deleted:

```
DELETE test FROM test AS t1, test2 WHERE ...
```

In MySQL 4.1, you must use the alias:

```
DELETE t1 FROM test AS t1, test2 WHERE ...
```

If you use such `DELETE` statements, the change in syntax means that a 4.0 master cannot replicate to 4.1 (or higher) slaves.

- It is safe to connect servers in a circular master/slave relationship if you use the `-log-slave-updates` option. That means that you can create a setup such as this:

```
A -> B -> C -> A
```

However, many statements do not work correctly in this kind of setup unless your client code is written to take care of the potential problems that can occur from updates that occur in different sequence on different servers.

Server IDs are encoded in binary log events, so server A knows when an event that it reads was originally created by itself and does not execute the event (unless server A was started with the `-replicate-same-server-id` option, which is meaningful only in rare cases). Thus, there are no infinite loops. This type of circular setup works only if you perform no conflicting updates between the tables. In other words, if you insert data in both A and C, you should never insert a row in A that may have a key that conflicts with a row inserted in C. You should also not update the same rows on two servers if the order in which the updates are applied is significant.

- If a statement on a slave produces an error, the slave SQL thread terminates, and the slave writes a message to its error log. You should then connect to the slave manually and determine the cause of the problem. (`SHOW SLAVE STATUS` is useful for this.) Then fix the problem (for example, you might need to create a non-existent table) and run `START SLAVE`.
- It is safe to shut down a master server and restart it later. When a slave loses its connection to the master, the slave tries to reconnect immediately and retries periodically if that fails. The default is to retry every 60 seconds. This may be changed with the `--master-connect-retry` option. A slave also is able to deal with network connectivity outages. However, the slave notices the network outage only after receiving no data from the master for `slave_net_timeout` seconds. If your outages are short, you may want to decrease `slave_net_timeout`. See [Section 5.2.2, “System Variables”](#).
- Shutting down the slave (cleanly) is also safe because it keeps track of where it left off. Unclean shutdowns might produce problems, especially if the disk cache was not flushed to disk before the system went down. Your system fault tolerance is greatly increased if you have a good uninterrupt-

ible power supply. Unclean shutdowns of the master may cause inconsistencies between the content of tables and the binary log in master; this can be avoided by using `InnoDB` tables and the `--innodb-safe-binlog` option on the master. See [Section 5.11.4, “The Binary Log”](#).

- A crash on the master side can result in the master's binary log having a final position less than the most recent position read by the slave, due to the master's binary log file not being flushed. This can cause the slave not to be able to replicate when the master comes back up. Setting `sync_binlog=1` in the master `my.cnf` file helps to minimize this problem because it causes the master to flush its binary log more frequently.
- Due to the non-transactional nature of `MyISAM` tables, it is possible to have a statement that only partially updates a table and returns an error code. This can happen, for example, on a multiple-row insert that has one row violating a key constraint, or if a long update statement is killed after updating some of the rows. If that happens on the master, the slave thread exits and waits for the database administrator to decide what to do about it unless the error code is legitimate and execution of the statement results in the same error code on the slave. If this error code validation behavior is not desirable, some or all errors can be masked out (ignored) with the `--slave-skip-errors` option. This option is available starting with MySQL 3.23.47.
- If you update transactional tables from non-transactional tables inside a `BEGIN/COMMIT` sequence, updates to the binary log may be out of synchrony with table states if the non-transactional table is updated before the transaction commits. This occurs because the transaction is written to the binary log only when it is committed.
- Before version 4.0.15, any update to a non-transactional table is written to the binary log at once when the update is made, whereas transactional updates are written on `COMMIT` or not written at all if you use `ROLLBACK`. You must take this into account when updating both transactional tables and non-transactional tables within the same transaction. (This is true not only for replication, but also if you are using binary logging for backups.)

As of version 4.0.15, we changed the logging behavior for transactions that mix updates to transactional and non-transactional tables, which solves the problems (order of statements is good in the binary log, and all needed statements are written to the binary log even in case of `ROLLBACK`). The problem that remains is that when a second connection updates the non-transactional table while the first connection's transaction is not finished yet, incorrect ordering can still occur because the second connection's update is written immediately after it is done.

- When a 4.x slave replicates a `LOAD DATA INFILE` from a 3.23 master, the values of the `Exec_Master_Log_Pos` and `Relay_Log_Space` columns of `SHOW SLAVE STATUS` become incorrect. The inaccuracy in `Exec_Master_Log_Pos` causes problems when you stop and restart replication; so it is a good idea to correct the value before this, by doing `FLUSH LOGS` on the master.

The following table lists replication problems in MySQL 3.23 that are fixed in MySQL 4.0:

- `LOAD DATA INFILE` is handled properly, as long as the data file still resides on the master server at the time of update propagation.
- `LOAD DATA LOCAL INFILE` is no longer skipped on the slave as it was in 3.23.
- In 3.23, `RAND()` in updates does not replicate properly. Use `RAND(some_non_rand_expr)` if you are replicating updates with `RAND()`. You can, for example, use `UNIX_TIMESTAMP()` as the argument to `RAND()`.
- Floating-point values are approximate, so comparisons involving them are inexact. This is true for operations that use floating-point values explicitly, or values that are converted to floating-point implicitly. Comparisons of floating-point values might yield different results on master and slave serv-

ers due to differences in computer architecture, the compiler used to build MySQL, and so forth. See [Section 12.1.2, “Type Conversion in Expression Evaluation”](#), and [Section A.5.8, “Problems with Floating-Point Comparisons”](#).

## 6.8. Replication Startup Options

This section describes the options that you can use on slave replication servers. You can specify these options either on the command line or in an option file.

On the master and each slave, you must use the `server-id` option to establish a unique replication ID. For each server, you should pick a unique positive integer in the range from 1 to  $2^{32} - 1$ , and each ID must be different from every other ID. Example: `server-id=3`

Options that you can use on the master server for controlling binary logging are described in [Section 5.11.4, “The Binary Log”](#).

Some slave server replication options are handled in a special way, in the sense that each is ignored if a `master.info` file exists when the slave starts and contains a value for the option. The following options are handled this way:

- `--master-host`
- `--master-user`
- `--master-password`
- `--master-port`
- `--master-connect-retry`

As of MySQL 4.1.1, the following options also are handled specially:

- `--master-ssl`
- `--master-ssl-ca`
- `--master-ssl-capath`
- `--master-ssl-cert`
- `--master-ssl-cipher`
- `--master-ssl-key`

The `master.info` file format in 4.1.1 changed to include values corresponding to the SSL options. In addition, the 4.1.1 file format includes as its first line the number of lines in the file. (See [Section 6.3.4, “Replication Relay and Status Files”](#).) If you upgrade an older server to 4.1.1, the new server upgrades the `master.info` file to the new format automatically when it starts. However, if you downgrade a 4.1.1 or newer server to a version older than 4.1.1, you should manually remove the first line before starting the older server for the first time. Note that, in this case, the downgraded server no longer can use an SSL connection to communicate with the master.

If no `master.info` file exists when the slave server starts, it uses the values for those options that are specified in option files or on the command line. This occurs when you start the server as a replication

slave for the very first time, or when you have run `RESET SLAVE` and then have shut down and restarted the slave.

If the `master.info` file exists when the slave server starts, the server uses its contents and ignores any options that correspond to the values listed in the file. Thus, if you start the slave server with different values of the startup options that correspond to values in the `master.info` file, the different values have no effect, because the server continues to use the `master.info` file. To use different values, you must either restart after removing the `master.info` file or (preferably) use the `CHANGE MASTER TO` statement to reset the values while the slave is running.

Suppose that you specify this option in your `my.cnf` file:

```
[mysqld]
master-host=some_host
```

The first time you start the server as a replication slave, it reads and uses that option from the `my.cnf` file. The server then records the value in the `master.info` file. The next time you start the server, it reads the master host value from the `master.info` file only and ignores the value in the option file. If you modify the `my.cnf` file to specify a different master host of `some_other_host`, the change still has no effect. You should use `CHANGE MASTER TO` instead.

Because the server gives an existing `master.info` file precedence over the startup options just described, you might prefer not to use startup options for these values at all, and instead specify them by using the `CHANGE MASTER TO` statement. See [Section 13.6.2.1, “CHANGE MASTER TO Syntax”](#).

This example shows a more extensive use of startup options to configure a slave server:

```
[mysqld]
server-id=2
master-host=db-master.mycompany.com
master-port=3306
master-user=pertinax
master-password=freitag
master-connect-retry=60
report-host=db-slave.mycompany.com
```

The following list describes startup options for controlling replication. Many of these options can be reset while the server is running by using the `CHANGE MASTER TO` statement. Others, such as the `--replicate-*` options, can be set only when the slave server starts.

- `--log-slave-updates`

Normally, a slave does not log to its own binary log any updates that are received from a master server. This option tells the slave to log the updates performed by its SQL thread to its own binary log. For this option to have any effect, the slave must also be started with the `--log-bin` option to enable binary logging. `--log-slave-updates` is used when you want to chain replication servers. For example, you might want to set up replication servers using this arrangement:

```
A -> B -> C
```

Here, A serves as the master for the slave B, and B serves as the master for the slave C. For this to work, B must be both a master *and* a slave. You must start both A and B with `--log-bin` to enable binary logging, and B with the `--log-slave-updates` option so that updates received from A are logged by B to its binary log.

- `--log-warnings[=level]`

This option causes a server to print more messages to the error log about what it is doing. With respect to replication, the server generates warnings that it succeeded in reconnecting after a network/connection failure, and informs you as to how each slave thread started. This option is enabled by

default as of MySQL 4.0.19 and 4.1.2; to disable it, use `--skip-log-warnings`. As of MySQL 4.0.21 and 4.1.3, aborted connections are not logged to the error log unless the value is greater than 1.

Note that the effects of this option are not limited to replication. It produces warnings across a spectrum of server activities.

- `--master-connect-retry=seconds`

The number of seconds that the slave thread sleeps before trying to reconnect to the master in case the master goes down or the connection is lost. The value in the `master.info` file takes precedence if it can be read. If not set, the default is 60. Connection retries are not invoked until the slave times out reading data from the master according to the value of `--slave-net-timeout`. The number of reconnection attempts is limited by the `--master-retry-count` option.

- `--master-host=host_name`

The hostname or IP number of the master replication server. The value in `master.info` takes precedence if it can be read. If no master host is specified, the slave thread does not start.

- `--master-info-file=file_name`

The name to use for the file in which the slave records information about the master. The default name is `master.info` in the data directory.

- `--master-password=password`

The password of the account that the slave thread uses for authentication when it connects to the master. The value in the `master.info` file takes precedence if it can be read. If not set, an empty password is assumed.

- `--master-port=port_number`

The TCP/IP port number that the master is listening on. The value in the `master.info` file takes precedence if it can be read. If not set, the compiled-in setting is assumed (normally 3306).

- `--master-retry-count=count`

The number of times that the slave tries to connect to the master before giving up. Reconnects are attempted at intervals set by `--master-connect-retry` and reconnects are triggered when data reads by the slave time out according to the `--slave-net-timeout` option. The default value is 86400.

- `--master-ssl, --master-ssl-ca=file_name, --master-ssl-capath=directory_name, --master-ssl-cert=file_name, --master-ssl-cipher=cipher_list, --master-ssl-key=file_name`

These options are used for setting up a secure replication connection to the master server using SSL. Their meanings are the same as the corresponding `--ssl, --ssl-ca, --ssl-capath, --ssl-cert, --ssl-cipher, --ssl-key` options that are described in [Section 5.8.7.3, “SSL Command Options”](#). The values in the `master.info` file take precedence if they can be read.

These options are operational as of MySQL 4.1.1.

- `--master-user=user_name`

The username of the account that the slave thread uses for authentication when it connects to the master. This account must have the `REPLICATION SLAVE` privilege. `FILE` privilege instead.)



The value in the `master.info` file takes precedence if it can be read. If the master username is not set, the name `test` is assumed.

- `--max-relay-log-size=size`

The size at which the server rotates relay log files automatically. For more information, see [Section 6.3.4, “Replication Relay and Status Files”](#). Default is 1GB.

This option is available as of MySQL 4.0.14.

- `--read-only`

Cause the slave to allow no updates except from slave threads or from users with the `SUPER` privilege. This enables you to ensure that a slave server accepts no updates from clients.

This option is available as of MySQL 4.0.14.

- `--relay-log=file_name`

The name for the relay log. The default name is `host_name-relay-bin.nnnnnn`, where `host_name` is the name of the slave server host and `nnnnnn` indicates that relay logs are created in numbered sequence. You can specify the option to create hostname-independent relay log names, or if your relay logs tend to be big (and you don't want to decrease `max_relay_log_size`) and you need to put them in some area different from the data directory, or if you want to increase speed by balancing load between disks.

- `--relay-log-index=file_name`

The name to use for the relay log index file. The default name is `host_name-relay-bin.index` in the data directory, where `host_name` is the name of the slave server.

- `--relay-log-info-file=file_name`

The name to use for the file in which the slave records information about the relay logs. The default name is `relay-log.info` in the data directory.

- `--relay-log-purge={0|1}`

Disable or enable automatic purging of relay logs as soon as they are not needed any more. The default value is 1 (enabled). This is a global variable that can be changed dynamically with `SET GLOBAL relay_log_purge = N`.

This option is available as of MySQL 4.1.1.

- `--relay-log-space-limit=size`

This option places an upper limit on the total size in bytes of all relay logs on the slave. A value of 0 means “no limit.” This is useful for a slave server host that has limited disk space. When the limit is reached, the I/O thread stops reading binary log events from the master server until the SQL thread has caught up and deleted some unused relay logs. Note that this limit is not absolute: There are cases where the SQL thread needs more events before it can delete relay logs. In that case, the I/O thread exceeds the limit until it becomes possible for the SQL thread to delete some relay logs, because not doing so would cause a deadlock (which is what happens before MySQL 4.0.13). You should not set `--relay-log-space-limit` to less than twice the value of `--max-relay-log-size` (or `--max-binlog-size` if `--max-relay-log-size` is 0). In that case, there is a chance that the I/O thread waits for free space because `--relay-log-space-limit` is exceeded, but the SQL thread has no relay log to purge and is unable to satisfy the I/O thread. This forces the I/O thread to temporarily ignore `--relay-log-space-limit`.

- `--replicate-do-db=db_name`

Tell the slave to restrict replication to statements where the default database (that is, the one selected by `USE`) is `db_name`. To specify more than one database, use this option multiple times, once for each database. Note that this does not replicate cross-database statements such as `UPDATE some_db.some_table SET foo='bar'` while having selected a different database or no database.

An example of what does not work as you might expect: If the slave is started with `--replicate-do-db=sales` and you issue the following statements on the master, the `UPDATE` statement is *not* replicated:

```
USE prices;
UPDATE sales.january SET amount=amount+1000;
```

The main reason for this “check-just-the-default-database” behavior is that it’s difficult from the statement alone to know whether it should be replicated (for example, if you are using multiple-table `DELETE` or multiple-table `UPDATE` statements that go across multiple databases). It is also faster to check only the default database rather than all databases if there is no need.

If you need cross-database updates to work, make sure that you have MySQL 3.23.28 or later, and use `--replicate-wild-do-table=db_name.%` instead. See [Section 6.9, “How Servers Evaluate Replication Rules”](#).

- `--replicate-do-table=db_name.tbl_name`

Tell the slave thread to restrict replication to the specified table. To specify more than one table, use this option multiple times, once for each table. This works for cross-database updates, in contrast to `--replicate-do-db`. See [Section 6.9, “How Servers Evaluate Replication Rules”](#).

- `--replicate-ignore-db=db_name`

Tells the slave to not replicate any statement where the default database (that is, the one selected by `USE`) is `db_name`. To specify more than one database to ignore, use this option multiple times, once for each database. You should not use this option if you are using cross-database updates and you do not want these updates to be replicated. See [Section 6.9, “How Servers Evaluate Replication Rules”](#).

An example of what does not work as you might expect: If the slave is started with `--replicate-ignore-db=sales` and you issue the following statements on the master, the `UPDATE` statement is *not* replicated:

```
USE prices;
UPDATE sales.january SET amount=amount+1000;
```

If you need cross-database updates to work, use `--replicate-wild-ignore-table=db_name.%` instead. See [Section 6.9, “How Servers Evaluate Replication Rules”](#).

- `--replicate-ignore-table=db_name.tbl_name`

Tells the slave thread to not replicate any statement that updates the specified table, even if any other tables might be updated by the same statement. To specify more than one table to ignore, use this option multiple times, once for each table. This works for cross-database updates, in contrast to `--replicate-ignore-db`. See [Section 6.9, “How Servers Evaluate Replication Rules”](#).

- `--replicate-rewrite-db=from_name->to_name`

Tells the slave to translate the default database (that is, the one selected by `USE`) to `to_name` if it was `from_name` on the master. Only statements involving tables are affected (not statements such

as `CREATE DATABASE`, `DROP DATABASE`, and `ALTER DATABASE`), and only if *from\_name* is the default database on the master. This does not work for cross-database updates. The database name translation is done *before* the `--replicate-*` rules are tested.

If you use this option on the command line and the `'>'` character is special to your command interpreter, quote the option value. For example:

```
shell> mysqld --replicate-rewrite-db="olddb->newdb"
```

- `--replicate-same-server-id`

To be used on slave servers. Usually you should use the default setting of 0, to prevent infinite loops caused by circular replication. If set to 1, the slave does not skip events having its own server ID. Normally, this is useful only in rare configurations. Cannot be set to 1 if `--log-slave-updates` is used. Be careful that starting from MySQL 4.1, by default the slave I/O thread does not even write binary log events to the relay log if they have the slave's server id (this optimization helps save disk usage compared to 4.0). So if you want to use `--replicate-same-server-id` in 4.1 versions, be sure to start the slave with this option before you make the slave read its own events that you want the slave SQL thread to execute.

- `--replicate-wild-do-table=db_name.tbl_name`

Tells the slave thread to restrict replication to statements where any of the updated tables match the specified database and table name patterns. Patterns can contain the `'%'` and `'_'` wildcard characters, which have the same meaning as for the `LIKE` pattern-matching operator. To specify more than one table, use this option multiple times, once for each table. This works for cross-database updates. See [Section 6.9, "How Servers Evaluate Replication Rules"](#).

Example: `--replicate-wild-do-table=foo%.bar%` replicates only updates that use a table where the database name starts with `foo` and the table name starts with `bar`.

If the table name pattern is `%`, it matches any table name and the option also applies to database-level statements (`CREATE DATABASE`, `DROP DATABASE`, and `ALTER DATABASE`). For example, if you use `--replicate-wild-do-table=foo%.%`, database-level statements are replicated if the database name matches the pattern `foo%`.

To include literal wildcard characters in the database or table name patterns, escape them with a backslash. For example, to replicate all tables of a database that is named `my_owndb`, but not replicate tables from the `mylowndb` database, you should escape the `'_'` and `'%'` characters like this: `--replicate-wild-do-table=my\_owndb`. If you're using the option on the command line, you might need to double the backslashes or quote the option value, depending on your command interpreter. For example, with the `bash` shell, you would need to type `--replicate-wild-do-table=my\\_owndb`.

- `--replicate-wild-ignore-table=db_name.tbl_name`

Tells the slave thread not to replicate a statement where any table matches the given wildcard pattern. To specify more than one table to ignore, use this option multiple times, once for each table. This works for cross-database updates. See [Section 6.9, "How Servers Evaluate Replication Rules"](#).

Example: `--replicate-wild-ignore-table=foo%.bar%` does not replicate updates that use a table where the database name starts with `foo` and the table name starts with `bar`.

For information about how matching works, see the description of the `--replicate-wild-do-table` option. The rules for including literal wildcard characters in the option value are the same as for `--replicate-wild-ignore-table` as well.

- `--report-host=slave_name`

The hostname or IP number of the slave to be reported to the master during slave registration. This value appears in the output of `SHOW SLAVE HOSTS` on the master server. Leave the value unset if you do not want the slave to register itself with the master. Note that it is not sufficient for the master to simply read the IP number of the slave from the TCP/IP socket after the slave connects. Due to NAT and other routing issues, that IP may not be valid for connecting to the slave from the master or other hosts.

This option is available as of MySQL 4.0.0.

- `--report-port=slave_port_num`

The TCP/IP port number for connecting to the slave, to be reported to the master during slave registration. Set this only if the slave is listening on a non-default port or if you have a special tunnel from the master or other clients to the slave. If you are not sure, do not use this option.

This option is available as of MySQL 4.0.0.

- `--skip-slave-start`

Tells the slave server not to start the slave threads when the server starts. To start the threads later, use a `START SLAVE` statement.

- `--slave_compressed_protocol={0|1}`

If this option is set to 1, use compression for the slave/master protocol if both the slave and the master support it. Default is 0 (no compression).

- `--slave-load-tmpdir=file_name`

The name of the directory where the slave creates temporary files. This option is by default equal to the value of the `tmpdir` system variable. When the slave SQL thread replicates a `LOAD DATA INFILE` statement, it extracts the file to be loaded from the relay log into temporary files, and then loads these into the table. If the file loaded on the master is huge, the temporary files on the slave are huge, too. Therefore, it might be advisable to use this option to tell the slave to put temporary files in a directory located in some filesystem that has a lot of available space. In that case, the relay logs are huge as well, so you might also want to use the `--relay-log` option to place the relay logs in that filesystem.

The directory specified by this option should be located in a disk-based filesystem (not a memory-based filesystem) because the temporary files used to replicate `LOAD DATA INFILE` must survive machine restarts. The directory also should not be one that is cleared by the operating system during the system startup process.

- `--slave-net-timeout=seconds`

The number of seconds to wait for more data from the master before the slave considers the connection broken, aborts the read, and tries to reconnect. The first retry occurs immediately after the timeout. The interval between retries is controlled by the `--master-connect-retry` option and the number of reconnection attempts is limited by the `--master-retry-count` option. The default is 3600 seconds (one hour).

- `--slave-skip-errors=[err_code1,err_code2,...|all]`

Normally, replication stops when an error occurs on the slave. This gives you the opportunity to resolve the inconsistency in the data manually. This option tells the slave SQL thread to continue replication when a statement returns any of the errors listed in the option value.

Do not use this option unless you fully understand why you are getting errors. If there are no bugs in

your replication setup and client programs, and no bugs in MySQL itself, an error that stops replication should never occur. Indiscriminate use of this option results in slaves becoming hopelessly out of synchrony with the master, with you having no idea why this has occurred.

For error codes, you should use the numbers provided by the error message in your slave error log and in the output of `SHOW SLAVE STATUS`. [Appendix B, Error Codes and Messages](#), lists server error codes.

You can also (but should not) use the very non-recommended value of `all` to cause the slave to ignore all error messages and keeps going regardless of what happens. Needless to say, if you use `all`, there are no guarantees regarding the integrity of your data. Please do not complain (or file bug reports) in this case if the slave's data is not anywhere close to what it is on the master. *You have been warned.*

Examples:

```
--slave-skip-errors=1062,1053
--slave-skip-errors=all
```

## 6.9. How Servers Evaluate Replication Rules

If a master server does not write a statement to its binary log, the statement is not replicated. If the server does log the statement, the statement is sent to all slaves and each slave determines whether to execute it or ignore it.

On the master side, decisions about which statements to log are based on the `--binlog-do-db` and `--binlog-ignore-db` options that control binary logging. For a description of the rules that servers use in evaluating these options, see [Section 5.11.4, “The Binary Log”](#).

On the slave side, decisions about whether to execute or ignore statements received from the master are made according to the `--replicate-*` options that the slave was started with. (See [Section 6.8, “Replication Startup Options”](#).) The slave evaluates these options using the following procedure, which first checks the database-level options and then the table-level options.

In the simplest case, when there are no `--replicate-*` options, the procedure yields the result that the slave executes all statements that it receives from the master. Otherwise, the result depends on the particular options given. In general, to make it easier to determine what effect an option set will have, it is recommended that you avoid mixing “do” and “ignore” options, or wildcard and non-wildcard options.

### Stage 1. Check the database options.

At this stage, the slave checks whether there are any `--replicate-do-db` or `--replicate-ignore-db` options that specify database-specific conditions:

- *No*: Permit the statement and proceed to the table-checking stage.
- *Yes*: Test the options using the same rules as for the `--binlog-do-db` and `--binlog-ignore-db` options to determine whether to permit or ignore the statement. What is the result of the test?
  - *Permit*: Do not execute the statement immediately. Defer the decision and proceed to the table-checking stage.
  - *Ignore*: Ignore the statement and exit.

This stage can permit a statement for further option-checking, or cause it to be ignored. However, statements that are permitted at this stage are not actually executed yet. Instead, they pass to the following stage that checks the table options.

### Stage 2. Check the table options.

If the slave reaches this point, it executes all statements if there are no table options. If there are “do” table options, the statement must match one of them if it is to be executed; otherwise, it is ignored. If there are any “ignore” options, all statements are executed except those that match any `ignore` option. The following steps describe how this evaluation occurs in more detail.

1. Are there any `--replicate-*-table` options?
  - *No*: There are no table restrictions, so all statements match. Execute the statement and exit.
  - *Yes*: There are table restrictions. Evaluate the tables to be updated against them. There might be multiple tables to update, so loop through the following steps for each table looking for a matching option (first the non-wild options, and then the wild options). Only tables that are to be updated are compared to the options. For example, if the statement is `INSERT INTO sales SELECT * FROM prices`, only `sales` is compared to the options). If several tables are to be updated (multiple-table statement), the first table that matches “do” or “ignore” wins. That is, the server checks the first table against the options. If no decision could be made, it checks the second table against the options, and so on.
2. Are there any `--replicate-do-table` options?
  - *No*: Proceed to the next step.
  - *Yes*: Does the table match any of them?
    - *No*: Proceed to the next step.
    - *Yes*: Execute the statement and exit.
3. Are there any `--replicate-ignore-table` options?
  - *No*: Proceed to the next step.
  - *Yes*: Does the table match any of them?
    - *No*: Proceed to the next step.
    - *Yes*: Ignore the statement and exit.
4. Are there any `--replicate-wild-do-table` options?
  - *No*: Proceed to the next step.
  - *Yes*: Does the table match any of them?
    - *No*: Proceed to the next step.
    - *Yes*: Execute the statement and exit.
5. Are there any `--replicate-wild-ignore-table` options?
  - *No*: Proceed to the next step.
  - *Yes*: Does the table match any of them?

- *No*: Proceed to the next step.
  - *Yes*: Ignore the statement and exit.
6. No `--replicate-*-table` option was matched. Is there another table to test against these options?
- *No*: We have now tested all tables to be updated and could not match any option. Are there `--replicate-do-table` or `--replicate-wild-do-table` options?
  - *No*: There were no “do” table options, so no explicit “do” match is required. Execute the statement and exit.
  - *Yes*: There were “do” table options, so the statement is executed only with an explicit match to one of them. Ignore the statement and exit.
- *Yes*: Loop.

Examples:

- No `--replicate-*` options at all

The slave executes all statements that it receives from the master.

- `--replicate-*-db` options, but no table options

The slave permits or ignores statements using the database options. Then it executes all statements permitted by those options because there are no table restrictions.

- `--replicate-*-table` options, but no database options

All statements are permitted at the database-checking stage because there are no database conditions. The slave executes or ignores statements based on the table options.

- A mix of database and table options

The slave permits or ignores statements using the database options. Then it evaluates all statements permitted by those options according to the table options. In some cases, this process can yield what might seem a counterintuitive result. Consider the following set of options:

```
[mysqld]
replicate-do-db      = db1
replicate-do-table  = db2.mytbl2
```

Suppose that `db1` is the default database and the slave receives this statement:

```
INSERT INTO mytbl1 VALUES(1,2,3);
```

The database is `db1`, which matches the `--replicate-do-db` option at the database-checking stage. The algorithm then proceeds to the table-checking stage. If there were no table options, the statement would be executed. However, because the options include a “do” table option, the statement must match if it is to be executed. The statement does not match, so it is ignored. (The same would happen for any table in `db1`.)

## 6.10. Replication FAQ

**Q:** How do I configure a slave if the master is running and I do not want to stop it?

**A:** There are several possibilities. If you have taken a snapshot backup of the master at some point and recorded the binary log filename and offset (from the output of `SHOW MASTER STATUS`) corresponding to the snapshot, use the following procedure:

1. Make sure that the slave is assigned a unique server ID.
2. Execute the following statement on the slave, filling in appropriate values for each option:

```
mysql> CHANGE MASTER TO
->     MASTER_HOST='master_host_name',
->     MASTER_USER='master_user_name',
->     MASTER_PASSWORD='master_pass',
->     MASTER_LOG_FILE='recorded_log_file_name',
->     MASTER_LOG_POS=recorded_log_position;
```

3. Execute `START SLAVE` on the slave.

If you do not have a backup of the master server, here is a quick procedure for creating one. All steps should be performed on the master host.

1. Issue this statement to acquire a global read lock:

```
mysql> FLUSH TABLES WITH READ LOCK;
```

2. With the lock still in place, execute this command (or a variation of it):

```
shell> tar zcf /tmp/backup.tar.gz /var/lib/mysql
```

3. Issue this statement and record the output, which you will need later:

```
mysql> SHOW MASTER STATUS;
```

4. Release the lock:

```
mysql> UNLOCK TABLES;
```

An alternative to using the preceding procedure to make a binary copy is to make an SQL dump of the master. To do this, you can use `mysqldump --master-data` on your master and later load the SQL dump into your slave. However, this is slower than making a binary copy.

Regardless of which of the two methods you use, afterward follow the instructions for the case when you have a snapshot and have recorded the log filename and offset. You can use the same snapshot to set up several slaves. Once you have the snapshot of the master, you can wait to set up a slave as long as the binary logs of the master are left intact. The two practical limitations on the length of time you can wait are the amount of disk space available to retain binary logs on the master and the length of time it takes the slave to catch up.

**Q:** Does the slave need to be connected to the master all the time?

**A:** No, it does not. The slave can go down or stay disconnected for hours or even days, and then reconnect and catch up on updates. For example, you can set up a master/slave relationship over a dial-up link



where the link is up only sporadically and for short periods of time. The implication of this is that, at any given time, the slave is not guaranteed to be in synchrony with the master unless you take some special measures.

**Q:** How do I know how late a slave is compared to the master? In other words, how do I know the date of the last statement replicated by the slave?

**A:** If the slave is 4.1.1 or newer, read the `Seconds_Behind_Master` column in `SHOW SLAVE STATUS`. For older versions, the following applies. This is possible only if `SHOW SLAVE STATUS` on the slave shows that the SQL thread is running (or for MySQL 3.23, that the slave thread is running), and that the thread has executed at least one event from the master. See [Section 6.3, “Replication Implementation Details”](#).

When the slave SQL thread executes an event read from the master, it modifies its own time to the event timestamp. (This is why `TIMESTAMP` is well replicated.) In the `Time` column in the output of `SHOW PROCESSLIST`, the number of seconds displayed for the slave SQL thread is the number of seconds between the timestamp of the last replicated event and the real time of the slave machine. You can use this to determine the date of the last replicated event. Note that if your slave has been disconnected from the master for one hour, and then reconnects, you may immediately see `Time` values like 3600 for the slave SQL thread in `SHOW PROCESSLIST`. This is because the slave is executing statements that are one hour old.

**Q:** How do I force the master to block updates until the slave catches up?

**A:** Use the following procedure:

1. On the master, execute these statements:

```
mysql> FLUSH TABLES WITH READ LOCK;  
mysql> SHOW MASTER STATUS;
```

Record the replication coordinates (the log filename and offset) from the output of the `SHOW` statement.

2. On the slave, issue the following statement, where the arguments to the `MASTER_POS_WAIT()` function are the replication coordinate values obtained in the previous step:

```
mysql> SELECT MASTER_POS_WAIT('log_name', log_offset);
```

The `SELECT` statement blocks until the slave reaches the specified log file and offset. At that point, the slave is in synchrony with the master and the statement returns.

3. On the master, issue the following statement to allow the master to begin processing updates again:

```
mysql> UNLOCK TABLES;
```

**Q:** What issues should I be aware of when setting up two-way replication?

**A:** MySQL replication currently does not support any locking protocol between master and slave to guarantee the atomicity of a distributed (cross-server) update. In other words, it is possible for client A to make an update to co-master 1, and in the meantime, before it propagates to co-master 2, client B could make an update to co-master 2 that makes the update of client A work differently than it did on co-master 1. Thus, when the update of client A makes it to co-master 2, it produces tables that are different from what you have on co-master 1, even after all the updates from co-master 2 have also propagated. This means that you should not chain two servers together in a two-way replication relationship unless you are sure that your updates can safely happen in any order, or unless you take care of mis-ordered updates somehow in the client code.

You should also realize that two-way replication actually does not improve performance very much (if at all) as far as updates are concerned. Each server must do the same number of updates, just as you would have a single server do. The only difference is that there is a little less lock contention, because the updates originating on another server are serialized in one slave thread. Even this benefit might be offset by network delays.

**Q:** How can I use replication to improve performance of my system?

**A:** You should set up one server as the master and direct all writes to it. Then configure as many slaves as you have the budget and rackspace for, and distribute the reads among the master and the slaves. You can also start the slaves with the `--skip-innodb`, `--skip-bdb`, `--low-priority-updates`, and `--delay-key-write=ALL` options to get speed improvements on the slave end. In this case, the slave uses non-transactional [MyISAM](#) tables instead of [InnoDB](#) and [BDB](#) tables to get more speed by eliminating transactional overhead.

**Q:** What should I do to prepare client code in my own applications to use performance-enhancing replication?

**A:** If the part of your code that is responsible for database access has been properly abstracted/modularized, converting it to run with a replicated setup should be very smooth and easy. Change the implementation of your database access to send all writes to the master, and to send reads to either the master or a slave. If your code does not have this level of abstraction, setting up a replicated system gives you the opportunity and motivation to it clean up. Start by creating a wrapper library or module that implements the following functions:

- `safe_writer_connect()`
- `safe_reader_connect()`
- `safe_reader_statement()`
- `safe_writer_statement()`

`safe_` in each function name means that the function takes care of handling all error conditions. You can use different names for the functions. The important thing is to have a unified interface for connecting for reads, connecting for writes, doing a read, and doing a write.

Then convert your client code to use the wrapper library. This may be a painful and scary process at first, but it pays off in the long run. All applications that use the approach just described are able to take advantage of a master/slave configuration, even one involving multiple slaves. The code is much easier to maintain, and adding troubleshooting options is trivial. You need modify only one or two functions; for example, to log how long each statement took, or which statement among those issued gave you an error.

If you have written a lot of code, you may want to automate the conversion task by using the [replace](#) utility that comes with standard MySQL distributions, or just write your own conversion script. Ideally, your code uses consistent programming style conventions. If not, then you are probably better off rewriting it anyway, or at least going through and manually regularizing it to use a consistent style.

**Q:** When and how much can MySQL replication improve the performance of my system?

**A:** MySQL replication is most beneficial for a system that processes frequent reads and infrequent writes. In theory, by using a single-master/multiple-slave setup, you can scale the system by adding more slaves until you either run out of network bandwidth, or your update load grows to the point that the master cannot handle it.

To determine how many slaves you can use before the added benefits begin to level out, and how much

you can improve performance of your site, you need to know your query patterns, and to determine empirically by benchmarking the relationship between the throughput for reads (reads per second, or `reads`) and for writes (`writes`) on a typical master and a typical slave. The example here shows a rather simplified calculation of what you can get with replication for a hypothetical system.

Let's say that system load consists of 10% writes and 90% reads, and we have determined by benchmarking that `reads` is  $1200 - 2 \times \text{writes}$ . In other words, the system can do 1,200 reads per second with no writes, the average write is twice as slow as the average read, and the relationship is linear. Let us suppose that the master and each slave have the same capacity, and that we have one master and  $N$  slaves. Then we have for each server (master or slave):

$$\text{reads} = 1200 - 2 \times \text{writes}$$

$$\text{reads} = 9 \times \text{writes} / (N + 1) \text{ (reads are split, but writes go to all servers)}$$

$$9 \times \text{writes} / (N + 1) + 2 \times \text{writes} = 1200$$

$$\text{writes} = 1200 / (2 + 9/(N+1))$$

The last equation indicates the maximum number of writes for  $N$  slaves, given a maximum possible read rate of 1,200 per minute and a ratio of nine reads per write.

This analysis yields the following conclusions:

- If  $N = 0$  (which means we have no replication), our system can handle about  $1200/11 = 109$  writes per second.
- If  $N = 1$ , we get up to 184 writes per second.
- If  $N = 8$ , we get up to 400 writes per second.
- If  $N = 17$ , we get up to 480 writes per second.
- Eventually, as  $N$  approaches infinity (and our budget negative infinity), we can get very close to 600 writes per second, increasing system throughput about 5.5 times. However, with only eight servers, we increase it nearly four times.


Note that these computations assume infinite network bandwidth and neglect several other factors that could be significant on your system. In many cases, you may not be able to perform a computation similar to the one just shown that accurately predicts what will happen on your system if you add  $N$  replication slaves. However, answering the following questions should help you decide whether and by how much replication will improve the performance of your system:

- What is the read/write ratio on your system?
- How much more write load can one server handle if you reduce the reads?
- For how many slaves do you have bandwidth available on your network?

**Q:** How can I use replication to provide redundancy or high availability?

**A:** With the currently available features, you would have to set up a master and a slave (or several slaves), and to write a script that monitors the master to check whether it is up. Then instruct your applications and the slaves to change master in case of failure. Some suggestions:

- To tell a slave to change its master, use the `CHANGE MASTER TO` statement.
- A good way to keep your applications informed as to the location of the master is by having a dynamic DNS entry for the master. With `bind` you can use `nsupdate` to dynamically update your DNS.
- Run your slaves with the `--log-bin` option and without `--log-slave-updates`. In this way, the slave is ready to become a master as soon as you issue `STOP SLAVE; RESET MASTER`, and `CHANGE MASTER TO` statement on the other slaves. For example, assume that you have the following setup:



In this diagram, `M` means the master, `S` the slaves, `WC` the clients issuing database writes and reads; clients that issue only database reads are not represented, because they need not switch. `S1`, `S2`, and `S3` are slaves running with `--log-bin` and without `--log-slave-updates`. Because updates received by a slave from the master are not logged in the binary log unless `--log-slave-updates` is specified, the binary log on each slave is empty initially. If for some reason `M` becomes unavailable, you can pick one of the slaves to become the new master. For example, if you pick `S1`, all `WC` should be redirected to `S1`, which will log updates to its binary log. `S2` and `S3` should then replicate from `S1`.


The reason for running the slave without `--log-slave-updates` is to prevent slaves from receiving updates twice in case you cause one of the slaves to become the new master. Suppose that `S1` has `--log-slave-updates` enabled. Then it will write updates that it receives from `M` to its own binary log. When `S2` changes from `M` to `S1` as its master, it may receive updates from `S1` that it has already received from `M`.

Make sure that all slaves have processed any statements in their relay log. On each slave, issue `STOP SLAVE IO_THREAD`, then check the output of `SHOW PROCESSLIST` until you see `Has read all relay log`. When this is true for all slaves, they can be reconfigured to the new setup. On the slave `S1` being promoted to become the master, issue `STOP SLAVE` and `RESET MASTER`.

On the other slaves `S2` and `S3`, use `STOP SLAVE` and `CHANGE MASTER TO MASTER_HOST='S1'` (where `'S1'` represents the real hostname of `S1`). To `CHANGE MASTER`, add all information about how to connect to `S1` from `S2` or `S3` (`user`, `password`, `port`). In `CHANGE MASTER`, there is no need to specify the name of `S1`'s binary log or binary log position to read from: We know it is the first binary log and position 4, which are the defaults for `CHANGE MASTER`. Finally, use `START SLAVE` on `S2` and `S3`.

Then instruct all `WC` to direct their statements to `S1`. From that point on, all updates statements sent by `WC` to `S1` are written to the binary log of `S1`, which then contains every update statement sent to `S1` since `M` died.

The result is this configuration:



+-----+

When **M** is up again, you must issue on it the same `CHANGE MASTER` as that issued on **S2** and **S3**, so that **M** becomes a slave of **S1** and picks up all the **WC** writes that it missed while it was down. To make **M** a master again (because it is the most powerful machine, for example), use the preceding procedure as if **S1** was unavailable and **M** was to be the new master. During this procedure, do not forget to run `RESET MASTER` on **M** before making **S1**, **S2**, and **S3** slaves of **M**. Otherwise, they may pick up old **WC** writes from before the point at which **M** became unavailable.

Note that there is no synchronization between the different slaves to a master. Some slaves might be ahead of others. This means that the concept outlined in the previous example might not work. In practice, however, the relay logs of different slaves will most likely not be far behind the master, so it would work, anyway (but there is no guarantee).

**Q:** How do I prevent `GRANT` and `REVOKE` statements from replicating to slave machines?

**A:** Start the server with the `--replicate-wild-ignore-table=mysql.%` option.

**Q:** Does replication work on mixed operating systems (for example, the master runs on Linux while slaves run on Mac OS X and Windows)?

**A:** Yes.

**Q:** Does replication work on mixed hardware architectures (for example, the master runs on a 64-bit machine while slaves run on 32-bit machines)?

**A:** Yes.

## 6.11. Troubleshooting Replication

If you have followed the instructions, and your replication setup is not working, the first thing to do is *check the error log for messages*. Many users have lost time by not doing this soon enough after encountering problems.

If you cannot tell from the error log what the problem was, try the following techniques:

- Verify that the master has binary logging enabled by issuing a `SHOW MASTER STATUS` statement. If logging is enabled, `Position` is non-zero. If binary logging is not enabled, verify that you are running the master with the `--log-bin` and `--server-id` options.
- Verify that the slave is running. Use `SHOW SLAVE STATUS` to check whether the `Slave_IO_Running` and `Slave_SQL_Running` values are both `Yes`. If not, verify the options that were used when starting the slave server. For example, `--skip-slave-start` prevents the slave threads from starting until you issue a `START SLAVE` statement.
- If the slave is running, check whether it established a connection to the master. Use `SHOW PROCESSLIST`, find the I/O and SQL threads and check their `State` column to see what they display. See [Section 6.3, “Replication Implementation Details”](#). If the I/O thread state says `Connecting to master`, verify the privileges for the replication user on the master, the master hostname, your DNS setup, whether the master is actually running, and whether it is reachable from the slave.
- If the slave was running previously but has stopped, the reason usually is that some statement that succeeded on the master failed on the slave. This should never happen if you have taken a proper snapshot of the master, and never modified the data on the slave outside of the slave thread. If the slave stops unexpectedly, it is a bug or you have encountered one of the known replication limita-

tions described in [Section 6.7, “Replication Features and Known Problems”](#). If it is a bug, see [Section 6.12, “How to Report Replication Bugs or Problems”](#), for instructions on how to report it.

- If a statement that succeeded on the master refuses to run on the slave, try the following procedure if it is not feasible to do a full database resynchronization by deleting the slave's databases and copying a new snapshot from the master:

1. Determine whether the affected table on the slave is different from the master table. Try to understand how this happened. Then make the slave's table identical to the master's and run `START SLAVE`.
2. If the preceding step does not work or does not apply, try to understand whether it would be safe to make the update manually (if needed) and then ignore the next statement from the master.
3. If you decide that you can skip the next statement from the master, issue the following statements:

```
mysql> SET GLOBAL SQL_SLAVE_SKIP_COUNTER = N;  
mysql> START SLAVE;
```

The value of *N* should be 1 if the next statement from the master does not use `AUTO_INCREMENT` or `LAST_INSERT_ID()`. Otherwise, the value should be 2. The reason for using a value of 2 for statements that use `AUTO_INCREMENT` or `LAST_INSERT_ID()` is that they take two events in the binary log of the master.

4. If you are sure that the slave started out perfectly synchronized with the master, and that no one has updated the tables involved outside of the slave thread, then presumably the discrepancy is the result of a bug. If you are running the most recent version of MySQL, please report the problem. If you are running an older version, try upgrading to the latest production release to determine whether the problem persists.

## 6.12. How to Report Replication Bugs or Problems

When you have determined that there is no user error involved, and replication still either does not work at all or is unstable, it is time to send us a bug report. We need to obtain as much information as possible from you to be able to track down the bug. Please spend some time and effort in preparing a good bug report.

If you have a repeatable test case that demonstrates the bug, please enter it into our bugs database using the instructions given in [Section 1.8, “How to Report Bugs or Problems”](#). If you have a “phantom” problem (one that you cannot duplicate at will), use the following procedure:

1. Verify that no user error is involved. For example, if you update the slave outside of the slave thread, the data goes out of synchrony, and you can have unique key violations on updates. In this case, the slave thread stops and waits for you to clean up the tables manually to bring them into synchrony. *This is not a replication problem. It is a problem of outside interference causing replication to fail.*
2. Run the slave with the `--log-slave-updates` and `--log-bin` options. These options cause the slave to log the updates that it receives from the master into its own binary logs.
3. Save all evidence before resetting the replication state. If we have no information or only sketchy information, it becomes difficult or impossible for us to track down the problem. The evidence you should collect is:

- All binary logs from the master
  - All binary logs from the slave
  - The output of `SHOW MASTER STATUS` from the master at the time you discovered the problem
  - The output of `SHOW SLAVE STATUS` from the slave at the time you discovered the problem
  - Error logs from the master and the slave
4. Use `mysqlbinlog` to examine the binary logs. The following should be helpful to find the problem statement. `log_pos` and `log_file` are the `Master_Log_File` and `Read_Master_Log_Pos` values from `SHOW SLAVE STATUS`.

```
shell> mysqlbinlog -j log_pos log_file | head
```

After you have collected the evidence for the problem, try to isolate it as a separate test case first. Then enter the problem with as much information as possible into our bugs database using the instructions at [Section 1.8, “How to Report Bugs or Problems”](#).

---

# Chapter 7. Optimization

Optimization is a complex task because ultimately it requires understanding of the entire system to be optimized. Although it may be possible to perform some local optimizations with little knowledge of your system or application, the more optimal you want your system to become, the more you must know about it.

This chapter tries to explain and give some examples of different ways to optimize MySQL. Remember, however, that there are always additional ways to make the system even faster, although they may require increasing effort to achieve.

## 7.1. Optimization Overview

The most important factor in making a system fast is its basic design. You must also know what kinds of processing your system is doing, and what its bottlenecks are. In most cases, system bottlenecks arise from these sources:

- Disk seeks. It takes time for the disk to find a piece of data. With modern disks, the mean time for this is usually lower than 10ms, so we can in theory do about 100 seeks a second. This time improves slowly with new disks and is very hard to optimize for a single table. The way to optimize seek time is to distribute the data onto more than one disk.
- Disk reading and writing. When the disk is at the correct position, we need to read the data. With modern disks, one disk delivers at least 10–20MB/s throughput. This is easier to optimize than seeks because you can read in parallel from multiple disks.
- CPU cycles. When we have the data in main memory, we need to process it to get our result. Having small tables compared to the amount of memory is the most common limiting factor. But with small tables, speed is usually not the problem.
- Memory bandwidth. When the CPU needs more data than can fit in the CPU cache, main memory bandwidth becomes a bottleneck. This is an uncommon bottleneck for most systems, but one to be aware of.

### 7.1.1. MySQL Design Limitations and Tradeoffs

When using the [MyISAM](#) storage engine, MySQL uses extremely fast table locking that allows multiple readers or a single writer. The biggest problem with this storage engine occurs when you have a steady stream of mixed updates and slow selects on a single table. If this is a problem for certain tables, you can use another storage engine for them. See [Chapter 14, Storage Engines and Table Types](#).

MySQL can work with both transactional and non-transactional tables. To make it easier to work smoothly with non-transactional tables (which cannot roll back if something goes wrong), MySQL has the following rules. Note that these rules apply *only* when you use the [IGNORE](#) specifier for [INSERT](#) or [UPDATE](#).

- All columns have default values.
- If you insert an inappropriate or out-of-range value into a column, MySQL sets the column to the “best possible value” instead of reporting an error. For numerical values, this is 0, the smallest possible value or the largest possible value. For strings, this is either the empty string or as much of the string as can be stored in the column.



- All calculated expressions return a value that can be used instead of signaling an error condition. For example, 1/0 returns `NULL`.

To change the preceding behaviors, you can enable stricter data handling by setting the server SQL mode appropriately. For more information about data handling, see [Section 1.9.6, “How MySQL Deals with Constraints”](#), [Section 5.2.5, “SQL Modes”](#), and [Section 13.2.4, “INSERT Syntax”](#).

## 7.1.2. Designing Applications for Portability

Because all SQL servers implement different parts of standard SQL, it takes work to write portable database applications. It is very easy to achieve portability for very simple selects and inserts, but becomes more difficult the more capabilities you require. If you want an application that is fast with many database systems, it becomes even more difficult.

All database systems have some weak points. That is, they have different design compromises that lead to different behavior.

To make a complex application portable, you need to determine which SQL servers it must work with, and then determine what features those servers support. You can use the MySQL `crash-me` program to find functions, types, and limits that you can use with a selection of database servers. `crash-me` does not check for every possible feature, but it is still reasonably comprehensive, performing about 450 tests. An example of the type of information `crash-me` can provide is that you should not use column names that are longer than 18 characters if you want to be able to use Informix or DB2.

The `crash-me` program and the MySQL benchmarks are all very database independent. By taking a look at how they are written, you can get a feeling for what you must do to make your own applications database independent. The programs can be found in the `sql-bench` directory of MySQL source distributions. They are written in Perl and use the DBI database interface. Use of DBI in itself solves part of the portability problem because it provides database-independent access methods. See [Section 7.1.4, “The MySQL Benchmark Suite”](#).

If you strive for database independence, you need to get a good feeling for each SQL server's bottlenecks. For example, MySQL is very fast in retrieving and updating rows for `MyISAM` tables, but has a problem in mixing slow readers and writers on the same table. Oracle, on the other hand, has a big problem when you try to access rows that you have recently updated (until they are flushed to disk). Transactional database systems in general are not very good at generating summary tables from log tables, because in this case row locking is almost useless.

To make your application *really* database independent, you should define an easily extendable interface through which you manipulate your data. For example, C++ is available on most systems, so it makes sense to use a C++ class-based interface to the databases.

If you use some feature that is specific to a given database system (such as the `REPLACE` statement, which is specific to MySQL), you should implement the same feature for other SQL servers by coding an alternative method. Although the alternative might be slower, it enables the other servers to perform the same tasks.

With MySQL, you can use the `/*! */` syntax to add MySQL-specific keywords to a statement. The code inside `/* */` is treated as a comment (and ignored) by most other SQL servers. For information about writing comments, see [Section 9.4, “Comment Syntax”](#).

If high performance is more important than exactness, as for some Web applications, it is possible to create an application layer that caches all results to give you even higher performance. By letting old results expire after a while, you can keep the cache reasonably fresh. This provides a method to handle high load spikes, in which case you can dynamically increase the cache size and set the expiration timeout higher until things get back to normal.

In this case, the table creation information should contain information about the initial cache size and how often the table should normally be refreshed.

An attractive alternative to implementing an application cache is to use the MySQL query cache. By enabling the query cache, the server handles the details of determining whether a query result can be re-used. This simplifies your application. See [Section 5.13, “The MySQL Query Cache”](#).

### 7.1.3. What We Have Used MySQL For

This section describes an early application for MySQL.

During MySQL initial development, the features of MySQL were made to fit our largest customer, which handled data warehousing for a couple of the largest retailers in Sweden.

From all stores, we got weekly summaries of all bonus card transactions, and were expected to provide useful information for the store owners to help them find how their advertising campaigns were affecting their own customers.

The volume of data was quite huge (about seven million summary transactions per month), and we had data for 4–10 years that we needed to present to the users. We got weekly requests from our customers, who wanted instant access to new reports from this data.

We solved this problem by storing all information per month in compressed “transaction tables.” We had a set of simple macros that generated summary tables grouped by different criteria (product group, customer id, store, and so on) from the tables in which the transactions were stored. The reports were Web pages that were dynamically generated by a small Perl script. This script parsed a Web page, executed the SQL statements in it, and inserted the results. We would have used PHP or `mod_perl` instead, but they were not available at the time.

For graphical data, we wrote a simple tool in C that could process SQL query results and produce GIF images based on those results. This tool also was dynamically executed from the Perl script that parses the Web pages.

In most cases, a new report could be created simply by copying an existing script and modifying the SQL query that it used. In some cases, we needed to add more columns to an existing summary table or generate a new one. This also was quite simple because we kept all transaction-storage tables on disk. (This amounted to about 50GB of transaction tables and 200GB of other customer data.)

We also let our customers access the summary tables directly with ODBC so that the advanced users could experiment with the data themselves.

This system worked well and we had no problems handling the data with quite modest Sun Ultra SPARCstation hardware (2×200MHz). Eventually the system was migrated to Linux.

### 7.1.4. The MySQL Benchmark Suite

This benchmark suite is meant to tell any user what operations a given SQL implementation performs well or poorly. You can get a good idea for how the benchmarks work by looking at the code and results in the `sql-bench` directory in any MySQL source distribution.

Note that this benchmark is single-threaded, so it measures the minimum time for the operations performed. We plan to add multi-threaded tests to the benchmark suite in the future.

To use the benchmark suite, the following requirements must be satisfied:

- The benchmark suite is provided with MySQL source distributions. You can either download a released distribution from <http://dev.mysql.com/downloads/>, or use the current development source

tree. (See [Section 2.9.3, “Installing from the Development Source Tree”](#).)

- The benchmark scripts are written in Perl and use the Perl DBI module to access database servers, so DBI must be installed. You also need the server-specific DBD drivers for each of the servers you want to test. For example, to test MySQL, PostgreSQL, and DB2, you must have the `DBD::mysql`, `DBD::Pg`, and `DBD::DB2` modules installed. See [Section 2.14, “Perl Installation Notes”](#).

After you obtain a MySQL source distribution, you can find the benchmark suite located in its `sql-bench` directory. To run the benchmark tests, build MySQL, and then change location into the `sql-bench` directory and execute the `run-all-tests` script:

```
shell> cd sql-bench
shell> perl run-all-tests --server=server_name
```

`server_name` should be the name of one of the supported servers. To get a list of all options and supported servers, invoke this command:

```
shell> perl run-all-tests --help
```

The `crash-me` script also is located in the `sql-bench` directory. `crash-me` tries to determine what features a database system supports and what its capabilities and limitations are by actually running queries. For example, it determines:

- What data types are supported
- How many indexes are supported
- What functions are supported
- How big a query can be
- How big a `VARCHAR` column can be

You can find the results from `crash-me` for many different database servers at <http://dev.mysql.com/tech-resources/crash-me.php>. For more information about benchmark results, visit <http://dev.mysql.com/tech-resources/benchmarks/>.

## 7.1.5. Using Your Own Benchmarks

You should definitely benchmark your application and database to find out where the bottlenecks are. After fixing one bottleneck (or by replacing it with a “dummy” module), you can proceed to identify the next bottleneck. Even if the overall performance for your application currently is acceptable, you should at least make a plan for each bottleneck and decide how to solve it if someday you really need the extra performance.

For examples of portable benchmark programs, look at those in the MySQL benchmark suite. See [Section 7.1.4, “The MySQL Benchmark Suite”](#). You can take any program from this suite and modify it for your own needs. By doing this, you can try different solutions to your problem and test which really is fastest for you.

Another free benchmark suite is the Open Source Database Benchmark, available at <http://osdb.sourceforge.net/>.

It is very common for a problem to occur only when the system is very heavily loaded. We have had many customers who contact us when they have a (tested) system in production and have encountered

load problems. In most cases, performance problems turn out to be due to issues of basic database design (for example, table scans are not good under high load) or problems with the operating system or libraries. Most of the time, these problems would be much easier to fix if the systems were not already in production.

To avoid problems like this, you should put some effort into benchmarking your whole application under the worst possible load. You can use Super Smack, available at <http://jeremy.zawodny.com/mysql/super-smack/>. As suggested by its name, it can bring a system to its knees, so make sure to use it only on your development systems.

## 7.2. Optimizing **SELECT** and Other Statements

First, one factor affects all statements: The more complex your permissions setup, the more overhead you have. Using simpler permissions when you issue **GRANT** statements enables MySQL to reduce permission-checking overhead when clients execute statements. For example, if you do not grant any table-level or column-level privileges, the server need not ever check the contents of the `tables_priv` and `columns_priv` tables. Similarly, if you place no resource limits on any accounts, the server does not have to perform resource counting. If you have a very high statement-processing load, it may be worth the time to use a simplified grant structure to reduce permission-checking overhead.

If your problem is with a specific MySQL expression or function, you can perform a timing test by invoking the `BENCHMARK()` function using the `mysql` client program. Its syntax is `BENCHMARK(loop_count, expression)`. The return value is always zero, but `mysql` prints a line displaying approximately how long the statement took to execute. For example:

```
mysql> SELECT BENCHMARK(1000000,1+1);
+-----+
| BENCHMARK(1000000,1+1) |
+-----+
|                          0 |
+-----+
1 row in set (0.32 sec)
```

This result was obtained on a Pentium II 400MHz system. It shows that MySQL can execute 1,000,000 simple addition expressions in 0.32 seconds on that system.

All MySQL functions should be highly optimized, but there may be some exceptions. `BENCHMARK()` is an excellent tool for finding out if some function is a problem for your queries.

### 7.2.1. Optimizing Queries with **EXPLAIN**

```
EXPLAIN tbl_name
```

Or:

```
EXPLAIN [EXTENDED] SELECT select_options
```

The **EXPLAIN** statement can be used either as a synonym for **DESCRIBE** or as a way to obtain information about how MySQL executes a **SELECT** statement:

- `EXPLAIN tbl_name` is synonymous with `DESCRIBE tbl_name` or `SHOW COLUMNS FROM tbl_name`.
- When you precede a **SELECT** statement with the keyword **EXPLAIN**, MySQL displays information from the optimizer about the query execution plan. That is, MySQL explains how it would process the **SELECT**, including information about how tables are joined and in which order.

This section describes the second use of `EXPLAIN` for obtaining query execution plan information. For a description of the `DESCRIBE` and `SHOW COLUMNS` statements, see [Section 13.3.1, “DESCRIBE Syntax”](#), and [Section 13.5.4.3, “SHOW COLUMNS Syntax”](#).

With the help of `EXPLAIN`, you can see where you should add indexes to tables to get a faster `SELECT` that uses indexes to find rows. You can also use `EXPLAIN` to check whether the optimizer joins the tables in an optimal order. To force the optimizer to use a join order corresponding to the order in which the tables are named in the `SELECT` statement, begin the statement with `SELECT STRAIGHT_JOIN` rather than just `SELECT`.

If you have a problem with indexes not being used when you believe that they should be, you should run `ANALYZE TABLE` to update table statistics such as cardinality of keys, that can affect the choices the optimizer makes. See [Section 13.5.2.1, “ANALYZE TABLE Syntax”](#).

`EXPLAIN` returns a row of information for each table used in the `SELECT` statement. The tables are listed in the output in the order that MySQL would read them while processing the query. MySQL resolves all joins using a *single-sweep multi-join* method. This means that MySQL reads a row from the first table, and then finds a matching row in the second table, the third table, and so on. When all tables are processed, MySQL outputs the selected columns and backtracks through the table list until a table is found for which there are more matching rows. The next row is read from this table and the process continues with the next table.

In MySQL version 4.1, the `EXPLAIN` output format was changed to work better with constructs such as `UNION` statements, subqueries, and derived tables. Most notable is the addition of two new columns: `id` and `select_type`. You do not see these columns when using servers older than MySQL 4.1. `EXPLAIN` syntax also was augmented to allow the `EXTENDED` keyword. When this keyword is used, `EXPLAIN` produces extra information that can be viewed by issuing a `SHOW WARNINGS` statement following the `EXPLAIN` statement. This information displays how the optimizer qualifies table and column names in the `SELECT` statement, what the `SELECT` looks like after the application of rewriting and optimization rules, and possibly other notes about the optimization process.

Each output row from `EXPLAIN` provides information about one table, and each row contains the following columns:

- `id`

The `SELECT` identifier. This is the sequential number of the `SELECT` within the query.

- `select_type`

The type of `SELECT`, which can be any of those shown in the following table:

<code>SIMPLE</code>	Simple <code>SELECT</code> (not using <code>UNION</code> or subqueries)
<code>PRIMARY</code>	Outermost <code>SELECT</code>
<code>UNION</code>	Second or later <code>SELECT</code> statement in a <code>UNION</code>
<code>DEPENDENT UNION</code>	Second or later <code>SELECT</code> statement in a <code>UNION</code> , dependent on outer query
<code>UNION RESULT</code>	Result of a <code>UNION</code> .
<code>SUBQUERY</code>	First <code>SELECT</code> in subquery
<code>DEPENDENT SUB-QUERY</code>	First <code>SELECT</code> in subquery, dependent on outer query
<code>DERIVED</code>	Derived table <code>SELECT</code> (subquery in <code>FROM</code> clause)
<code>UNCACHEABLE SUB-QUERY</code>	A subquery for which the result cannot be cached and must be re-evaluated for each row of the outer query

`DEPENDENT` typically signifies the use of a correlated subquery. See [Section 13.2.8.7, “Correlated Subqueries”](#).

“DEPENDENT SUBQUERY” evaluation differs from `UNCACHEABLE SUBQUERY` evaluation. For “DEPENDENT SUBQUERY”, the subquery is re-evaluated only once for each set of different values of the variables from its outer context. For `UNCACHEABLE SUBQUERY`, the subquery is re-evaluated for each row of the outer context. Cacheability of subqueries is subject to the restrictions detailed in [Section 5.13.1, “How the Query Cache Operates”](#). For example, referring to user variables makes a subquery uncacheable.

- `table`

The table to which the row of output refers.

- `type`

The join type. The different join types are listed here, ordered from the best type to the worst:

- `system`

The table has only one row (= system table). This is a special case of the `const` join type.

- `const`

The table has at most one matching row, which is read at the start of the query. Because there is only one row, values from the column in this row can be regarded as constants by the rest of the optimizer. `const` tables are very fast because they are read only once.

`const` is used when you compare all parts of a `PRIMARY KEY` or `UNIQUE` index to constant values. In the following queries, `tbl_name` can be used as a `const` table:

```
SELECT * FROM tbl_name WHERE primary_key=1;

SELECT * FROM tbl_name
WHERE primary_key_part1=1 AND primary_key_part2=2;
```

- `eq_ref`

One row is read from this table for each combination of rows from the previous tables. Other than the `system` and `const` types, this is the best possible join type. It is used when all parts of an index are used by the join and the index is a `PRIMARY KEY` or `UNIQUE` index.

`eq_ref` can be used for indexed columns that are compared using the `=` operator. The comparison value can be a constant or an expression that uses columns from tables that are read before this table. In the following examples, MySQL can use an `eq_ref` join to process `ref_table`:

```
SELECT * FROM ref_table,other_table
WHERE ref_table.key_column=other_table.column;

SELECT * FROM ref_table,other_table
WHERE ref_table.key_column_part1=other_table.column
AND ref_table.key_column_part2=1;
```

- `ref`

All rows with matching index values are read from this table for each combination of rows from the previous tables. `ref` is used if the join uses only a leftmost prefix of the key or if the key is not a `PRIMARY KEY` or `UNIQUE` index (in other words, if the join cannot select a single row based on the key value). If the key that is used matches only a few rows, this is a good join type.

`ref` can be used for indexed columns that are compared using the `=` or `<=>` operator. In the following examples, MySQL can use a `ref` join to process `ref_table`:

```
SELECT * FROM ref_table WHERE key_column=expr;

SELECT * FROM ref_table,other_table
  WHERE ref_table.key_column=other_table.column;

SELECT * FROM ref_table,other_table
  WHERE ref_table.key_column_part1=other_table.column
  AND ref_table.key_column_part2=1;
```

- `ref_or_null`

This join type is like `ref`, but with the addition that MySQL does an extra search for rows that contain `NULL` values. This join type optimization was added for MySQL 4.1.1 and is used mostly when resolving subqueries. In the following examples, MySQL can use a `ref_or_null` join to process `ref_table`:

```
SELECT * FROM ref_table
  WHERE key_column=expr OR key_column IS NULL;
```

See [Section 7.2.6, “IS NULL Optimization”](#).

- `unique_subquery`

This type replaces `ref` for some `IN` subqueries of the following form:

```
value IN (SELECT primary_key FROM single_table WHERE some_expr)
```

`unique_subquery` is just an index lookup function that replaces the subquery completely for better efficiency.

- `index_subquery`

This join type is similar to `unique_subquery`. It replaces `IN` subqueries, but it works for non-unique indexes in subqueries of the following form:

```
value IN (SELECT key_column FROM single_table WHERE some_expr)
```

- `range`

Only rows that are in a given range are retrieved, using an index to select the rows. The `key` column in the output row indicates which index is used. The `key_len` contains the longest key part that was used. The `ref` column is `NULL` for this type.

`range` can be used when a key column is compared to a constant using any of the `=`, `<>`, `>`, `>=`, `<`, `<=`, `IS NULL`, `<=>`, `BETWEEN`, or `IN` operators:

```
SELECT * FROM tbl_name
  WHERE key_column = 10;

SELECT * FROM tbl_name
  WHERE key_column BETWEEN 10 and 20;

SELECT * FROM tbl_name
  WHERE key_column IN (10,20,30);

SELECT * FROM tbl_name
  WHERE key_part1= 10 AND key_part2 IN (10,20,30);
```

- `index`

This join type is the same as `ALL`, except that only the index tree is scanned. This usually is faster than `ALL` because the index file usually is smaller than the data file.

MySQL can use this join type when the query uses only columns that are part of a single index.

- `ALL`

A full table scan is done for each combination of rows from the previous tables. This is normally not good if the table is the first table not marked `const`, and usually *very* bad in all other cases. Normally, you can avoid `ALL` by adding indexes that allow row retrieval from the table based on constant values or column values from earlier tables.

- `possible_keys`

The `possible_keys` column indicates which indexes MySQL can choose from use to find the rows in this table. Note that this column is totally independent of the order of the tables as displayed in the output from `EXPLAIN`. That means that some of the keys in `possible_keys` might not be usable in practice with the generated table order.

If this column is `NULL`, there are no relevant indexes. In this case, you may be able to improve the performance of your query by examining the `WHERE` clause to check whether it refers to some column or columns that would be suitable for indexing. If so, create an appropriate index and check the query with `EXPLAIN` again. See [Section 13.1.2, “ALTER TABLE Syntax”](#).

To see what indexes a table has, use `SHOW INDEX FROM tbl_name`.

- `key`

The `key` column indicates the key (index) that MySQL actually decided to use. The key is `NULL` if no index was chosen. To force MySQL to use or ignore an index listed in the `possible_keys` column, use `FORCE INDEX`, `USE INDEX`, or `IGNORE INDEX` in your query. See [Section 13.2.7, “SELECT Syntax”](#).

For `MyISAM` and `BDB` tables, running `ANALYZE TABLE` helps the optimizer choose better indexes. For `MyISAM` tables, `myisamchk --analyze` does the same. See [Section 13.5.2.1, “ANALYZE TABLE Syntax”](#), and [Section 5.9.4, “Table Maintenance and Crash Recovery”](#).

- `key_len`

The `key_len` column indicates the length of the key that MySQL decided to use. The length is `NULL` if the `key` column says `NULL`. Note that the value of `key_len` enables you to determine how many parts of a multiple-part key MySQL actually uses.

- `ref`

The `ref` column shows which columns or constants are compared to the index named in the `key` column to select rows from the table.

- `rows`

The `rows` column indicates the number of rows MySQL believes it must examine to execute the query.

- `Extra`

This column contains additional information about how MySQL resolves the query. Here is an ex-



planation of the values that can appear in this column:

- [Distinct](#)

MySQL is looking for distinct values, so it stops searching for more rows for the current row combination after it has found the first matching row.

- [Not exists](#)

MySQL was able to do a [LEFT JOIN](#) optimization on the query and does not examine more rows in this table for the previous row combination after it finds one row that matches the [LEFT JOIN](#) criteria. Here is an example of the type of query that can be optimized this way:

```
SELECT * FROM t1 LEFT JOIN t2 ON t1.id=t2.id
WHERE t2.id IS NULL;
```

Assume that `t2.id` is defined as [NOT NULL](#). In this case, MySQL scans `t1` and looks up the rows in `t2` using the values of `t1.id`. If MySQL finds a matching row in `t2`, it knows that `t2.id` can never be [NULL](#), and does not scan through the rest of the rows in `t2` that have the same `id` value. In other words, for each row in `t1`, MySQL needs to do only a single lookup in `t2`, regardless of how many rows actually match in `t2`.

- [range checked for each record \(index map: N\)](#)

MySQL found no good index to use, but found that some of indexes might be used after column values from preceding tables are known. For each row combination in the preceding tables, MySQL checks whether it is possible to use a [range](#) access method to retrieve rows. The applicability criteria are as described in [Section 7.2.5, “Range Optimization”](#), with the exception that all column values for the preceding table are known and considered to be constants.

This is not very fast, but is faster than performing a join with no index at all.

- [Using filesort](#)

MySQL must do an extra pass to find out how to retrieve the rows in sorted order. The sort is done by going through all rows according to the join type and storing the sort key and pointer to the row for all rows that match the [WHERE](#) clause. The keys then are sorted and the rows are retrieved in sorted order. See [Section 7.2.9, “ORDER BY Optimization”](#).

- [Using index](#)

The column information is retrieved from the table using only information in the index tree without having to do an additional seek to read the actual row. This strategy can be used when the query uses only columns that are part of a single index.

- [Using temporary](#)

To resolve the query, MySQL needs to create a temporary table to hold the result. This typically happens if the query contains [GROUP BY](#) and [ORDER BY](#) clauses that list columns differently.

- [Using where](#)

A [WHERE](#) clause is used to restrict which rows to match against the next table or send to the client. Unless you specifically intend to fetch or examine all rows from the table, you may have something wrong in your query if the [Extra](#) value is not [Using where](#) and the table join type is [ALL](#) or [index](#).

If you want to make your queries as fast as possible, you should look out for [Extra](#) values of

Using `filesort` and `Using temporary`.

- Using `index` for `group-by`

Similar to the `Using index` way of accessing a table, `Using index for group-by` indicates that MySQL found an index that can be used to retrieve all columns of a `GROUP BY` or `DISTINCT` query without any extra disk access to the actual table. Additionally, the index is used in the most efficient way so that for each group, only a few index entries are read. For details, see [Section 7.2.10, “GROUP BY Optimization”](#).

You can get a good indication of how good a join is by taking the product of the values in the `rows` column of the `EXPLAIN` output. This should tell you roughly how many rows MySQL must examine to execute the query. If you restrict queries with the `max_join_size` system variable, this row product also is used to determine which multiple-table `SELECT` statements to execute and which to abort. See [Section 7.5.2, “Tuning Server Parameters”](#).

The following example shows how a multiple-table join can be optimized progressively based on the information provided by `EXPLAIN`.

Suppose that you have the `SELECT` statement shown here and that you plan to examine it using `EXPLAIN`:

```
EXPLAIN SELECT tt.TicketNumber, tt.TimeIn,
              tt.ProjectReference, tt.EstimatedShipDate,
              tt.ActualShipDate, tt.ClientID,
              tt.ServiceCodes, tt.RepetitiveID,
              tt.CurrentProcess, tt.CurrentDPPerson,
              tt.RecordVolume, tt.DPPrinted, et.COUNTRY,
              et_1.COUNTRY, do.CUSTNAME
FROM tt, et, et AS et_1, do
WHERE tt.SubmitTime IS NULL
      AND tt.ActualPC = et.EMPLOYID
      AND tt.AssignedPC = et_1.EMPLOYID
      AND tt.ClientID = do.CUSTNMBR;
```

For this example, make the following assumptions:

- The columns being compared have been declared as follows:

Table	Column	Data Type
tt	ActualPC	CHAR(10)
tt	AssignedPC	CHAR(10)
tt	ClientID	CHAR(10)
et	EMPLOYID	CHAR(15)
do	CUSTNMBR	CHAR(15)

- The tables have the following indexes:

Table	Index
tt	ActualPC
tt	AssignedPC
tt	ClientID
et	EMPLOYID (primary key)

do	CUSTNMBR (primary key)
----	------------------------

- The `tt.ActualPC` values are not evenly distributed.

Initially, before any optimizations have been performed, the `EXPLAIN` statement produces the following information:

```

table type possible_keys key key_len ref rows Extra
et ALL PRIMARY NULL NULL NULL 74
do ALL PRIMARY NULL NULL NULL 2135
et_1 ALL PRIMARY NULL NULL NULL 74
tt ALL AssignedPC, NULL NULL NULL 3872
      ClientID,
      ActualPC
      range checked for each record (key map: 35)

```

Because `type` is `ALL` for each table, this output indicates that MySQL is generating a Cartesian product of all the tables; that is, every combination of rows. This takes quite a long time, because the product of the number of rows in each table must be examined. For the case at hand, this product is  $74 \times 2135 \times 74 \times 3872 = 45,268,558,720$  rows. If the tables were bigger, you can only imagine how long it would take.

One problem here is that MySQL can use indexes on columns more efficiently if they are declared as the same type and size. (For `ISAM` tables, indexes may not be used at all unless the columns are declared the same.) In this context, `VARCHAR` and `CHAR` are considered the same if they are declared as the same size. `tt.ActualPC` is declared as `CHAR(10)` and `et.EMPLOYID` is `CHAR(15)`, so there is a length mismatch.

To fix this disparity between column lengths, use `ALTER TABLE` to lengthen `ActualPC` from 10 characters to 15 characters:

```
mysql> ALTER TABLE tt MODIFY ActualPC VARCHAR(15);
```

Now `tt.ActualPC` and `et.EMPLOYID` are both `VARCHAR(15)`. Executing the `EXPLAIN` statement again produces this result:

```

table type possible_keys key key_len ref rows Extra
tt ALL AssignedPC, NULL NULL NULL 3872 Using
      ClientID, where
      ActualPC
do ALL PRIMARY NULL NULL NULL 2135
  range checked for each record (key map: 1)
et_1 ALL PRIMARY NULL NULL NULL 74
  range checked for each record (key map: 1)
et eq_ref PRIMARY PRIMARY 15 tt.ActualPC 1

```

This is not perfect, but is much better: The product of the `rows` values is less by a factor of 74. This version executes in a couple of seconds.

A second alteration can be made to eliminate the column length mismatches for the `tt.AssignedPC = et_1.EMPLOYID` and `tt.ClientID = do.CUSTNMBR` comparisons:

```
mysql> ALTER TABLE tt MODIFY AssignedPC VARCHAR(15),
->      MODIFY ClientID VARCHAR(15);
```

After that modification, `EXPLAIN` produces the output shown here:

```

table type possible_keys key key_len ref rows Extra
et ALL PRIMARY NULL NULL NULL 74
tt ref AssignedPC, ActualPC 15 et.EMPLOYID 52 Using

```

		ClientID, ActualPC				where
et_1	eq_ref	PRIMARY	PRIMARY	15	tt.AssignedPC	1
do	eq_ref	PRIMARY	PRIMARY	15	tt.ClientID	1

At this point, the query is optimized almost as well as possible. The remaining problem is that, by default, MySQL assumes that values in the `tt.ActualPC` column are evenly distributed, and that is not the case for the `tt` table. Fortunately, it is easy to tell MySQL to analyze the key distribution:

```
mysql> ANALYZE TABLE tt;
```

With the additional index information, the join is perfect and `EXPLAIN` produces this result:

table	type	possible_keys	key	key_len	ref	rows	Extra
tt	ALL	AssignedPC ClientID, ActualPC	NULL	NULL	NULL	3872	Using where
et	eq_ref	PRIMARY	PRIMARY	15	tt.ActualPC	1	
et_1	eq_ref	PRIMARY	PRIMARY	15	tt.AssignedPC	1	
do	eq_ref	PRIMARY	PRIMARY	15	tt.ClientID	1	

Note that the `rows` column in the output from `EXPLAIN` is an educated guess from the MySQL join optimizer. You should check whether the numbers are even close to the truth by comparing the `rows` product with the actual number of rows that the query returns. If the numbers are quite different, you might get better performance by using `STRAIGHT_JOIN` in your `SELECT` statement and trying to list the tables in a different order in the `FROM` clause.

## 7.2.2. Estimating Query Performance

In most cases, you can estimate query performance by counting disk seeks. For small tables, you can usually find a row in one disk seek (because the index is probably cached). For bigger tables, you can estimate that, using B-tree indexes, you need this many seeks to find a row:  $\log(\text{row\_count}) / \log(\text{index\_block\_length} / 3 \times 2 / (\text{index\_length} + \text{data\_pointer\_length})) + 1$ .

In MySQL, an index block is usually 1,024 bytes and the data pointer is usually four bytes. For a 500,000-row table with an index length of three bytes (the size of `MEDIUMINT`), the formula indicates  $\log(500,000) / \log(1024 / 3 \times 2 / (3 + 4)) + 1 = 4$  seeks.

This index would require storage of about  $500,000 \times 7 \times 3/2 = 5.2\text{MB}$  (assuming a typical index buffer fill ratio of 2/3), so you probably have much of the index in memory and so need only one or two calls to read data to find the row.

For writes, however, you need four seek requests to find where to place a new index value and normally two seeks to update the index and write the row.

Note that the preceding discussion does not mean that your application performance slowly degenerates by  $\log N$ . As long as everything is cached by the OS or the MySQL server, things become only marginally slower as the table gets bigger. After the data gets too big to be cached, things start to go much slower until your applications are only bound by disk-seeks (which increase by  $\log N$ ). To avoid this, increase the key cache size as the data grows. For `MyISAM` tables, the key cache size is controlled by the `key_buffer_size` system variable. See [Section 7.5.2, “Tuning Server Parameters”](#).

## 7.2.3. Speed of `SELECT` Queries

In general, when you want to make a slow `SELECT . . . WHERE` query faster, the first thing to check is whether you can add an index. All references between different tables should usually be done with indexes. You can use the `EXPLAIN` statement to determine which indexes are used for a `SELECT`. See [Section 7.2.1, “Optimizing Queries with `EXPLAIN`”](#), and [Section 7.4.5, “How MySQL Uses Indexes”](#).

Some general tips for speeding up queries on `MyISAM` tables:

- To help MySQL better optimize queries, use `ANALYZE TABLE` or run `myisamchk --analyze` on a table after it has been loaded with data. This updates a value for each index part that indicates the average number of rows that have the same value. (For unique indexes, this is always 1.) MySQL uses this to decide which index to choose when you join two tables based on a non-constant expression. You can check the result from the table analysis by using `SHOW INDEX FROM tbl_name` and examining the `Cardinality` value. `myisamchk --description --verbose` shows index distribution information.
- To sort an index and data according to an index, use `myisamchk --sort-index - --sort-records=1` (assuming that you want to sort on index 1). This is a good way to make queries faster if you have a unique index from which you want to read all rows in order according to the index. The first time you sort a large table this way, it may take a long time.

## 7.2.4. WHERE Clause Optimization

This section discusses optimizations that can be made for processing `WHERE` clauses. The examples use `SELECT` statements, but the same optimizations apply for `WHERE` clauses in `DELETE` and `UPDATE` statements.

Work on the MySQL optimizer is ongoing, so this section is incomplete. MySQL performs a great many optimizations, not all of which are documented here.

Some of the optimizations performed by MySQL follow:

- Removal of unnecessary parentheses:

```
((a AND b) AND c OR ((a AND b) AND (c AND d)))
-> (a AND b AND c) OR (a AND b AND c AND d)
```

- Constant folding:

```
(a<b AND b=c) AND a=5
-> b>5 AND b=c AND a=5
```

- Constant condition removal (needed because of constant folding):

```
(B>=5 AND B=5) OR (B=6 AND 5=5) OR (B=7 AND 5=6)
-> B=5 OR B=6
```

- Constant expressions used by indexes are evaluated only once.
- `COUNT( *)` on a single table without a `WHERE` is retrieved directly from the table information for `MyISAM` and `MEMORY (HASH)` tables. This is also done for any `NOT NULL` expression when used with only one table.
- Early detection of invalid constant expressions. MySQL quickly detects that some `SELECT` statements are impossible and returns no rows.
- `HAVING` is merged with `WHERE` if you do not use `GROUP BY` or aggregate functions (`COUNT( )`, `MIN( )`, and so on).
- For each table in a join, a simpler `WHERE` is constructed to get a fast `WHERE` evaluation for the table and also to skip rows as soon as possible.

- All constant tables are read first before any other tables in the query. A constant table is any of the following:
  - An empty table or a table with one row.
  - A table that is used with a `WHERE` clause on a `PRIMARY KEY` or a `UNIQUE` index, where all index parts are compared to constant expressions and are defined as `NOT NULL`.

All of the following tables are used as constant tables:

```
SELECT * FROM t WHERE primary_key=1;
SELECT * FROM t1,t2
WHERE t1.primary_key=1 AND t2.primary_key=t1.id;
```

- The best join combination for joining the tables is found by trying all possibilities. If all columns in `ORDER BY` and `GROUP BY` clauses come from the same table, that table is preferred first when joining.
- If there is an `ORDER BY` clause and a different `GROUP BY` clause, or if the `ORDER BY` or `GROUP BY` contains columns from tables other than the first table in the join queue, a temporary table is created.
- If you use the `SQL_SMALL_RESULT` option, MySQL uses an in-memory temporary table.
- Each table index is queried, and the best index is used unless the optimizer believes that it is more efficient to use a table scan. At one time, a scan was used based on whether the best index spanned more than 30% of the table, but a fixed percentage no longer determines the choice between using an index or a scan. The optimizer now is more complex and bases its estimate on additional factors such as table size, number of rows, and I/O block size.
- In some cases, MySQL can read rows from the index without even consulting the data file. If all columns used from the index are numeric, only the index tree is used to resolve the query.
- Before each row is output, those that do not match the `HAVING` clause are skipped.

Some examples of queries that are very fast:

```
SELECT COUNT(*) FROM tbl_name;
SELECT MIN(key_part1),MAX(key_part1) FROM tbl_name;
SELECT MAX(key_part2) FROM tbl_name
WHERE key_part1=constant;
SELECT ... FROM tbl_name
ORDER BY key_part1,key_part2,... LIMIT 10;
SELECT ... FROM tbl_name
ORDER BY key_part1 DESC, key_part2 DESC, ... LIMIT 10;
```

MySQL resolves the following queries using only the index tree, assuming that the indexed columns are numeric:

```
SELECT key_part1,key_part2 FROM tbl_name WHERE key_part1=val;
SELECT COUNT(*) FROM tbl_name
WHERE key_part1=val1 AND key_part2=val2;
SELECT key_part2 FROM tbl_name GROUP BY key_part1;
```

The following queries use indexing to retrieve the rows in sorted order without a separate sorting pass:

```
SELECT ... FROM tbl_name
ORDER BY key_part1,key_part2,... ;

SELECT ... FROM tbl_name
ORDER BY key_part1 DESC, key_part2 DESC, ... ;
```

## 7.2.5. Range Optimization

The [range](#) access method uses a single index to retrieve a subset of table rows that are contained within one or several index value intervals. It can be used for a single-part or multiple-part index. The following sections give a detailed description of how intervals are extracted from the [WHERE](#) clause.

### 7.2.5.1. The Range Access Method for Single-Part Indexes

For a single-part index, index value intervals can be conveniently represented by corresponding conditions in the [WHERE](#) clause, so we speak of *range conditions* rather than “intervals.”

The definition of a range condition for a single-part index is as follows:

- For both [BTREE](#) and [HASH](#) indexes, comparison of a key part with a constant value is a range condition when using the `=`, `<=>`, `IN`, `IS NULL`, or `IS NOT NULL` operators.
- For [BTREE](#) indexes, comparison of a key part with a constant value is a range condition when using the `>`, `<`, `>=`, `<=`, `BETWEEN`, `!=`, or `<>` operators, or `LIKE 'pattern'` (where `'pattern'` does not start with a wildcard).
- For all types of indexes, multiple range conditions combined with `OR` or `AND` form a range condition.

“Constant value” in the preceding descriptions means one of the following:

- A constant from the query string
- A column of a [const](#) or [system](#) table from the same join
- The result of an uncorrelated subquery
- Any expression composed entirely from subexpressions of the preceding types

Here are some examples of queries with range conditions in the [WHERE](#) clause:

```
SELECT * FROM t1
WHERE key_col > 1
AND key_col < 10;

SELECT * FROM t1
WHERE key_col = 1
OR key_col IN (15,18,20);

SELECT * FROM t1
WHERE key_col LIKE 'ab%'
OR key_col BETWEEN 'bar' AND 'foo';
```

Note that some non-constant values may be converted to constants during the constant propagation phase.

MySQL tries to extract range conditions from the [WHERE](#) clause for each of the possible indexes. Dur-

ing the extraction process, conditions that cannot be used for constructing the range condition are dropped, conditions that produce overlapping ranges are combined, and conditions that produce empty ranges are removed.

Consider the following statement, where `key1` is an indexed column and `nonkey` is not indexed:

```
SELECT * FROM t1 WHERE
  (key1 < 'abc' AND (key1 LIKE 'abcde%' OR key1 LIKE '%b')) OR
  (key1 < 'bar' AND nonkey = 4) OR
  (key1 < 'uux' AND key1 > 'z');
```

The extraction process for key `key1` is as follows:

1. Start with original `WHERE` clause:

```
(key1 < 'abc' AND (key1 LIKE 'abcde%' OR key1 LIKE '%b')) OR
(key1 < 'bar' AND nonkey = 4) OR
(key1 < 'uux' AND key1 > 'z')
```

2. Remove `nonkey = 4` and `key1 LIKE '%b'` because they cannot be used for a range scan. The correct way to remove them is to replace them with `TRUE`, so that we do not miss any matching rows when doing the range scan. Having replaced them with `TRUE`, we get:

```
(key1 < 'abc' AND (key1 LIKE 'abcde%' OR TRUE)) OR
(key1 < 'bar' AND TRUE) OR
(key1 < 'uux' AND key1 > 'z')
```

3. Collapse conditions that are always true or false:

- `(key1 LIKE 'abcde%' OR TRUE)` is always true
- `(key1 < 'uux' AND key1 > 'z')` is always false

Replacing these conditions with constants, we get:

```
(key1 < 'abc' AND TRUE) OR (key1 < 'bar' AND TRUE) OR (FALSE)
```

Removing unnecessary `TRUE` and `FALSE` constants, we obtain:

```
(key1 < 'abc') OR (key1 < 'bar')
```

4. Combining overlapping intervals into one yields the final condition to be used for the range scan:

```
(key1 < 'bar')
```

In general (and as demonstrated by the preceding example), the condition used for a range scan is less restrictive than the `WHERE` clause. MySQL performs an additional check to filter out rows that satisfy the range condition but not the full `WHERE` clause.

The range condition extraction algorithm can handle nested `AND/OR` constructs of arbitrary depth, and its output does not depend on the order in which conditions appear in `WHERE` clause.

## 7.2.5.2. The Range Access Method for Multiple-Part Indexes

Range conditions on a multiple-part index are an extension of range conditions for a single-part index. A range condition on a multiple-part index restricts index rows to lie within one or several key tuple intervals. Key tuple intervals are defined over a set of key tuples, using ordering from the index.



For example, consider a multiple-part index defined as `key1(key_part1, key_part2, key_part3)`, and the following set of key tuples listed in key order:

```
key_part1 key_part2 key_part3
NULL      1         'abc'
NULL      1         'xyz'
NULL      2         'foo'
1         1         'abc'
1         1         'xyz'
1         2         'abc'
2         1         'aaa'
```

The condition `key_part1 = 1` defines this interval:

```
(1,-inf,-inf) <= (key_part1,key_part2,key_part3) < (1,+inf,+inf)
```

The interval covers the 4th, 5th, and 6th tuples in the preceding data set and can be used by the range access method.

By contrast, the condition `key_part3 = 'abc'` does not define a single interval and cannot be used by the range access method.

The following descriptions indicate how range conditions work for multiple-part indexes in greater detail.

- For **HASH** indexes, each interval containing identical values can be used. This means that the interval can be produced only for conditions in the following form:

```
key_part1 cmp const1
AND key_part2 cmp const2
AND ...
AND key_partN cmp constN;
```

Here, `const1`, `const2`, ... are constants, `cmp` is one of the `=`, `<=>`, or `IS NULL` comparison operators, and the conditions cover all index parts. (That is, there are `N` conditions, one for each part of an `N`-part index.) For example, the following is a range condition for a three-part **HASH** index:

```
key_part1 = 1 AND key_part2 IS NULL AND key_part3 = 'foo'
```

For the definition of what is considered to be a constant, see [Section 7.2.5.1, “The Range Access Method for Single-Part Indexes”](#).

- For a **BTREE** index, an interval might be usable for conditions combined with **AND**, where each condition compares a key part with a constant value using `=`, `<=>`, `IS NULL`, `>`, `<`, `>=`, `<=`, `!=`, `<>`, `BETWEEN`, or `LIKE 'pattern'` (where `'pattern'` does not start with a wildcard). An interval can be used as long as it is possible to determine a single key tuple containing all rows that match the condition (or two intervals if `<>` or `!=` is used). For example, for this condition:

```
key_part1 = 'foo' AND key_part2 >= 10 AND key_part3 > 10
```

The single interval is:

```
('foo',10,10) < (key_part1,key_part2,key_part3) < ('foo',+inf,+inf)
```

It is possible that the created interval contains more rows than the initial condition. For example, the preceding interval includes the value `('foo', 11, 0)`, which does not satisfy the original condition.

- If conditions that cover sets of rows contained within intervals are combined with `OR`, they form a condition that covers a set of rows contained within the union of their intervals. If the conditions are combined with `AND`, they form a condition that covers a set of rows contained within the intersection of their intervals. For example, for this condition on a two-part index:

```
(key_part1 = 1 AND key_part2 < 2) OR (key_part1 > 5)
```

The intervals are:

```
(1,-inf) < (key_part1,key_part2) < (1,2)
(5,-inf) < (key_part1,key_part2)
```

In this example, the interval on the first line uses one key part for the left bound and two key parts for the right bound. The interval on the second line uses only one key part. The `key_len` column in the `EXPLAIN` output indicates the maximum length of the key prefix used.

In some cases, `key_len` may indicate that a key part was used, but that might be not what you would expect. Suppose that `key_part1` and `key_part2` can be `NULL`. Then the `key_len` column displays two key part lengths for the following condition:

```
key_part1 >= 1 AND key_part2 < 2
```

But, in fact, the condition is converted to this:

```
key_part1 >= 1 AND key_part2 IS NOT NULL
```

Section 7.2.5.1, “The Range Access Method for Single-Part Indexes”, describes how optimizations are performed to combine or eliminate intervals for range conditions on a single-part index. Analogous steps are performed for range conditions on multiple-part indexes.

## 7.2.6. IS NULL Optimization

MySQL can perform the same optimization on `col_name IS NULL` that it can use for `col_name = constant_value`. For example, MySQL can use indexes and ranges to search for `NULL` with `IS NULL`.

Examples:

```
SELECT * FROM tbl_name WHERE key_col IS NULL;
SELECT * FROM tbl_name WHERE key_col <=> NULL;
SELECT * FROM tbl_name
WHERE key_col=const1 OR key_col=const2 OR key_col IS NULL;
```

If a `WHERE` clause includes a `col_name IS NULL` condition for a column that is declared as `NOT NULL`, that expression is optimized away. This optimization does not occur in cases when the column might produce `NULL` anyway; for example, if it comes from a table on the right side of a `LEFT JOIN`.

MySQL 4.1.1 and up can also optimize the combination `col_name = expr AND col_name IS NULL`, a form that is common in resolved subqueries. `EXPLAIN` shows `ref_or_null` when this optimization is used.

This optimization can handle one `IS NULL` for any key part.

Some examples of queries that are optimized, assuming that there is an index on columns `a` and `b` of table `t2`:

```
SELECT * FROM t1 WHERE t1.a=expr OR t1.a IS NULL;

SELECT * FROM t1, t2 WHERE t1.a=t2.a OR t2.a IS NULL;

SELECT * FROM t1, t2
  WHERE (t1.a=t2.a OR t2.a IS NULL) AND t2.b=t1.b;

SELECT * FROM t1, t2
  WHERE t1.a=t2.a AND (t2.b=t1.b OR t2.b IS NULL);

SELECT * FROM t1, t2
  WHERE (t1.a=t2.a AND t2.a IS NULL AND ...)
  OR (t1.a=t2.a AND t2.a IS NULL AND ...);
```

`ref_or_null` works by first doing a read on the reference key, and then a separate search for rows with a `NULL` key value.

Note that the optimization can handle only one `IS NULL` level. In the following query, MySQL uses key lookups only on the expression `(t1.a=t2.a AND t2.a IS NULL)` and is not able to use the key part on `b`:

```
SELECT * FROM t1, t2
  WHERE (t1.a=t2.a AND t2.a IS NULL)
  OR (t1.b=t2.b AND t2.b IS NULL);
```

## 7.2.7. DISTINCT Optimization

`DISTINCT` combined with `ORDER BY` needs a temporary table in many cases.

Because `DISTINCT` may use `GROUP BY`, you should be aware of how MySQL works with columns in `ORDER BY` or `HAVING` clauses that are not part of the selected columns. See [Section 12.10.3, “GROUP BY and HAVING with Hidden Fields”](#).

In most cases, a `DISTINCT` clause can be considered as a special case of `GROUP BY`. For example, the following two queries are equivalent:

```
SELECT DISTINCT c1, c2, c3 FROM t1 WHERE c1 > const;

SELECT c1, c2, c3 FROM t1 WHERE c1 > const GROUP BY c1, c2, c3;
```

Due to this equivalence, the optimizations applicable to `GROUP BY` queries can be also applied to queries with a `DISTINCT` clause. Thus, for more details on the optimization possibilities for `DISTINCT` queries, see [Section 7.2.10, “GROUP BY Optimization”](#).

When combining `LIMIT row_count` with `DISTINCT`, MySQL stops as soon as it finds `row_count` unique rows.

If you do not use columns from all tables named in a query, MySQL stops scanning any unused tables as soon as it finds the first match. In the following case, assuming that `t1` is used before `t2` (which you can check with `EXPLAIN`), MySQL stops reading from `t2` (for any particular row in `t1`) when it finds the first row in `t2`:

```
SELECT DISTINCT t1.a FROM t1, t2 where t1.a=t2.a;
```

## 7.2.8. LEFT JOIN and RIGHT JOIN Optimization

MySQL implements a `A LEFT JOIN B join_condition` as follows:

- Table `B` is set to depend on table `A` and all tables on which `A` depends.

- Table *A* is set to depend on all tables (except *B*) that are used in the `LEFT JOIN` condition.
- The `LEFT JOIN` condition is used to decide how to retrieve rows from table *B*. (In other words, any condition in the `WHERE` clause is not used.)
- All standard join optimizations are performed, with the exception that a table is always read after all tables on which it depends. If there is a circular dependence, MySQL issues an error.
- All standard `WHERE` optimizations are performed.
- If there is a row in *A* that matches the `WHERE` clause, but there is no row in *B* that matches the `ON` condition, an extra *B* row is generated with all columns set to `NULL`.
- If you use `LEFT JOIN` to find rows that do not exist in some table and you have the following test: `col_name IS NULL` in the `WHERE` part, where `col_name` is a column that is declared as `NOT NULL`, MySQL stops searching for more rows (for a particular key combination) after it has found one row that matches the `LEFT JOIN` condition.

The implementation of `RIGHT JOIN` is analogous to that of `LEFT JOIN` with the roles of the tables reversed.

The join optimizer calculates the order in which tables should be joined. The table read order forced by `LEFT JOIN` or `STRAIGHT_JOIN` helps the join optimizer do its work much more quickly, because there are fewer table permutations to check. Note that this means that if you do a query of the following type, MySQL does a full scan on *b* because the `LEFT JOIN` forces it to be read before *d*:

```
SELECT *
  FROM a JOIN b LEFT JOIN c ON (c.key=a.key) LEFT JOIN d ON (d.key=a.key)
 WHERE b.key=d.key;
```

The fix in this case is reverse the order in which *a* and *b* are listed in the `FROM` clause:

```
SELECT *
  FROM b JOIN a LEFT JOIN c ON (c.key=a.key) LEFT JOIN d ON (d.key=a.key)
 WHERE b.key=d.key;
```

Starting from 4.0.14, for a `LEFT JOIN`, if the `WHERE` condition is always false for the generated `NULL` row, the `LEFT JOIN` is changed to a normal join. For example, the `WHERE` clause would be false in the following query if `t2.column1` were `NULL`:

```
SELECT * FROM t1 LEFT JOIN t2 ON (column1) WHERE t2.column2=5;
```

Therefore, it is safe to convert the query to a normal join:

```
SELECT * FROM t1, t2 WHERE t2.column2=5 AND t1.column1=t2.column1;
```

This can be made faster because MySQL can use table *t2* before table *t1* if doing so would result in a better query plan. To force a specific table order, use `STRAIGHT_JOIN`.

## 7.2.9. ORDER BY Optimization

In some cases, MySQL can use an index to satisfy an `ORDER BY` clause without doing any extra sorting.

The index can also be used even if the `ORDER BY` does not match the index exactly, as long as all of the unused portions of the index and all the extra `ORDER BY` columns are constants in the `WHERE` clause. The following queries use the index to resolve the `ORDER BY` part:

```

SELECT * FROM t1
ORDER BY key_part1, key_part2, ... ;

SELECT * FROM t1
WHERE key_part1=constant
ORDER BY key_part2;

SELECT * FROM t1
ORDER BY key_part1 DESC, key_part2 DESC;

SELECT * FROM t1
WHERE key_part1=1
ORDER BY key_part1 DESC, key_part2 DESC;

```

In some cases, MySQL *cannot* use indexes to resolve the `ORDER BY`, although it still uses indexes to find the rows that match the `WHERE` clause. These cases include the following:

- You use `ORDER BY` on different keys:

```
SELECT * FROM t1 ORDER BY key1, key2;
```

- You use `ORDER BY` on non-consecutive parts of a key:

```
SELECT * FROM t1 WHERE key2=constant ORDER BY key_part2;
```

- You mix `ASC` and `DESC`:

```
SELECT * FROM t1 ORDER BY key_part1 DESC, key_part2 ASC;
```

- The key used to fetch the rows is not the same as the one used in the `ORDER BY`:

```
SELECT * FROM t1 WHERE key2=constant ORDER BY key1;
```

- You are joining many tables, and the columns in the `ORDER BY` are not all from the first non-constant table that is used to retrieve rows. (This is the first table in the `EXPLAIN` output that does not have a `const` join type.)
- You have different `ORDER BY` and `GROUP BY` expressions.
- The type of table index used does not store rows in order. For example, this is true for a `HASH` index in a `MEMORY` table.

With `EXPLAIN SELECT ... ORDER BY`, you can check whether MySQL can use indexes to resolve the query. It cannot if you see `Using filesort` in the `Extra` column. See [Section 7.2.1, “Optimizing Queries with EXPLAIN”](#).

Prior to MySQL 4.1, in those cases where MySQL must sort the result, it uses the following `filesort` algorithm:

1. Read all rows according to key or by table scanning. Rows that do not match the `WHERE` clause are skipped.
2. For each row, store a pair of values in a buffer (the sort key and the row pointer). The size of the buffer is the value of the `sort_buffer_size` system variable.
3. When the buffer gets full, run a `qsort` (quicksort) on it and store the result in a temporary file. Save a pointer to the sorted block. (If all pairs fit into the sort buffer, no temporary file is created.)

4. Repeat the preceding steps until all rows have been read.
5. Do a multi-merge of up to `MERGEBUFF` (7) regions to one block in another temporary file. Repeat until all blocks from the first file are in the second file.
6. Repeat the following until there are fewer than `MERGEBUFF2` (15) blocks left.
7. On the last multi-merge, only the pointer to the row (the last part of the sort key) is written to a result file.
8. Read the rows in sorted order by using the row pointers in the result file. To optimize this, we read in a big block of row pointers, sort them, and use them to read the rows in sorted order into a row buffer. The size of the buffer is the value of the `read_rnd_buffer_size` system variable. The code for this step is in the `sql/records.cc` source file.

One problem with this approach is that it reads rows twice: One time when evaluating the `WHERE` clause, and again after sorting the pair values. And even if the rows were accessed successively the first time (for example, if a table scan is done), the second time they are accessed randomly. (The sort keys are ordered, but the row positions are not.)

In MySQL 4.1 and up, a `filesort` optimization is used that records not only the sort key value and row position, but also the columns required for the query. This avoids reading the rows twice. The modified `filesort` algorithm works like this:

1. Read the rows that match the `WHERE` clause.
2. For each row, record a tuple of values consisting of the sort key value and row position, and also the columns required for the query.
3. Sort the tuples by sort key value
4. Retrieve the rows in sorted order, but read the required columns directly from the sorted tuples rather than by accessing the table a second time.

Using the modified `filesort` algorithm, the tuples are longer than the pairs used in the original method, and fewer of them fit in the sort buffer (the size of which is given by `sort_buffer_size`). As a result, it is possible for the extra I/O to make the modified approach slower, not faster. To avoid a slowdown, the optimization is used only if the total size of the extra columns in the sort tuple does not exceed the value of the `max_length_for_sort_data` system variable. (A symptom of setting the value of this variable too high is that you should see high disk activity and low CPU activity.)

If you want to increase `ORDER BY` speed, check whether you can get MySQL to use indexes rather than an extra sorting phase. If this is not possible, you can try the following strategies:

- Increase the size of the `sort_buffer_size` variable.
- Increase the size of the `read_rnd_buffer_size` variable.
- Change `tmpdir` to point to a dedicated filesystem with large amounts of empty space. If you use MySQL 4.1 or later, this option accepts several paths that are used in round-robin fashion. Paths should be separated by colon characters (':') on Unix and semicolon characters (;) on Windows, NetWare, and OS/2. You can use this feature to spread the load across several directories. *Note:* The paths should be for directories in filesystems that are located on different *physical* disks, not different partitions on the same disk.

By default, MySQL sorts all `GROUP BY col1, col2, ...` queries as if you specified `ORDER BY col1, col2, ...` in the query as well. If you include an `ORDER BY` clause explicitly that contains the same column list, MySQL optimizes it away without any speed penalty, although the sorting still occurs. If a query includes `GROUP BY` but you want to avoid the overhead of sorting the result, you can suppress sorting by specifying `ORDER BY NULL`. For example:

```
INSERT INTO foo
SELECT a, COUNT(*) FROM bar GROUP BY a ORDER BY NULL;
```

## 7.2.10. GROUP BY Optimization

The most general way to satisfy a `GROUP BY` clause is to scan the whole table and create a new temporary table where all rows from each group are consecutive, and then use this temporary table to discover groups and apply aggregate functions (if any). In some cases, MySQL is able to do much better than that and to avoid creation of temporary tables by using index access.

The most important preconditions for using indexes for `GROUP BY` are that all `GROUP BY` columns reference attributes from the same index, and that the index stores its keys in order (for example, this is a `BTREE` index, and not a `HASH` index). Whether use of temporary tables can be replaced by index access also depends on which parts of an index are used in a query, the conditions specified for these parts, and the selected aggregate functions.

There are two ways to execute a `GROUP BY` query via index access, as detailed in the following sections. In the first method, the grouping operation is applied together with all range predicates (if any). The second method first performs a range scan, and then groups the resulting tuples.

### 7.2.10.1. Tight index scan

A tight index scan may be either a full index scan or a range index scan, depending on the query conditions.

When the conditions for a loose index scan are not met, it is still possible to avoid creation of temporary tables for `GROUP BY` queries. If there are range conditions in the `WHERE` clause, this method reads only the keys that satisfy these conditions. Otherwise, it performs an index scan. Because this method reads all keys in each range defined by the `WHERE` clause, or scans the whole index if there are no range conditions, we term it a *tight index scan*. Notice that with a tight index scan, the grouping operation is performed only after all keys that satisfy the range conditions have been found.

For this method to work, it is sufficient that there is a constant equality condition for all columns in a query referring to parts of the key coming before or in between parts of the `GROUP BY` key. The constants from the equality conditions fill in any “gaps” in the search keys so that it is possible to form complete prefixes of the index. These index prefixes then can be used for index lookups. If we require sorting of the `GROUP BY` result, and it is possible to form search keys that are prefixes of the index, MySQL also avoids extra sorting operations because searching with prefixes in an ordered index already retrieves all the keys in order.

The following queries do not work with the loose index scan access method described earlier, but still work with the tight index scan access method (assuming that there is an index `idx(c1, c2, c3)` on table `t1(c1, c2, c3, c4)`).

- There is a gap in the `GROUP BY`, but it is covered by the condition `c2 = 'a'`:

```
SELECT c1, c2, c3 FROM t1 WHERE c2 = 'a' GROUP BY c1, c3;
```

- The `GROUP BY` does not begin with the first part of the key, but there is a condition that provides a constant for that part:

```
SELECT c1, c2, c3 FROM t1 WHERE c1 = 'a' GROUP BY c2, c3;
```

## 7.2.11. LIMIT Optimization

In some cases, MySQL handles a query differently when you are using `LIMIT row_count` and not using `HAVING`:

- If you are selecting only a few rows with `LIMIT`, MySQL uses indexes in some cases when normally it would prefer to do a full table scan.
- If you use `LIMIT row_count` with `ORDER BY`, MySQL ends the sorting as soon as it has found the first `row_count` rows of the sorted result, rather than sorting the entire result. If ordering is done by using an index, this is very fast. If a filesort must be done, all rows that match the query without the `LIMIT` clause must be selected, and most or all of them must be sorted, before it can be ascertained that the first `row_count` rows have been found. In either case, after the initial rows have been found, there is no need to sort any remainder of the result set, and MySQL does not do so.
- When combining `LIMIT row_count` with `DISTINCT`, MySQL stops as soon as it finds `row_count` unique rows.
- In some cases, a `GROUP BY` can be resolved by reading the key in order (or doing a sort on the key) and then calculating summaries until the key value changes. In this case, `LIMIT row_count` does not calculate any unnecessary `GROUP BY` values.
- As soon as MySQL has sent the required number of rows to the client, it aborts the query unless you are using `SQL_CALC_FOUND_ROWS`.
- `LIMIT 0` quickly returns an empty set. This can be useful for checking the validity of a query. When using one of the MySQL APIs, it can also be employed for obtaining the types of the result columns. (This trick does not work in the MySQL Monitor (the `mysql` program), which merely displays `Empty set` in such cases; you should instead use `SHOW COLUMNS` or `DESCRIBE` for this purpose.)
- When the server uses temporary tables to resolve the query, it uses the `LIMIT row_count` clause to calculate how much space is required.

## 7.2.12. How to Avoid Table Scans

The output from `EXPLAIN` shows `ALL` in the `type` column when MySQL uses a table scan to resolve a query. This usually happens under the following conditions:

- The table is so small that it is faster to perform a table scan than to bother with a key lookup. This is common for tables with fewer than 10 rows and a short row length.
- There are no usable restrictions in the `ON` or `WHERE` clause for indexed columns.
- You are comparing indexed columns with constant values and MySQL has calculated (based on the index tree) that the constants cover too large a part of the table and that a table scan would be faster. See [Section 7.2.4, “WHERE Clause Optimization”](#).
- You are using a key with low cardinality (many rows match the key value) through another column. In this case, MySQL assumes that by using the key it probably will do many key lookups and that a table scan would be faster.



For small tables, a table scan often is appropriate and the performance impact is negligible. For large tables, try the following techniques to avoid having the optimizer incorrectly choose a table scan:

- Use `ANALYZE TABLE tbl_name` to update the key distributions for the scanned table. See [Section 13.5.2.1, “ANALYZE TABLE Syntax”](#).
- Use `FORCE INDEX` for the scanned table to tell MySQL that table scans are very expensive compared to using the given index:

```
SELECT * FROM t1, t2 FORCE INDEX (index_for_column)
WHERE t1.col_name=t2.col_name;
```

See [Section 13.2.7, “SELECT Syntax”](#).

- Start `mysqld` with the `--max-seeks-for-key=1000` option or use `SET max_seeks_for_key=1000` to tell the optimizer to assume that no key scan causes more than 1,000 key seeks. See [Section 5.2.2, “System Variables”](#).

## 7.2.13. Speed of `INSERT` Statements

The time required for inserting a row is determined by the following factors, where the numbers indicate approximate proportions:

- Connecting: (3)
- Sending query to server: (2)
- Parsing query: (2)
- Inserting row: (1 × size of row)
- Inserting indexes: (1 × number of indexes)
- Closing: (1)

This does not take into consideration the initial overhead to open tables, which is done once for each concurrently running query.

The size of the table slows down the insertion of indexes by  $\log N$ , assuming B-tree indexes.

You can use the following methods to speed up inserts:

- If you are inserting many rows from the same client at the same time, use `INSERT` statements with multiple `VALUES` lists to insert several rows at a time. This is considerably faster (many times faster in some cases) than using separate single-row `INSERT` statements. If you are adding data to a non-empty table, you can tune the `bulk_insert_buffer_size` variable to make data insertion even faster. See [Section 5.2.2, “System Variables”](#).
- If you are inserting a lot of rows from different clients, you can get higher speed by using the `INSERT DELAYED` statement. See [Section 13.2.4.2, “INSERT DELAYED Syntax”](#).
- For a `MyISAM` table, you can use concurrent inserts to add rows at the same time that `SELECT` statements are running if there are no deleted rows in middle of the table. See [Section 7.3.3, “Concurrent Inserts”](#).

- When loading a table from a text file, use `LOAD DATA INFILE`. This is usually 20 times faster than using `INSERT` statements. See [Section 13.2.5, “LOAD DATA INFILE Syntax”](#).
- With some extra work, it is possible to make `LOAD DATA INFILE` run even faster for a `MyISAM` table when the table has many indexes. Use the following procedure:
  1. Optionally create the table with `CREATE TABLE`.
  2. Execute a `FLUSH TABLES` statement or a `mysqladmin flush-tables` command.
  3. Use `myisamchk --keys-used=0 -rq /path/to/db/tbl_name`. This removes all use of indexes for the table.
  4. Insert data into the table with `LOAD DATA INFILE`. This does not update any indexes and therefore is very fast.
  5. If you intend only to read from the table in the future, use `myisampack` to compress it. See [Section 14.1.3.3, “Compressed Table Characteristics”](#).
  6. Re-create the indexes with `myisamchk -rq /path/to/db/tbl_name`. This creates the index tree in memory before writing it to disk, which is much faster than updating the index during `LOAD DATA INFILE` because it avoids lots of disk seeks. The resulting index tree is also perfectly balanced.
  7. Execute a `FLUSH TABLES` statement or a `mysqladmin flush-tables` command.

Note that `LOAD DATA INFILE` performs the preceding optimization automatically if the `MyISAM` table into which you insert data is empty. The main difference is that you can let `myisamchk` allocate much more temporary memory for the index creation than you might want the server to allocate for index re-creation when it executes the `LOAD DATA INFILE` statement.

As of MySQL 4.0, you can also disable or enable the indexes for a `MyISAM` table by using the following statements rather than `myisamchk`. If you use these statements, you can skip the `FLUSH TABLE` operations:

```
ALTER TABLE tbl_name DISABLE KEYS;
ALTER TABLE tbl_name ENABLE KEYS;
```

- To speed up `INSERT` operations that are performed with multiple statements for non-transactional tables, lock your tables:

```
LOCK TABLES a WRITE;
INSERT INTO a VALUES (1,23),(2,34),(4,33);
INSERT INTO a VALUES (8,26),(6,29);
...
UNLOCK TABLES;
```

This benefits performance because the index buffer is flushed to disk only once, after all `INSERT` statements have completed. Normally, there would be as many index buffer flushes as there are `INSERT` statements. Explicit locking statements are not needed if you can insert all rows with a single `INSERT`.

To obtain faster insertions, for transactional tables, you should use `START TRANSACTION` and `COMMIT` instead of `LOCK TABLES`.

Locking also lowers the total time for multiple-connection tests, although the maximum wait time for individual connections might go up because they wait for locks. For example:

1. Connection 1 does 1000 inserts

2. Connections 2, 3, and 4 do 1 insert
3. Connection 5 does 1000 inserts

If you do not use locking, connections 2, 3, and 4 finish before 1 and 5. If you use locking, connections 2, 3, and 4 probably do not finish before 1 or 5, but the total time should be about 40% faster.

`INSERT`, `UPDATE`, and `DELETE` operations are very fast in MySQL, but you can obtain better overall performance by adding locks around everything that does more than about five inserts or updates in a row. If you do very many inserts in a row, you could do a `LOCK TABLES` followed by an `UNLOCK TABLES` once in a while (each 1,000 rows or so) to allow other threads access to the table. This would still result in a nice performance gain.

`INSERT` is still much slower for loading data than `LOAD DATA INFILE`, even when using the strategies just outlined.

- To increase performance for `MyISAM` tables, for both `LOAD DATA INFILE` and `INSERT`, enlarge the key cache by increasing the `key_buffer_size` system variable. See [Section 7.5.2, “Tuning Server Parameters”](#).

## 7.2.14. Speed of `UPDATE` Statements

An update statement is optimized like a `SELECT` query with the additional overhead of a write. The speed of the write depends on the amount of data being updated and the number of indexes that are updated. Indexes that are not changed do not get updated.

Another way to get fast updates is to delay updates and then do many updates in a row later. Performing multiple updates together is much quicker than doing one at a time if you lock the table.

For a `MyISAM` table that uses dynamic row format, updating a row to a longer total length may split the row. If you do this often, it is very important to use `OPTIMIZE TABLE` occasionally. See [Section 13.5.2.5, “OPTIMIZE TABLE Syntax”](#).

## 7.2.15. Speed of `DELETE` Statements

The time required to delete individual rows is exactly proportional to the number of indexes. To delete rows more quickly, you can increase the size of the key cache by increasing the `key_buffer_size` system variable. See [Section 7.5.2, “Tuning Server Parameters”](#).

To delete all rows from a table, `TRUNCATE TABLE tbl_name` is faster than `DELETE FROM tbl_name`. See [Section 13.2.9, “TRUNCATE Syntax”](#).

## 7.2.16. Other Optimization Tips

This section lists a number of miscellaneous tips for improving query processing speed:

- Use persistent connections to the database to avoid connection overhead. If you cannot use persistent connections and you are initiating many new connections to the database, you may want to change the value of the `thread_cache_size` variable. See [Section 7.5.2, “Tuning Server Parameters”](#).
- Always check whether all your queries really use the indexes that you have created in the tables. In MySQL, you can do this with the `EXPLAIN` statement. See [Section 7.2.1, “Optimizing Queries with EXPLAIN”](#).

- Try to avoid complex `SELECT` queries on `MyISAM` tables that are updated frequently, to avoid problems with table locking that occur due to contention between readers and writers.
- With `MyISAM` tables that have no deleted rows in the middle, you can insert rows at the end at the same time that another query is reading from the table. If it is important to be able to do this, you should consider using the table in ways that avoid deleting rows. Another possibility is to run `OPTIMIZE TABLE` to defragment the table after you have deleted a lot of rows from it. See [Section 14.1, “The MyISAM Storage Engine”](#).
- To fix any compression issues that may have occurred with `ARCHIVE` tables, you can use `OPTIMIZE TABLE`. See [Section 14.7, “The ARCHIVE Storage Engine”](#).
- Use `ALTER TABLE ... ORDER BY expr1, expr2, ...` if you usually retrieve rows in `expr1, expr2, ...` order. By using this option after extensive changes to the table, you may be able to get higher performance.
- In some cases, it may make sense to introduce a column that is “hashed” based on information from other columns. If this column is short and reasonably unique, it may be much faster than a “wide” index on many columns. In MySQL, it is very easy to use this extra column:

```
SELECT * FROM tbl_name
WHERE hash_col=MD5(CONCAT(col1,col2))
AND col1='constant' AND col2='constant';
```

- For `MyISAM` tables that change frequently, you should try to avoid all variable-length columns (`VARCHAR`, `BLOB`, and `TEXT`). The table uses dynamic row format if it includes even a single variable-length column. See [Chapter 14, Storage Engines and Table Types](#).
- It is normally not useful to split a table into different tables just because the rows become large. In accessing a row, the biggest performance hit is the disk seek needed to find the first byte of the row. After finding the data, most modern disks can read the entire row fast enough for most applications. The only cases where splitting up a table makes an appreciable difference is if it is a `MyISAM` table using dynamic row format that you can change to a fixed row size, or if you very often need to scan the table but do not need most of the columns. See [Chapter 14, Storage Engines and Table Types](#).
- If you often need to calculate results such as counts based on information from a lot of rows, it may be preferable to introduce a new table and update the counter in real time. An update of the following form is very fast:

```
UPDATE tbl_name SET count_col=count_col+1 WHERE key_col=constant;
```

This is very important when you use MySQL storage engines such as `MyISAM` and `ISAM` that have only table-level locking (multiple readers with single writers). This also gives better performance with most database systems, because the row locking manager in this case has less to do.

- If you need to collect statistics from large log tables, use summary tables instead of scanning the entire log table. Maintaining the summaries should be much faster than trying to calculate statistics “live.” Regenerating new summary tables from the logs when things change (depending on business decisions) is faster than changing the running application.
- If possible, you should classify reports as “live” or as “statistical,” where data needed for statistical reports is created only from summary tables that are generated periodically from the live data.
- Take advantage of the fact that columns have default values. Insert values explicitly only when the value to be inserted differs from the default. This reduces the parsing that MySQL must do and improves the insert speed.
- In some cases, it is convenient to pack and store data into a `BLOB` column. In this case, you must

provide code in your application to pack and unpack information, but this may save a lot of accesses at some stage. This is practical when you have data that does not conform well to a rows-and-columns table structure.

- Normally, you should try to keep all data non-redundant (observing what is referred to in database theory as *third normal form*). However, there may be situations in which it can be advantageous to duplicate information or create summary tables to gain more speed.
- UDFs (user-defined functions) may be a good way to get more performance for some tasks. See [Section 19.2, “Adding New Functions to MySQL”](#), for more information.
- You can always gain something by caching queries or answers in your application and then performing many inserts or updates together. If your database system supports table locks (as do MySQL and Oracle), this should help to ensure that the index cache is only flushed once after all updates. You can also take advantage of MySQL's query cache to achieve similar results; see [Section 5.13, “The MySQL Query Cache”](#).
- Use `INSERT DELAYED` when you do not need to know when your data is written. This reduces the overall insertion impact because many rows can be written with a single disk write.
- Use `INSERT LOW_PRIORITY` when you want to give `SELECT` statements higher priority than your inserts.
- Use `SELECT HIGH_PRIORITY` to get retrievals that jump the queue. That is, the `SELECT` is executed even if there is another client waiting to do a write.
- Use multiple-row `INSERT` statements to store many rows with one SQL statement. Many SQL servers support this, including MySQL.
- Use `LOAD DATA INFILE` to load large amounts of data. This is faster than using `INSERT` statements.
- Use `AUTO_INCREMENT` columns to generate unique values.
- Use `OPTIMIZE TABLE` once in a while to avoid fragmentation with dynamic-format `MyISAM` tables. See [Section 14.1.3, “MyISAM Table Storage Formats”](#).
- Use `MEMORY (HEAP)` tables when possible to get more speed. See [Section 14.4, “The MEMORY \(HEAP\) Storage Engine”](#). `MEMORY` tables are useful for non-critical data that is accessed often, such as information about the last displayed banner for users who don't have cookies enabled in their Web browser. User sessions are another alternative available in many Web application environments for handling volatile state data.
- With Web servers, images and other binary assets should normally be stored as files. That is, store only a reference to the file rather than the file itself in the database. Most Web servers are better at caching files than database contents, so using files is generally faster.
- Columns with identical information in different tables should be declared to have identical data types so that joins based on the corresponding columns will be faster. Before MySQL 3.23, you get slow joins otherwise.

Try to keep column names simple. For example, in a table named `customer`, use a column name of `name` instead of `customer_name`. To make your names portable to other SQL servers, you should keep them shorter than 18 characters.

- If you need really high speed, you should take a look at the low-level interfaces for data storage that the different SQL servers support. For example, by accessing the MySQL `MyISAM` storage engine directly, you could get a speed increase of two to five times compared to using the SQL interface. To be able to do this, the data must be on the same server as the application, and usually it should only

be accessed by one process (because external file locking is really slow). One could eliminate these problems by introducing low-level `MyISAM` commands in the MySQL server (this could be one easy way to get more performance if needed). By carefully designing the database interface, it should be quite easy to support this type of optimization.

- If you are using numerical data, it is faster in many cases to access information from a database (using a live connection) than to access a text file. Information in the database is likely to be stored in a more compact format than in the text file, so accessing it involves fewer disk accesses. You also save code in your application because you need not parse your text files to find line and column boundaries.
- Replication can provide a performance benefit for some operations. You can distribute client retrievals among replication servers to split up the load. To avoid slowing down the master while making backups, you can make backups using a slave server. See [Chapter 6, Replication](#).
- Declaring a `MyISAM` table with the `DELAY_KEY_WRITE=1` table option makes index updates faster because they are not flushed to disk until the table is closed. The downside is that if something kills the server while such a table is open, you should ensure that the table is okay by running the server with the `--myisam-recover` option, or by running `myisamchk` before restarting the server. (However, even in this case, you should not lose anything by using `DELAY_KEY_WRITE`, because the key information can always be generated from the data rows.)

## 7.3. Locking Issues

### 7.3.1. Locking Methods

MySQL uses table-level locking for `ISAM`, `MyISAM`, and `MEMORY (HEAP)` tables, page-level locking for `BDB` tables, and row-level locking for `InnoDB` tables.

In many cases, you can make an educated guess about which locking type is best for an application, but generally it is difficult to say that a given lock type is better than another. Everything depends on the application and different parts of an application may require different lock types.

To decide whether you want to use a storage engine with row-level locking, you should look at what your application does and what mix of select and update statements it uses. For example, most Web applications perform many selects, relatively few deletes, updates based mainly on key values, and inserts into a few specific tables. The base MySQL `MyISAM` setup is very well tuned for this.

Table locking in MySQL is deadlock-free for storage engines that use table-level locking. Deadlock avoidance is managed by always requesting all needed locks at once at the beginning of a query and always locking the tables in the same order.

The table-locking method MySQL uses for `WRITE` locks works as follows:

- If there are no locks on the table, put a write lock on it.
- Otherwise, put the lock request in the write lock queue.

The table-locking method MySQL uses for `READ` locks works as follows:

- If there are no write locks on the table, put a read lock on it.
- Otherwise, put the lock request in the read lock queue.

When a lock is released, the lock is made available to the threads in the write lock queue and then to the threads in the read lock queue. This means that if you have many updates for a table, `SELECT` statements wait until there are no more updates.

Starting in MySQL 3.23.33, you can analyze the table lock contention on your system by checking the `Table_locks_waited` and `Table_locks_immediate` status variables:

```
mysql> SHOW STATUS LIKE 'Table%';
+-----+-----+
| Variable_name | Value |
+-----+-----+
| Table_locks_immediate | 1151552 |
| Table_locks_waited | 15324 |
+-----+-----+
```

As of MySQL 3.23.7 (3.23.25 for Windows), if a `MyISAM` table contains no free blocks in the middle, rows always are inserted at the end of the data file. In this case, you can freely mix concurrent `INSERT` and `SELECT` statements for a `MyISAM` table without locks. That is, you can insert rows into a `MyISAM` table at the same time other clients are reading from it. (Holes can result from rows having been deleted from or updated in the middle of the table. If there are holes, concurrent inserts are disabled but are re-enabled automatically when all holes have been filled with new data.)

If you want to perform many `INSERT` and `SELECT` operations on a table when concurrent inserts are not possible, you can insert rows in a temporary table and update the real table with the rows from the temporary table once in a while. This can be done with the following code:

```
mysql> LOCK TABLES real_table WRITE, insert_table WRITE;
mysql> INSERT INTO real_table SELECT * FROM insert_table;
mysql> TRUNCATE TABLE insert_table;
mysql> UNLOCK TABLES;
```

`InnoDB` uses row locks and `BDB` uses page locks. For these two storage engines, deadlocks are possible because they automatically acquire locks during the processing of SQL statements, not at the start of the transaction.

Advantages of row-level locking:

- Fewer lock conflicts when accessing different rows in many threads.
- Fewer changes for rollbacks.
- Possible to lock a single row for a long time.

Disadvantages of row-level locking:

- Requires more memory than page-level or table-level locks.
- Slower than page-level or table-level locks when used on a large part of the table because you must acquire many more locks.
- Definitely much slower than other locks if you often do `GROUP BY` operations on a large part of the data or if you must scan the entire table frequently.

Table locks are superior to page-level or row-level locks in the following cases:

- Most statements for the table are reads.

- A mix of reads and writes, where writes are updates or deletes for a single row that can be fetched with one key read:

```
UPDATE tbl_name SET column=value WHERE unique_key_col=key_value;
DELETE FROM tbl_name WHERE unique_key_col=key_value;
```

- `SELECT` combined with concurrent `INSERT` statements, and very few `UPDATE` or `DELETE` statements.
- Many scans or `GROUP BY` operations on the entire table without any writers.

With higher-level locks, you can more easily tune applications by supporting locks of different types, because the lock overhead is less than for row-level locks.

Options other than row-level or page-level locking:

- Versioning (such as that used in MySQL for concurrent inserts) where it is possible to have one writer at the same time as many readers. This means that the database or table supports different views for the data depending on when access begins. Other common terms for this are “time travel,” “copy on write,” or “copy on demand.”
- Copy on demand is in many cases superior to page-level or row-level locking. However, in the worst case, it can use much more memory than using normal locks.
- Instead of using row-level locks, you can employ application-level locks, such as `GET_LOCK()` and `RELEASE_LOCK()` in MySQL. These are advisory locks, so they work only in well-behaved applications. (See [Section 12.9.4, “Miscellaneous Functions”](#).)

## 7.3.2. Table Locking Issues

To achieve a very high lock speed, MySQL uses table locking (instead of page, row, or column locking) for all storage engines except `InnoDB` and `BDB`.

For `InnoDB` and `BDB` tables, MySQL uses only table locking if you explicitly lock the table with `LOCK TABLES`. For these storage engines, we recommend that you not use `LOCK TABLES` at all, because `InnoDB` uses automatic row-level locking and `BDB` uses page-level locking to ensure transaction isolation.

For large tables, table locking is much better than row locking for most applications, but there are some pitfalls:

- Table locking enables many threads to read from a table at the same time, but if a thread wants to write to a table, it must first get exclusive access. During the update, all other threads that want to access this particular table must wait until the update is done.
- Table updates normally are considered to be more important than table retrievals, so they are given higher priority. This should ensure that updates to a table are not “starved” even if there is heavy `SELECT` activity for the table.
- Table locking causes problems in cases such as when a thread is waiting because the disk is full and free space needs to become available before the thread can proceed. In this case, all threads that want to access the problem table are also put in a waiting state until more disk space is made available.

Table locking is also disadvantageous under the following scenario:



- A client issues a `SELECT` that takes a long time to run.
- Another client then issues an `UPDATE` on the same table. This client waits until the `SELECT` is finished.
- Another client issues another `SELECT` statement on the same table. Because `UPDATE` has higher priority than `SELECT`, this `SELECT` waits for the `UPDATE` to finish, *and* for the first `SELECT` to finish.

The following items describe some ways to avoid or reduce contention caused by table locking:

- Try to get the `SELECT` statements to run faster so that they lock tables for a shorter time. You might have to create some summary tables to do this.
- Start `mysqld` with `--low-priority-updates`. This gives all statements that update (modify) a table lower priority than `SELECT` statements. In this case, the second `SELECT` statement in the preceding scenario would execute before the `UPDATE` statement, and would not need to wait for the first `SELECT` to finish.
- You can specify that all updates issued in a specific connection should be done with low priority by using the `SET LOW_PRIORITY_UPDATES=1` statement. See [Section 13.5.3, “SET Syntax”](#).
- You can give a specific `INSERT`, `UPDATE`, or `DELETE` statement lower priority with the `LOW_PRIORITY` attribute.
- You can give a specific `SELECT` statement higher priority with the `HIGH_PRIORITY` attribute. See [Section 13.2.7, “SELECT Syntax”](#).
- Starting from MySQL 3.23.7, you can start `mysqld` with a low value for the `max_write_lock_count` system variable to force MySQL to temporarily elevate the priority of all `SELECT` statements that are waiting for a table after a specific number of inserts to the table occur. This allows `READ` locks after a certain number of `WRITE` locks.
- If you have problems with `INSERT` combined with `SELECT`, you might want to consider switching to `MyISAM` tables, which support concurrent `SELECT` and `INSERT` statements. (See [Section 7.3.3, “Concurrent Inserts”](#).)
- If you mix inserts and deletes on the same table, `INSERT DELAYED` may be of great help. See [Section 13.2.4.2, “INSERT DELAYED Syntax”](#).
- If you have problems with mixed `SELECT` and `DELETE` statements, the `LIMIT` option to `DELETE` may help. See [Section 13.2.1, “DELETE Syntax”](#).
- Using `SQL_BUFFER_RESULT` with `SELECT` statements can help to make the duration of table locks shorter. See [Section 13.2.7, “SELECT Syntax”](#).
- You could change the locking code in `mysys/thr_lock.c` to use a single queue. In this case, write locks and read locks would have the same priority, which might help some applications.

Here are some tips concerning table locks in MySQL:

- Concurrent users are not a problem if you do not mix updates with selects that need to examine many rows in the same table.
- You can use `LOCK TABLES` to increase speed, because many updates within a single lock is much

faster than updating without locks. Splitting table contents into separate tables may also help.

- If you encounter speed problems with table locks in MySQL, you may be able to improve performance by converting some of your tables to [InnoDB](#) or [BDB](#) tables. See [Section 14.2, “The InnoDB Storage Engine”](#), and [Section 14.5, “The BDB \(BerkeleyDB\) Storage Engine”](#).

### 7.3.3. Concurrent Inserts

For a [MyISAM](#) table, you can use concurrent inserts to add rows at the same time that [SELECT](#) statements are running if there are no deleted rows in middle of the table.

Under circumstances where concurrent inserts can be used, there is seldom any need to use the [DELAYED](#) modifier for [INSERT](#) statements. See [Section 13.2.4.2, “INSERT DELAYED Syntax”](#).

If you are using the update log or binary log, concurrent inserts are converted to normal inserts for [CREATE ... SELECT](#) or [INSERT ... SELECT](#) statements. This is done to ensure that you can re-create an exact copy of your tables by applying the log during a backup operation.

With [LOAD DATA INFILE](#), if you specify [CONCURRENT](#) with a [MyISAM](#) table that satisfies the condition for concurrent inserts (that is, it contains no free blocks in the middle), other threads can retrieve data from the table while [LOAD DATA](#) is executing. Using this option affects the performance of [LOAD DATA](#) a bit, even if no other thread is using the table at the same time.

## 7.4. Optimizing Database Structure

### 7.4.1. Design Choices

MySQL keeps row data and index data in separate files. Many (almost all) other database systems mix row and index data in the same file. We believe that the MySQL choice is better for a very wide range of modern systems.

Another way to store the row data is to keep the information for each column in a separate area (examples are [SDBM](#) and [Focus](#)). This causes a performance hit for every query that accesses more than one column. Because this degenerates so quickly when more than one column is accessed, we believe that this model is not good for general-purpose databases.

The more common case is that the index and data are stored together (as in [Oracle/Sybase](#), et al). In this case, you find the row information at the leaf page of the index. The good thing with this layout is that it, in many cases, depending on how well the index is cached, saves a disk read. The bad things with this layout are:

- Table scanning is much slower because you have to read through the indexes to get at the data.
- You cannot use only the index table to retrieve data for a query.
- You use more space because you must duplicate indexes from the nodes (you cannot store the row in the nodes).
- Deletes degenerate the table over time (because indexes in nodes are usually not updated on delete).
- It is more difficult to cache only the index data.

### 7.4.2. Make Your Data as Small as Possible

One of the most basic optimizations is to design your tables to take as little space on the disk as possible. This can result in huge improvements because disk reads are faster, and smaller tables normally require less main memory while their contents are being actively processed during query execution. Indexing also is a lesser resource burden if done on smaller columns.

MySQL supports many different storage engines (table types) and row formats. For each table, you can decide which storage and indexing method to use. Choosing the proper table format for your application may give you a big performance gain. See [Chapter 14, \*Storage Engines and Table Types\*](#).

You can get better performance for a table and minimize storage space by using the techniques listed here:

- Use the most efficient (smallest) data types possible. MySQL has many specialized types that save disk space and memory. For example, use the smaller integer types if possible to get smaller tables. `MEDIUMINT` is often a better choice than `INT` because a `MEDIUMINT` column uses 25% less space.
- Declare columns to be `NOT NULL` if possible. It makes everything faster and you save one bit per column. If you really need `NULL` in your application, you should definitely use it. Just avoid having it on all columns by default.
- For `MyISAM` tables, if you do not have any variable-length columns (`VARCHAR`, `TEXT`, or `BLOB` columns), a fixed-size row format is used. This is faster but unfortunately may waste some space. See [Section 14.1.3, “MyISAM Table Storage Formats”](#). You can hint that you want to have fixed length rows even if you have `VARCHAR` columns with the `CREATE TABLE` option `ROW_FORMAT=FIXED`.
- The primary index of a table should be as short as possible. This makes identification of each row easy and efficient.
- Create only the indexes that you really need. Indexes are good for retrieval but bad when you need to store data quickly. If you access a table mostly by searching on a combination of columns, create an index on them. The first part of the index should be the column most used. If you *always* use many columns when selecting from the table, you should use the column with more duplicates first to obtain better compression of the index.
- If it is very likely that a string column has a unique prefix on the first number of characters, it's better to index only this prefix, using MySQL's support for creating an index on the leftmost part of the column (see [Section 13.1.4, “CREATE INDEX Syntax”](#)). Shorter indexes are faster, not only because they require less disk space, but because they give also you more hits in the index cache, and thus fewer disk seeks. See [Section 7.5.2, “Tuning Server Parameters”](#).
- In some circumstances, it can be beneficial to split into two a table that is scanned very often. This is especially true if it is a dynamic-format table and it is possible to use a smaller static format table that can be used to find the relevant rows when scanning the table.

### 7.4.3. Column Indexes

All MySQL data types can be indexed. Use of indexes on the relevant columns is the best way to improve the performance of `SELECT` operations.

The maximum number of indexes per table and the maximum index length is defined per storage engine. See [Chapter 14, \*Storage Engines and Table Types\*](#). All storage engines support at least 16 indexes per table and a total index length of at least 256 bytes. Most storage engines have higher limits.

The `MyISAM` and (as of MySQL 4.0.14) `InnoDB` storage engines also support indexing on `BLOB` and `TEXT` columns. When indexing a `BLOB` or `TEXT` column, you *must* specify a prefix length for the index.

For example:

```
CREATE TABLE test (blob_col BLOB, INDEX(blob_col(10)));
```

Prefixes can be up to 1000 bytes long (767 bytes for [InnoDB](#) tables). (Before MySQL 4.1.2, the limit is 255 bytes for all tables.) Note that prefix limits are measured in bytes, whereas the prefix length in [CREATE TABLE](#) statements is interpreted as number of characters. *Be sure to take this into account when specifying a prefix length for a column that uses a multi-byte character set.*

As of MySQL 3.23.23, you can also create [FULLTEXT](#) indexes. They are used for full-text searches. Only the [MyISAM](#) storage engine supports [FULLTEXT](#) indexes and only for [CHAR](#), [VARCHAR](#), and [TEXT](#) columns. Indexing always takes place over the entire column and partial (column prefix) indexing is not supported. For details, see [Section 12.7, “Full-Text Search Functions”](#).

As of MySQL 4.1.0, you can create indexes on spatial data types. Spatial indexes use R-trees. Currently, only [MyISAM](#) supports indexes on spatial types.

The [MEMORY \(HEAP\)](#) storage engine uses [HASH](#) indexes by default. It also supports [BTREE](#) indexes as of MySQL 4.1.0.

## 7.4.4. Multiple-Column Indexes

MySQL can create composite indexes (that is, indexes on multiple columns). An index may consist of up to 15 columns. For certain data types, you can index a prefix of the column (see [Section 7.4.3, “Column Indexes”](#)).

A multiple-column index can be considered a sorted array containing values that are created by concatenating the values of the indexed columns.

MySQL uses multiple-column indexes in such a way that queries are fast when you specify a known quantity for the first column of the index in a [WHERE](#) clause, even if you do not specify values for the other columns.

Suppose that a table has the following specification:

```
CREATE TABLE test (
  id          INT NOT NULL,
  last_name   CHAR(30) NOT NULL,
  first_name  CHAR(30) NOT NULL,
  PRIMARY KEY (id),
  INDEX name  (last_name,first_name)
);
```

The [name](#) index is an index over the [last\\_name](#) and [first\\_name](#) columns. The index can be used for queries that specify values in a known range for [last\\_name](#), or for both [last\\_name](#) and [first\\_name](#). Therefore, the [name](#) index is used in the following queries:

```
SELECT * FROM test WHERE last_name='Widenius';
SELECT * FROM test
  WHERE last_name='Widenius' AND first_name='Michael';
SELECT * FROM test
  WHERE last_name='Widenius'
  AND (first_name='Michael' OR first_name='Monty');
SELECT * FROM test
  WHERE last_name='Widenius'
  AND first_name >='M' AND first_name < 'N';
```

However, the [name](#) index is *not* used in the following queries:

```
SELECT * FROM test WHERE first_name='Michael';
```

```
SELECT * FROM test
WHERE last_name='Widenius' OR first_name='Michael';
```

The manner in which MySQL uses indexes to improve query performance is discussed further in [Section 7.4.5, “How MySQL Uses Indexes”](#).

## 7.4.5. How MySQL Uses Indexes

Indexes are used to find rows with specific column values quickly. Without an index, MySQL must begin with the first row and then read through the entire table to find the relevant rows. The larger the table, the more this costs. If the table has an index for the columns in question, MySQL can quickly determine the position to seek to in the middle of the data file without having to look at all the data. If a table has 1,000 rows, this is at least 100 times faster than reading sequentially. If you need to access most of the rows, it is faster to read sequentially, because this minimizes disk seeks.

Most MySQL indexes ([PRIMARY KEY](#), [UNIQUE](#), [INDEX](#), and [FULLTEXT](#)) are stored in B-trees. Exceptions are that indexes on spatial data types use R-trees, and that [MEMORY \(HEAP\)](#) tables support hash indexes.

Strings are automatically prefix- and end-space compressed. See [Section 13.1.4, “CREATE INDEX Syntax”](#).

In general, indexes are used as described in the following discussion. Characteristics specific to hash indexes (as used in [MEMORY](#) tables) are described at the end of this section.

MySQL uses indexes for these operations:

- To find the rows matching a [WHERE](#) clause quickly.
- To eliminate rows from consideration. If there is a choice between multiple indexes, MySQL normally uses the index that finds the smallest number of rows.
- To retrieve rows from other tables when performing joins.
- To find the [MIN\(\)](#) or [MAX\(\)](#) value for a specific indexed column *key\_col*. This is optimized by a preprocessor that checks whether you are using [WHERE key\\_part\\_N = constant](#) on all key parts that occur before *key\_col* in the index. In this case, MySQL does a single key lookup for each [MIN\(\)](#) or [MAX\(\)](#) expression and replaces it with a constant. If all expressions are replaced with constants, the query returns at once. For example:

```
SELECT MIN(key_part2),MAX(key_part2)
FROM tbl_name WHERE key_part1=10;
```

- To sort or group a table if the sorting or grouping is done on a leftmost prefix of a usable key (for example, [ORDER BY key\\_part1, key\\_part2](#)). If all key parts are followed by [DESC](#), the key is read in reverse order. See [Section 7.2.9, “ORDER BY Optimization”](#).
- In some cases, a query can be optimized to retrieve values without consulting the data rows. If a query uses only columns from a table that are numeric and that form a leftmost prefix for some key, the selected values may be retrieved from the index tree for greater speed:

```
SELECT key_part3 FROM tbl_name
WHERE key_part1=1
```

Suppose that you issue the following [SELECT](#) statement:

```
mysql> SELECT * FROM tbl_name WHERE col1=val1 AND col2=val2;
```

If a multiple-column index exists on `col1` and `col2`, the appropriate rows can be fetched directly. If separate single-column indexes exist on `col1` and `col2`, the optimizer tries to find the most restrictive index by deciding which index finds fewer rows and using that index to fetch the rows.

If the table has a multiple-column index, any leftmost prefix of the index can be used by the optimizer to find rows. For example, if you have a three-column index on (`col1`, `col2`, `col3`), you have indexed search capabilities on (`col1`), (`col1`, `col2`), and (`col1`, `col2`, `col3`).

MySQL cannot use a partial index if the columns do not form a leftmost prefix of the index. Suppose that you have the `SELECT` statements shown here:

```
SELECT * FROM tbl_name WHERE col1=val1;
SELECT * FROM tbl_name WHERE col1=val1 AND col2=val2;

SELECT * FROM tbl_name WHERE col2=val2;
SELECT * FROM tbl_name WHERE col2=val2 AND col3=val3;
```

If an index exists on (`col1`, `col2`, `col3`), only the first two queries use the index. The third and fourth queries do involve indexed columns, but (`col2`) and (`col2`, `col3`) are not leftmost prefixes of (`col1`, `col2`, `col3`).

A B-tree index can be used for column comparisons in expressions that use the `=`, `>`, `>=`, `<`, `<=`, or `BETWEEN` operators. The index also can be used for `LIKE` comparisons if the argument to `LIKE` is a constant string that does not start with a wildcard character. For example, the following `SELECT` statements use indexes:

```
SELECT * FROM tbl_name WHERE key_col LIKE 'Patrick%';
SELECT * FROM tbl_name WHERE key_col LIKE 'Pat%ck%';
```

In the first statement, only rows with `'Patrick' <= key_col < 'Patricl'` are considered. In the second statement, only rows with `'Pat' <= key_col < 'Pau'` are considered.

The following `SELECT` statements do not use indexes:

```
SELECT * FROM tbl_name WHERE key_col LIKE '%Patrick%';
SELECT * FROM tbl_name WHERE key_col LIKE other_col;
```

In the first statement, the `LIKE` value begins with a wildcard character. In the second statement, the `LIKE` value is not a constant.

MySQL 4.0 and later versions perform an additional `LIKE` optimization. If you use `... LIKE 'string%'` and `string` is longer than three characters, MySQL uses the *Turbo Boyer-Moore algorithm* to initialize the pattern for the string and then uses this pattern to perform the search more quickly.

A search using `col_name IS NULL` employs indexes if `col_name` is indexed.

Any index that does not span all `AND` levels in the `WHERE` clause is not used to optimize the query. In other words, to be able to use an index, a prefix of the index must be used in every `AND` group.

The following `WHERE` clauses use indexes:

```
... WHERE index_part1=1 AND index_part2=2 AND other_column=3
/* index = 1 OR index = 2 */
... WHERE index=1 OR A=10 AND index=2
/* optimized like "index_part1='hello'" */
... WHERE index_part1='hello' AND index_part3=5
/* Can use index on index1 but not on index2 or index3 */
... WHERE index1=1 AND index2=2 OR index1=3 AND index3=3;
```

These [WHERE](#) clauses do *not* use indexes:

```

/* index_part1 is not used */
... WHERE index_part2=1 AND index_part3=2

/* Index is not used in both parts of the WHERE clause */
... WHERE index=1 OR A=10

/* No index spans all rows */
... WHERE index_part1=1 OR index_part2=10

```

Sometimes MySQL does not use an index, even if one is available. One circumstance under which this occurs is when the optimizer estimates that using the index would require MySQL to access a very large percentage of the rows in the table. (In this case, a table scan is likely to be much faster because it requires fewer seeks.) However, if such a query uses [LIMIT](#) to retrieve only some of the rows, MySQL uses an index anyway, because it can much more quickly find the few rows to return in the result.

Hash indexes have somewhat different characteristics from those just discussed:

- They are used only for equality comparisons that use the = or <=> operators (but are *very* fast). They are not used for comparison operators such as < that find a range of values.
- The optimizer cannot use a hash index to speed up [ORDER BY](#) operations. (This type of index cannot be used to search for the next entry in order.)
- MySQL cannot determine approximately how many rows there are between two values (this is used by the range optimizer to decide which index to use). This may affect some queries if you change a [MyISAM](#) table to a hash-indexed [MEMORY](#) table.
- Only whole keys can be used to search for a row. (With a B-tree index, any leftmost prefix of the key can be used to find rows.)

## 7.4.6. The [MyISAM](#) Key Cache

To minimize disk I/O, the [MyISAM](#) storage engine exploits a strategy that is used by many database management systems. It employs a cache mechanism to keep the most frequently accessed table blocks in memory:

- For index blocks, a special structure called the *key cache* (or *key buffer*) is maintained. The structure contains a number of block buffers where the most-used index blocks are placed.
- For data blocks, MySQL uses no special cache. Instead it relies on the native operating system filesystem cache.

This section first describes the basic operation of the [MyISAM](#) key cache. Then it discusses changes made in MySQL 4.1 that improve key cache performance and that enable you to better control cache operation:

- Access to the key cache no longer is serialized among threads. Multiple threads can access the cache concurrently.
- You can set up multiple key caches and assign table indexes to specific caches.

The key cache mechanism also is used for `ISAM` tables. However, the significance of this fact is on the wane. `ISAM` table use has been decreasing since MySQL 3.23 when `MyISAM` was introduced. MySQL 4.1 carries this trend further; the `ISAM` storage engine is disabled by default. (Subsequent MySQL release series have no support at all for `ISAM`.)

To control the size of the key cache, use the `key_buffer_size` system variable. If this variable is set equal to zero, no key cache is used. The key cache also is not used if the `key_buffer_size` value is too small to allocate the minimal number of block buffers (8).

When the key cache is not operational, index files are accessed using only the native filesystem buffering provided by the operating system. (In other words, table index blocks are accessed using the same strategy as that employed for table data blocks.)

An index block is a contiguous unit of access to the `MyISAM` index files. Usually the size of an index block is equal to the size of nodes of the index B-tree. (Indexes are represented on disk using a B-tree data structure. Nodes at the bottom of the tree are leaf nodes. Nodes above the leaf nodes are non-leaf nodes.)

All block buffers in a key cache structure are the same size. This size can be equal to, greater than, or less than the size of a table index block. Usually one these two values is a multiple of the other.

When data from any table index block must be accessed, the server first checks whether it is available in some block buffer of the key cache. If it is, the server accesses data in the key cache rather than on disk. That is, it reads from the cache or writes into it rather than reading from or writing to disk. Otherwise, the server chooses a cache block buffer containing a different table index block (or blocks) and replaces the data there by a copy of required table index block. As soon as the new index block is in the cache, the index data can be accessed.

If it happens that a block selected for replacement has been modified, the block is considered “dirty.” In this case, prior to being replaced, its contents are flushed to the table index from which it came.

Usually the server follows an *LRU (Least Recently Used)* strategy: When choosing a block for replacement, it selects the least recently used index block. To make this choice easier, the key cache module maintains a special queue (*LRU chain*) of all used blocks. When a block is accessed, it is placed at the end of the queue. When blocks need to be replaced, blocks at the beginning of the queue are the least recently used and become the first candidates for eviction.

### 7.4.6.1. Shared Key Cache Access

Prior to MySQL 4.1, access to the key cache is serialized: No two threads can access key cache buffers simultaneously. The server processes a request for an index block only after it has finished processing the previous request. As a result, a request for an index block not present in any key cache buffer blocks access by other threads while a buffer is being updated to contain the requested index block.

Starting from version 4.1.0, the server supports shared access to the key cache:

- A buffer that is not being updated can be accessed by multiple threads.
- A buffer that is being updated causes threads that need to use it to wait until the update is complete.
- Multiple threads can initiate requests that result in cache block replacements, as long as they do not interfere with each other (that is, as long as they need different index blocks, and thus cause different cache blocks to be replaced).

Shared access to the key cache enables the server to improve throughput significantly.



## 7.4.6.2. Multiple Key Caches

Shared access to the key cache improves performance but does not eliminate contention among threads entirely. They still compete for control structures that manage access to the key cache buffers. To reduce key cache access contention further, MySQL 4.1.1 also provides multiple key caches. This feature enables you to assign different table indexes to different key caches.

Where there are multiple key caches, the server must know which cache to use when processing queries for a given `MyISAM` table. By default, all `MyISAM` table indexes are cached in the default key cache. To assign table indexes to a specific key cache, use the `CACHE INDEX` statement (see [Section 13.5.5.1, “CACHE INDEX Syntax”](#)). For example, the following statement assigns indexes from the tables `t1`, `t2`, and `t3` to the key cache named `hot_cache`:

```
mysql> CACHE INDEX t1, t2, t3 IN hot_cache;
```

Table	Op	Msg_type	Msg_text
test.t1	assign_to_keycache	status	OK
test.t2	assign_to_keycache	status	OK
test.t3	assign_to_keycache	status	OK

**Note:** If the server has been built with the `ISAM` storage engine enabled, `ISAM` tables use the key cache mechanism. However, `ISAM` indexes use only the default key cache and cannot be reassigned to a different cache.

The key cache referred to in a `CACHE INDEX` statement can be created by setting its size with a `SET GLOBAL` parameter setting statement or by using server startup options. For example:

```
mysql> SET GLOBAL keycache1.key_buffer_size=128*1024;
```

To destroy a key cache, set its size to zero:

```
mysql> SET GLOBAL keycache1.key_buffer_size=0;
```

Note that you cannot destroy the default key cache. Any attempt to do this will be ignored:

```
mysql> SET GLOBAL key_buffer_size = 0;
```

```
mysql> SHOW VARIABLES LIKE 'key_buffer_size';
```

Variable_name	Value
key_buffer_size	8384512

Key cache variables are structured system variables that have a name and components. For `keycache1.key_buffer_size`, `keycache1` is the cache variable name and `key_buffer_size` is the cache component. See [Section 5.2.3.1, “Structured System Variables”](#), for a description of the syntax used for referring to structured key cache system variables.

By default, table indexes are assigned to the main (default) key cache created at the server startup. When a key cache is destroyed, all indexes assigned to it are reassigned to the default key cache.

For a busy server, we recommend a strategy that uses three key caches:

- A “hot” key cache that takes up 20% of the space allocated for all key caches. Use this for tables that are heavily used for searches but that are not updated.
- A “cold” key cache that takes up 20% of the space allocated for all key caches. Use this cache for

medium-sized, intensively modified tables, such as temporary tables.

- A “warm” key cache that takes up 60% of the key cache space. Employ this as the default key cache, to be used by default for all other tables.

One reason the use of three key caches is beneficial is that access to one key cache structure does not block access to the others. Statements that access tables assigned to one cache do not compete with statements that access tables assigned to another cache. Performance gains occur for other reasons as well:

- The hot cache is used only for retrieval queries, so its contents are never modified. Consequently, whenever an index block needs to be pulled in from disk, the contents of the cache block chosen for replacement need not be flushed first.
- For an index assigned to the hot cache, if there are no queries requiring an index scan, there is a high probability that the index blocks corresponding to non-leaf nodes of the index B-tree remain in the cache.
- An update operation most frequently executed for temporary tables is performed much faster when the updated node is in the cache and need not be read in from disk first. If the size of the indexes of the temporary tables are comparable with the size of cold key cache, the probability is very high that the updated node is in the cache.

`CACHE INDEX` sets up an association between a table and a key cache, but the association is lost each time the server restarts. If you want the association to take effect each time the server starts, one way to accomplish this is to use an option file: Include variable settings that configure your key caches, and an `init-file` option that names a file containing `CACHE INDEX` statements to be executed. For example:

```
key_buffer_size = 4G
hot_cache.key_buffer_size = 2G
cold_cache.key_buffer_size = 2G
init_file=/path/to/data-directory/mysql_d_init.sql
```

The statements in `mysql_d_init.sql` are executed each time the server starts. The file should contain one SQL statement per line. The following example assigns several tables each to `hot_cache` and `cold_cache`:

```
CACHE INDEX db1.t1, db1.t2, db2.t3 IN hot_cache
CACHE INDEX db1.t4, db2.t5, db2.t6 IN cold_cache
```

### 7.4.6.3. Midpoint Insertion Strategy

By default, the key cache management system of MySQL 4.1 uses the LRU strategy for choosing key cache blocks to be evicted, but it also supports a more sophisticated method called the *midpoint insertion strategy*.

When using the midpoint insertion strategy, the LRU chain is divided into two parts: a hot sub-chain and a warm sub-chain. The division point between two parts is not fixed, but the key cache management system takes care that the warm part is not “too short,” always containing at least `key_cache_division_limit` percent of the key cache blocks. `key_cache_division_limit` is a component of structured key cache variables, so its value is a parameter that can be set per cache.

When an index block is read from a table into the key cache, it is placed at the end of the warm sub-chain. After a certain number of hits (accesses of the block), it is promoted to the hot sub-chain. At

present, the number of hits required to promote a block (3) is the same for all index blocks.

A block promoted into the hot sub-chain is placed at the end of the chain. The block then circulates within this sub-chain. If the block stays at the beginning of the sub-chain for a long enough time, it is demoted to the warm chain. This time is determined by the value of the `key_cache_age_threshold` component of the key cache.

The threshold value prescribes that, for a key cache containing  $N$  blocks, the block at the beginning of the hot sub-chain not accessed within the last  $N \times \text{key\_cache\_age\_threshold} / 100$  hits is to be moved to the beginning of the warm sub-chain. It then becomes the first candidate for eviction, because blocks for replacement always are taken from the beginning of the warm sub-chain.

The midpoint insertion strategy allows you to keep more-valued blocks always in the cache. If you prefer to use the plain LRU strategy, leave the `key_cache_division_limit` value set to its default of 100.

The midpoint insertion strategy helps to improve performance when execution of a query that requires an index scan effectively pushes out of the cache all the index blocks corresponding to valuable high-level B-tree nodes. To avoid this, you must use a midpoint insertion strategy with the `key_cache_division_limit` set to much less than 100. Then valuable frequently hit nodes are preserved in the hot sub-chain during an index scan operation as well.

### 7.4.6.4. Index Preloading

If there are enough blocks in a key cache to hold blocks of an entire index, or at least the blocks corresponding to its non-leaf nodes, it makes sense to preload the key cache with index blocks before starting to use it. Preloading allows you to put the table index blocks into a key cache buffer in the most efficient way: by reading the index blocks from disk sequentially.

Without preloading, the blocks are still placed into the key cache as needed by queries. Although the blocks will stay in the cache, because there are enough buffers for all of them, they are fetched from disk in random order, and not sequentially.

To preload an index into a cache, use the `LOAD INDEX INTO CACHE` statement. For example, the following statement preloads nodes (index blocks) of indexes of the tables `t1` and `t2`:

```
mysql> LOAD INDEX INTO CACHE t1, t2 IGNORE LEAVES;
```

Table	Op	Msg_type	Msg_text
test.t1	preload_keys	status	OK
test.t2	preload_keys	status	OK

The `IGNORE LEAVES` modifier causes only blocks for the non-leaf nodes of the index to be preloaded. Thus, the statement shown preloads all index blocks from `t1`, but only blocks for the non-leaf nodes from `t2`.

If an index has been assigned to a key cache using a `CACHE INDEX` statement, preloading places index blocks into that cache. Otherwise, the index is loaded into the default key cache.

### 7.4.6.5. Key Cache Block Size

MySQL 4.1 introduces a new `key_cache_block_size` variable on a per-key cache basis. This variable specifies the size of the block buffers for a key cache. It is intended to allow tuning of the performance of I/O operations for index files.

The best performance for I/O operations is achieved when the size of read buffers is equal to the size of the native operating system I/O buffers. But setting the size of key nodes equal to the size of the I/O buf-

fer does not always ensure the best overall performance. When reading the big leaf nodes, the server pulls in a lot of unnecessary data, effectively preventing reading other leaf nodes.

Currently, you cannot control the size of the index blocks in a table. This size is set by the server when the `.MYI` index file is created, depending on the size of the keys in the indexes present in the table definition. In most cases, it is set equal to the I/O buffer size.

### 7.4.6.6. Restructuring a Key Cache

A key cache can be restructured at any time by updating its parameter values. For example:

```
mysql> SET GLOBAL cold_cache.key_buffer_size=4*1024*1024;
```

If you assign to either the `key_buffer_size` or `key_cache_block_size` key cache component a value that differs from the component's current value, the server destroys the cache's old structure and creates a new one based on the new values. If the cache contains any dirty blocks, the server saves them to disk before destroying and re-creating the cache. Restructuring does not occur if you change other key cache parameters.

When restructuring a key cache, the server first flushes the contents of any dirty buffers to disk. After that, the cache contents become unavailable. However, restructuring does not block queries that need to use indexes assigned to the cache. Instead, the server directly accesses the table indexes using native filesystem caching. Filesystem caching is not as efficient as using a key cache, so although queries execute, a slowdown can be anticipated. After the cache has been restructured, it becomes available again for caching indexes assigned to it, and the use of filesystem caching for the indexes ceases.

### 7.4.7. MyISAM Index Statistics Collection

Storage engines collect statistics about tables for use by the optimizer. Table statistics are based on value groups, where a value group is a set of rows with the same key prefix value. For optimizer purposes, an important statistic is the average value group size.

MySQL uses the average value group size in the following ways:

- To estimate how many rows must be read for each `ref` access
- To estimate how many row a partial join will produce; that is, the number of rows that an operation of this form will produce:

```
(...) JOIN tbl_name ON tbl_name.key = expr
```

As the average value group size for an index increases, the index is less useful for those two purposes because the average number of rows per lookup increases: For the index to be good for optimization purposes, it is best that each index value target a small number of rows in the table. When a given index value yields a large number of rows, the index is less useful and MySQL is less likely to use it.

The average value group size is related to table cardinality, which is the number of value groups. The `SHOW INDEX` statement displays a cardinality value based on  $N/S$ , where  $N$  is the number of rows in the table and  $S$  is the average value group size. That ratio yields an approximate number of value groups in the table.

For a join based on the `<=>` comparison operator, `NULL` is not treated differently from any other value: `NULL <=> NULL`, just as `N <=> N` for any other  $N$ .

However, for a join based on the `=` operator, `NULL` is different from non-`NULL` values: `expr1 = ex-`

*pr2* is not true when *expr1* or *expr2* (or both) are `NULL`. This affects `ref` accesses for comparisons of the form `tbl_name.key = expr`: MySQL will not access the table if the current value of *expr* is `NULL`, because the comparison cannot be true.

For `=` comparisons, it does not matter how many `NULL` values are in the table. For optimization purposes, the relevant value is the average size of the non-`NULL` value groups. However, MySQL does not currently allow that average size to be collected or used.

For `MyISAM` tables, you have some control over collection of table statistics by means of the `myisam_stats_method` system variable. This variable has two possible values, which differ as follows:

- When `myisam_stats_method` is `nulls_equal`, all `NULL` values are treated as identical (that is, they all form a single value group).

If the `NULL` value group size is much higher than the average non-`NULL` value group size, this method skews the average value group size upward. This makes index appear to the optimizer to be less useful than it really is for joins that look for non-`NULL` values. Consequently, the `nulls_equal` method may cause the optimizer not to use the index for `ref` accesses when it should.

- When `myisam_stats_method` is `nulls_unequal`, `NULL` values are not considered the same. Instead, each `NULL` value forms a separate value group of size 1.

If you have many `NULL` values, this method skews the average value group size downward. If the average non-`NULL` value group size is large, counting `NULL` values each as a group of size 1 causes the optimizer to overestimate the value of the index for joins that look for non-`NULL` values. Consequently, the `nulls_unequal` method may cause the optimizer to use this index for `ref` lookups when other methods may be better.

If you tend to use many joins that use `<=>` rather than `=`, `NULL` values are not special in comparisons and one `NULL` is equal to another. In this case, `nulls_equal` is the appropriate statistics method.

The `myisam_stats_method` system variable has global and session values. Setting the global value affects `MyISAM` statistics collection for all `MyISAM` tables. Setting the session value affects statistics collection only for the current client connection. This means that you can force a table's statistics to be regenerated with a given method without affecting other clients by setting the session value of `myisam_stats_method`.

To regenerate table statistics, you can use any of the following methods:

- Set `myisam_stats_method`, and then issue a `CHECK TABLE` statement
- Execute `myisamchk --stats_method=method_name --analyze`
- Change the table to cause its statistics to go out of date (for example, insert a row and then delete it), and then set `myisam_stats_method` and issue an `ANALYZE TABLE` statement

Some caveats regarding the use of `myisam_stats_method`:

- You can force table statistics to be collected explicitly, as just described. However, MySQL may also collect statistics automatically. For example, if during the course of executing statements for a table, some of those statements modify the table, MySQL may collect statistics. (This may occur for bulk inserts or deletes, or some `ALTER TABLE` statements, for example.) If this happens, the statistics are collected using whatever value `myisam_stats_method` has at the time. Thus, if you collect statistics using one method, but `myisam_stats_method` is set to the other method when a table's statistics are collected automatically later, the other method will be used.

- There is no way to tell which method was used to generate statistics for a given `MyISAM` table.
- `myisam_stats_method` applies only to `MyISAM` tables. Other storage engines have only one method for collecting table statistics. Usually it is closer to the `nulls_equal` method.

## 7.4.8. How MySQL Opens and Closes Tables

When you execute a `mysqladmin status` command, you should see something like this:

```
Uptime: 426 Running threads: 1 Questions: 11082
Reloads: 1 Open tables: 12
```

The `Open tables` value of 12 can be somewhat puzzling if you have only six tables.

MySQL is multi-threaded, so there may be many clients issuing queries for a given table simultaneously. To minimize the problem with multiple client threads having different states on the same table, the table is opened independently by each concurrent thread. This uses additional memory but normally increases performance. With `MyISAM` tables, one extra file descriptor is required for the data file for each client that has the table open. (By contrast, the index file descriptor is shared between all threads.)

The `table_cache`, `max_connections`, and `max_tmp_tables` system variables affect the maximum number of files the server keeps open. If you increase one or more of these values, you may run up against a limit imposed by your operating system on the per-process number of open file descriptors. Many operating systems allow you to increase the open-files limit, although the method varies widely from system to system. Consult your operating system documentation to determine whether it is possible to increase the limit and how to do so.

`table_cache` is related to `max_connections`. For example, for 200 concurrent running connections, you should have a table cache size of at least  $200 \times N$ , where  $N$  is the maximum number of tables per join in any of the queries which you execute. You must also reserve some extra file descriptors for temporary tables and files.

Make sure that your operating system can handle the number of open file descriptors implied by the `table_cache` setting. If `table_cache` is set too high, MySQL may run out of file descriptors and refuse connections, fail to perform queries, and be very unreliable. You also have to take into account that the `MyISAM` storage engine needs two file descriptors for each unique open table. You can increase the number of file descriptors available to MySQL using the `--open-files-limit` startup option to `mysqld`. See [Section A.2.17, “File Not Found”](#).

The cache of open tables is kept at a level of `table_cache` entries. The default value is 64; this can be changed with the `--table_cache` option to `mysqld`. Note that MySQL may temporarily open more tables than this to execute queries.

MySQL closes an unused table and removes it from the table cache under the following circumstances:

- When the cache is full and a thread tries to open a table that is not in the cache.
- When the cache contains more than `table_cache` entries and a table in the cache is no longer being used by any threads.
- When a table flushing operation occurs. This happens when someone issues a `FLUSH TABLES` statement or executes a `mysqladmin flush-tables` or `mysqladmin refresh` command.

When the table cache fills up, the server uses the following procedure to locate a cache entry to use:

- Tables that are not currently in use are released, beginning with the table least recently used.
- If a new table needs to be opened, but the cache is full and no tables can be released, the cache is temporarily extended as necessary.

When the cache is in a temporarily extended state and a table goes from a used to unused state, the table is closed and released from the cache.

A table is opened for each concurrent access. This means the table needs to be opened twice if two threads access the same table or if a thread accesses the table twice in the same query (for example, by joining the table to itself). Each concurrent open requires an entry in the table cache. The first open of any `MyISAM` table takes two file descriptors: one for the data file and one for the index file. Each additional use of the table takes only one file descriptor for the data file. The index file descriptor is shared among all threads.

If you are opening a table with the `HANDLER tbl_name OPEN` statement, a dedicated table object is allocated for the thread. This table object is not shared by other threads and is not closed until the thread calls `HANDLER tbl_name CLOSE` or the thread terminates. When this happens, the table is put back in the table cache (if the cache is not full). See [Section 13.2.3, “HANDLER Syntax”](#).

You can determine whether your table cache is too small by checking the `mysqld` status variable `Opened_tables`:

```
mysql> SHOW STATUS LIKE 'Opened_tables';
+-----+-----+
| Variable_name | Value |
+-----+-----+
| Opened_tables | 2741  |
+-----+-----+
```

If the value is very large, even when you have not issued many `FLUSH TABLES` statements, you should increase the table cache size. See [Section 5.2.2, “System Variables”](#), and [Section 5.2.4, “Status Variables”](#).

## 7.4.9. Drawbacks to Creating Many Tables in the Same Database

If you have many `MyISAM` or `ISAM` tables in a database directory, open, close, and create operations are slow. If you execute `SELECT` statements on many different tables, there is a little overhead when the table cache is full, because for every table that has to be opened, another must be closed. You can reduce this overhead by making the table cache larger.

## 7.5. Optimizing the MySQL Server

### 7.5.1. System Factors and Startup Parameter Tuning

We start with system-level factors, because some of these decisions must be made very early to achieve large performance gains. In other cases, a quick look at this section may suffice. However, it is always nice to have a sense of how much can be gained by changing factors that apply at this level.

The operating system to use is very important. To get the best use of multiple-CPU machines, you should use Solaris (because its threads implementation works well) or Linux (because the 2.4 and later kernels have good SMP support). Note that older Linux kernels have a 2GB filesize limit by default. If you have such a kernel and a need for files larger than 2GB, you should get the Large File Support (LFS) patch for the ext2 filesystem. Other filesystems such as ReiserFS and XFS do not have this 2GB limitation.

Before using MySQL in production, we advise you to test it on your intended platform.

Other tips:

- If you have enough RAM, you could remove all swap devices. Some operating systems use a swap device in some contexts even if you have free memory.
- Use the `--skip-external-locking` MySQL option to avoid external locking. This option is turned on by default as of MySQL 4.0. Before that, it is on by default when compiling with MIT-pthreads, because `flock()` is not fully supported by MIT-pthreads on all platforms. It is also on by default for Linux because Linux file locking is not yet safe.

Note that disabling external locking does not affect MySQL's functionality as long as you run only one server. Just remember to take down the server (or lock and flush the relevant tables) before you run `myisamchk`. On some systems it is mandatory to disable external locking because it does not work, anyway.

The only case in which you cannot disable external locking is when you run multiple MySQL *servers* (not clients) on the same data, or if you run `myisamchk` to check (not repair) a table without telling the server to flush and lock the tables first. Note that using multiple MySQL servers to access the same data concurrently is generally *not* recommended, except when using MySQL Cluster.

The `LOCK TABLES` and `UNLOCK TABLES` statements use internal locking, so you can use them even if external locking is disabled.

## 7.5.2. Tuning Server Parameters

You can determine the default buffer sizes used by the `mysqld` server with this command (prior to MySQL 4.1, omit `--verbose`):

```
shell> mysqld --verbose --help
```

This command produces a list of all `mysqld` options and configurable system variables. The output includes the default variable values and looks something like this:

```
back_log                current value: 5
bdb_cache_size          current value: 1048540
binlog_cache_size      current value: 32768
connect_timeout        current value: 5
delayed_insert_limit   current value: 100
delayed_insert_timeout current value: 300
delayed_queue_size     current value: 1000
flush_time             current value: 0
interactive_timeout    current value: 28800
join_buffer_size       current value: 131072
key_buffer_size        current value: 1048540
long_query_time        current value: 10
lower_case_table_names current value: 0
max_allowed_packet     current value: 1048576
max_binlog_cache_size  current value: 4294967295
max_connect_errors     current value: 10
max_connections        current value: 100
max_delayed_threads    current value: 20
max_heap_table_size    current value: 16777216
max_join_size          current value: 4294967295
max_sort_length        current value: 1024
max_tmp_tables         current value: 32
max_write_lock_count   current value: 4294967295
myisam_sort_buffer_size current value: 8388608
net_buffer_length      current value: 16384
net_read_timeout       current value: 30
net_retry_count        current value: 10
net_write_timeout      current value: 60
read_buffer_size       current value: 131072
```



```
read_rnd_buffer_size    current value: 262144
slow_launch_time       current value: 2
sort_buffer             current value: 2097116
table_cache            current value: 64
thread_concurrency     current value: 10
thread_stack           current value: 131072
tmp_table_size         current value: 1048576
wait_timeout           current value: 28800
```

For a `mysqld` server that is currently running, you can see the current values of its system variables by connecting to it and issuing this statement:

```
mysql> SHOW VARIABLES;
```

You can also see some statistical and status indicators for a running server by issuing this statement:

```
mysql> SHOW STATUS;
```

System variable and status information also can be obtained using `mysqladmin`:

```
shell> mysqladmin variables
shell> mysqladmin extended-status
```

For a full description of all system and status variables, see [Section 5.2.2, “System Variables”](#), and [Section 5.2.4, “Status Variables”](#).

MySQL uses algorithms that are very scalable, so you can usually run with very little memory. However, normally you get better performance by giving MySQL more memory.

When tuning a MySQL server, the two most important variables to configure are `key_buffer_size` and `table_cache`. You should first feel confident that you have these set appropriately before trying to change any other variables.

The following examples indicate some typical variable values for different runtime configurations. The examples use the `mysqld_safe` script and use `--var_name=value` syntax to set the variable `var_name` to the value `value`. This syntax is available as of MySQL 4.0. For older versions of MySQL, take the following differences into account:

- Use `safe_mysqld` rather than `mysqld_safe`.
- Set variables using `--set-variable=var_name=value` or `-O var_name=value` syntax.
- For variable names that end in `_size`, you may need to specify them without `_size`. For example, the old name for `sort_buffer_size` is `sort_buffer`. The old name for `read_buffer_size` is `record_buffer`. To see which variables your version of the server recognizes, use `mysqld --help`.

If you have at least 256MB of memory and many tables and want maximum performance with a moderate number of clients, you should use something like this:

```
shell> mysqld_safe --key_buffer_size=64M --table_cache=256 \
               --sort_buffer_size=4M --read_buffer_size=1M &
```

If you have only 128MB of memory and only a few tables, but you still do a lot of sorting, you can use something like this:

```
shell> mysqld_safe --key_buffer_size=16M --sort_buffer_size=1M
```

If there are very many simultaneous connections, swapping problems may occur unless `mysqld` has been configured to use very little memory for each connection. `mysqld` performs better if you have enough memory for all connections.

With little memory and lots of connections, use something like this:

```
shell> mysqld_safe --key_buffer_size=512K --sort_buffer_size=100K \  
--read_buffer_size=100K &
```

Or even this:

```
shell> mysqld_safe --key_buffer_size=512K --sort_buffer_size=16K \  
--table_cache=32 --read_buffer_size=8K \  
--net_buffer_length=1K &
```

If you are performing `GROUP BY` or `ORDER BY` operations on tables that are much larger than your available memory, you should increase the value of `read_rnd_buffer_size` to speed up the reading of rows following sorting operations.

You can make use of the example option files included with your MySQL distribution; see [Section 4.3.2.1, “Preconfigured Option Files”](#).

If you specify an option on the command line for `mysqld` or `mysqld_safe`, it remains in effect only for that invocation of the server. To use the option every time the server runs, put it in an option file.

To see the effects of a parameter change, do something like this (prior to MySQL 4.1, omit `-verbose`):

```
shell> mysqld --key_buffer_size=32M --verbose --help
```

The variable values are listed near the end of the output. Make sure that the `--verbose` and `--help` options are last. Otherwise, the effect of any options listed after them on the command line are not reflected in the output.

For information on tuning the `InnoDB` storage engine, see [Section 14.2.12, “InnoDB Performance Tuning Tips”](#).

## 7.5.3. How Compiling and Linking Affects the Speed of MySQL

Most of the following tests were performed on Linux with the MySQL benchmarks, but they should give some indication for other operating systems and workloads.

You obtain the fastest executables when you link with `-static`.

On Linux, it is best to compile the server with `pgcc` and `-O3`. You need about 200MB memory to compile `sql_yacc.cc` with these options, because `gcc` or `pgcc` needs a great deal of memory to make all functions inline. You should also set `CXX=gcc` when configuring MySQL to avoid inclusion of the `libstdc++` library, which is not needed. Note that with some versions of `pgcc`, the resulting binary runs only on true Pentium processors, even if you use the compiler option indicating that you want the resulting code to work on all x586-type processors (such as AMD).

By using a better compiler and compilation options, you can obtain a 10–30% speed increase in applications. This is particularly important if you compile the MySQL server yourself.

When we tested both the Cygnus CodeFusion and Fujitsu compilers, neither was sufficiently bug-free to allow MySQL to be compiled with optimizations enabled.

The standard MySQL binary distributions are compiled with support for all character sets. When you compile MySQL yourself, you should include support only for the character sets that you are going to use. This is controlled by the `--with-charset` option to `configure`.

Here is a list of some measurements that we have made:

- If you use `pgcc` and compile everything with `-O6`, the `mysqld` server is 1% faster than with `gcc 2.95.2`.
- If you link dynamically (without `-static`), the result is 13% slower on Linux. Note that you still can use a dynamically linked MySQL library for your client applications. It is the server that is most critical for performance.
- For a connection from a client to a server running on the same host, if you connect using TCP/IP rather than a Unix socket file, performance is 7.5% slower. (On Unix, if you connect to the hostname `localhost`, MySQL uses a socket file by default.)
- For TCP/IP connections from a client to a server, connecting to a remote server on another host is 8–11% slower than connecting to a server on the same host, even for connections over 100Mb/s Ethernet.
- When running our benchmark tests using secure connections (all data encrypted with internal SSL support) performance was 55% slower than with unencrypted connections.
- If you compile with `--with-debug=full`, most queries are 20% slower. Some queries may take substantially longer; for example, the MySQL benchmarks run 35% slower. If you use `--with-debug` (without `=full`), the speed decrease is only 15%. For a version of `mysqld` that has been compiled with `--with-debug=full`, you can disable memory checking at runtime by starting it with the `--skip-safemalloc` option. The execution speed should then be close to that obtained when configuring with `--with-debug`.
- On a Sun UltraSPARC-IIe, a server compiled with Forte 5.0 is 4% faster than one compiled with `gcc 3.2`.
- On a Sun UltraSPARC-IIe, a server compiled with Forte 5.0 is 4% faster in 32-bit mode than in 64-bit mode.
- Compiling with `gcc 2.95.2` for UltraSPARC with the `-mcpu=v8 -Wa, -xarch=v8plusa` options gives 4% more performance.
- On Solaris 2.5.1, MIT-pthreads is 8–12% slower than Solaris native threads on a single processor. With greater loads or more CPUs, the difference should be larger.
- Compiling on Linux-x86 using `gcc` without frame pointers (`-fomit-frame-pointer` or `-fomit-frame-pointer -ffixed-ebp`) makes `mysqld` 1–4% faster.

Binary MySQL distributions for Linux that are provided by MySQL AB used to be compiled with `pgcc`. We had to go back to regular `gcc` due to a bug in `pgcc` that would generate binaries that do not run on AMD. We will continue using `gcc` until that bug is resolved. In the meantime, if you have a non-AMD machine, you can build a faster binary by compiling with `pgcc`. The standard MySQL Linux binary is linked statically to make it faster and more portable.

## 7.5.4. How MySQL Uses Memory

The following list indicates some of the ways that the `mysqld` server uses memory. Where applicable, the name of the system variable relevant to the memory use is given:

- The key buffer is shared by all threads; its size is determined by the `key_buffer_size` variable. Other buffers used by the server are allocated as needed. See [Section 7.5.2, “Tuning Server Parameters”](#).
- Each connection uses some thread-specific space. The following list indicates these and which variables control their size:
  - A stack (default 64KB before MySQL 4.0.10 and 192KB thereafter, variable `thread_stack`)
  - A connection buffer (variable `net_buffer_length`)
  - A result buffer (variable `net_buffer_length`)

The connection buffer and result buffer both begin with a size given by `net_buffer_length` but are dynamically enlarged up to `max_allowed_packet` bytes as needed. The result buffer shrinks to `net_buffer_length` after each SQL statement. While a statement is running, a copy of the current statement string is also allocated.

- All threads share the same base memory.
- When a thread is no longer needed, the memory allocated to it is released and returned to the system unless the thread goes back into the thread cache. In that case, the memory remains allocated.
- Only compressed `ISAM` and `MyISAM` tables are memory mapped. This is because the 32-bit memory space of 4GB is not large enough for most big tables. When systems with a 64-bit address space become more common, we may add general support for memory mapping.
- Each request that performs a sequential scan of a table allocates a *read buffer* (variable `read_buffer_size`).
- When reading rows in an arbitrary sequence (for example, following a sort), a *random-read buffer* (variable `read_rnd_buffer_size`) may be allocated in order to avoid disk seeks.
- All joins are executed in a single pass, and most joins can be done without even using a temporary table. Most temporary tables are memory-based hash tables. Temporary tables with a large row length (calculated as the sum of all column lengths) or that contain `BLOB` columns are stored on disk.

One problem before MySQL 3.23.2 is that if an internal in-memory heap table exceeds the size of `tmp_table_size`, the error `The table tbl_name is full` occurs. From 3.23.2 on, this is handled automatically by changing the in-memory heap table to a disk-based `MyISAM` table as necessary. To work around this problem for older servers, you can increase the temporary table size by setting the `tmp_table_size` option to `mysqld`, or by setting the SQL option `SQL_BIG_TABLES` in the client program. See [Section 13.5.3, “SET Syntax”](#).

In MySQL 3.20, the maximum size of the temporary table is `record_buffer*16`; if you are using this version, you have to increase the value of `record_buffer`. You can also start `mysqld` with the `--big-tables` option to always store temporary tables on disk. However, this affects the speed of many complicated queries.

- Most requests that perform a sort allocate a sort buffer and zero to two temporary files depending on the result set size. See [Section A.4.4, “Where MySQL Stores Temporary Files”](#).
- Almost all parsing and calculating is done in a local memory store. No memory overhead is needed for small items, so the normal slow memory allocation and freeing is avoided. Memory is allocated only for unexpectedly large strings. This is done with `malloc()` and `free()`.
- For each `MyISAM` or `ISAM` table that is opened, the index file is opened once and the data file is opened once for each concurrently running thread. For each concurrent thread, a table structure,

column structures for each column, and a buffer of size  $3 \times N$  are allocated (where  $N$  is the maximum row length, not counting BLOB columns). A BLOB column requires five to eight bytes plus the length of the BLOB data. The MyISAM and ISAM storage engines maintain one extra row buffer for internal use.

- For each table having BLOB columns, a buffer is enlarged dynamically to read in larger BLOB values. If you scan a table, a buffer as large as the largest BLOB value is allocated.
- Handler structures for all in-use tables are saved in a cache and managed as a FIFO. By default, the cache has 64 entries. If a table has been used by two running threads at the same time, the cache contains two entries for the table. See Section 7.4.8, “How MySQL Opens and Closes Tables”.
- A FLUSH TABLES statement or `mysqladmin flush-tables` command closes all tables that are not in use at once and marks all in-use tables to be closed when the currently executing thread finishes. This effectively frees most in-use memory. FLUSH TABLES does not return until all tables have been closed.

`ps` and other system status programs may report that `mysqld` uses a lot of memory. This may be caused by thread stacks on different memory addresses. For example, the Solaris version of `ps` counts the unused memory between stacks as used memory. You can verify this by checking available swap with `swap -s`. We test `mysqld` with several memory-leakage detectors (both commercial and Open Source), so there should be no memory leaks.

## 7.5.5. How MySQL Uses DNS

When a new client connects to `mysqld`, `mysqld` spawns a new thread to handle the request. This thread first checks whether the hostname is in the hostname cache. If not, the thread attempts to resolve the hostname:

- If the operating system supports the thread-safe `gethostbyaddr_r()` and `gethostbyname_r()` calls, the thread uses them to perform hostname resolution.
- If the operating system does not support the thread-safe calls, the thread locks a mutex and calls `gethostbyaddr()` and `gethostbyname()` instead. In this case, no other thread can resolve hostnames that are not in the hostname cache until the first thread unlocks the mutex.

You can disable DNS hostname lookups by starting `mysqld` with the `--skip-name-resolve` option. However, in this case, you can use only IP numbers in the MySQL grant tables.

If you have a very slow DNS and many hosts, you can get more performance by either disabling DNS lookups with `--skip-name-resolve` or by increasing the `HOST_CACHE_SIZE` define (default value: 128) and recompiling `mysqld`.

You can disable the hostname cache by starting the server with the `--skip-host-cache` option. To clear the hostname cache, issue a `FLUSH HOSTS` statement or execute the `mysqladmin flush-hosts` command.

To disallow TCP/IP connections entirely, start `mysqld` with the `--skip-networking` option.

## 7.6. Disk Issues

- Disk seeks are a huge performance bottleneck. This problem becomes more apparent when the amount of data starts to grow so large that effective caching becomes impossible. For large databases where you access data more or less randomly, you can be sure that you need at least one disk seek to

read and a couple of disk seeks to write things. To minimize this problem, use disks with low seek times.

- Increase the number of available disk spindles (and thereby reduce the seek overhead) by either symlinking files to different disks or striping the disks:

- Using symbolic links

This means that, for `MyISAM` tables, you symlink the index file and data files from their usual location in the data directory to another disk (that may also be striped). This makes both the seek and read times better, assuming that the disk is not used for other purposes as well. See [Section 7.6.1, “Using Symbolic Links”](#).

- Striping

Striping means that you have many disks and put the first block on the first disk, the second block on the second disk, and the  $N$ -th block on the  $(N \text{ MOD } \textit{number\_of\_disks})$  disk, and so on. This means if your normal data size is less than the stripe size (or perfectly aligned), you get much better performance. Striping is very dependent on the operating system and the stripe size, so benchmark your application with different stripe sizes. See [Section 7.1.5, “Using Your Own Benchmarks”](#).

The speed difference for striping is *very* dependent on the parameters. Depending on how you set the striping parameters and number of disks, you may get differences measured in orders of magnitude. You have to choose to optimize for random or sequential access.

- For reliability, you may want to use RAID 0+1 (striping plus mirroring), but in this case, you need  $2 \times N$  drives to hold  $N$  drives of data. This is probably the best option if you have the money for it. However, you may also have to invest in some volume-management software to handle it efficiently.
- A good option is to vary the RAID level according to how critical a type of data is. For example, store semi-important data that can be regenerated on a RAID 0 disk, but store really important data such as host information and logs on a RAID 0+1 or RAID  $N$  disk. RAID  $N$  can be a problem if you have many writes, due to the time required to update the parity bits.
- On Linux, you can get much more performance by using `hdparm` to configure your disk's interface. (Up to 100% under load is not uncommon.) The following `hdparm` options should be quite good for MySQL, and probably for many other applications:

```
hdparm -m 16 -d 1
```

Note that performance and reliability when using this command depend on your hardware, so we strongly suggest that you test your system thoroughly after using `hdparm`. Please consult the `hdparm` manual page for more information. If `hdparm` is not used wisely, filesystem corruption may result, so back up everything before experimenting!

- You can also set the parameters for the filesystem that the database uses:

If you do not need to know when files were last accessed (which is not really useful on a database server), you can mount your filesystems with the `-o noatime` option. That skips updates to the last access time in inodes on the filesystem, which avoids some disk seeks.

On many operating systems, you can set a filesystem to be updated asynchronously by mounting it with the `-o async` option. If your computer is reasonably stable, this should give you more performance without sacrificing too much reliability. (This flag is on by default on Linux.)

## 7.6.1. Using Symbolic Links

You can move tables and databases from the database directory to other locations and replace them with symbolic links to the new locations. You might want to do this, for example, to move a database to a file system with more free space or increase the speed of your system by spreading your tables to different disk.

The recommended way to do this is simply to symlink databases to a different disk. Symlink tables only as a last resort.

### 7.6.1.1. Using Symbolic Links for Databases on Unix

On Unix, the way to symlink a database is first to create a directory on some disk where you have free space and then to create a symlink to it from the MySQL data directory.

```
shell> mkdir /dr1/databases/test
shell> ln -s /dr1/databases/test /path/to/datadir
```

MySQL does not support linking one directory to multiple databases. Replacing a database directory with a symbolic link works as long as you do not make a symbolic link between databases. Suppose that you have a database `db1` under the MySQL data directory, and then make a symlink `db2` that points to `db1`:

```
shell> cd /path/to/datadir
shell> ln -s db1 db2
```

The result is that, or any table `tbl_a` in `db1`, there also appears to be a table `tbl_a` in `db2`. If one client updates `db1.tbl_a` and another client updates `db2.tbl_a`, problems are likely to occur.

If you really need to do this, you can change one of the source files. The file to modify depends on your version of MySQL. For MySQL 4.0 and up, look for the following statement in the `mysys/my_symlink.c` file:

```
if (!(MyFlags & MY_RESOLVE_LINK) ||
    (!lstat(filename,&stat_buff) && S_ISLNK(stat_buff.st_mode)))
```

Before MySQL 4.0, look for this statement in the `mysys/mf_format.c` file:

```
if (flag & 32 || (!lstat(to,&stat_buff) && S_ISLNK(stat_buff.st_mode)))
```

Change the statement to this:

```
if (1)
```

### 7.6.1.2. Using Symbolic Links for Tables on Unix

Before MySQL 4.0, you should not symlink tables unless you are *very* careful with them. The problem is that if you run `ALTER TABLE`, `REPAIR TABLE`, or `OPTIMIZE TABLE` on a symlinked table, the symlinks are removed and replaced by the original files. This happens because these statements work by creating a temporary file in the database directory and replacing the original file with the temporary file when the statement operation is complete.

You should not symlink tables on systems that do not have a fully operational `realpath()` call. (Linux and Solaris support `realpath()`). You can check whether your system supports symbolic links by issuing a `SHOW VARIABLES LIKE 'have_symlink'` statement.

In MySQL 4.0, symlinks are fully supported only for `MyISAM` tables. For files used by tables for other storage engines, you may get strange problems if you try to use symbolic links.

The handling of symbolic links for `MyISAM` tables in MySQL 4.0 works the following way:

- In the data directory, you always have the table format (`.frm`) file, the data (`.MYD`) file, and the index (`.MYI`) file. The data file and index file can be moved elsewhere and replaced in the data directory by symlinks. The format file cannot.
- You can symlink the data file and the index file independently to different directories.
- You can instruct a running MySQL server to perform the symlinking by using the `DATA DIRECTORY` and `INDEX DIRECTORY` options to `CREATE TABLE`. See [Section 13.1.5, “CREATE TABLE Syntax”](#). Alternatively, symlinking can be accomplished manually from the command line using `ln -s` if `mysqld` is not running.
- `myisamchk` does not replace a symlink with the data file or index file. It works directly on the file to which the symlink points. Any temporary files are created in the directory where the data file or index file is located.
- **Note:** When you drop a table that is using symlinks, *both the symlink and the file to which the symlink points are dropped*. This is an extremely good reason why you should *not* run `mysqld` as the system `root` or allow system users to have write access to MySQL database directories.
- If you rename a table with `ALTER TABLE ... RENAME` and you do not move the table to another database, the symlinks in the database directory are renamed to the new names and the data file and index file are renamed accordingly.
- If you use `ALTER TABLE ... RENAME` to move a table to another database, the table is moved to the other database directory. The old symlinks and the files to which they pointed are deleted. In other words, the new table is not symlinked.
- If you are not using symlinks, you should use the `--skip-symbolic-links` option to `mysqld` to ensure that no one can use `mysqld` to drop or rename a file outside of the data directory.

`SHOW CREATE TABLE` does not report if a table has symbolic links prior to MySQL 4.0.15. This is also true for `mysqldump`, which uses `SHOW CREATE TABLE` to generate `CREATE TABLE` statements.

Table symlink operations that are not supported up through MySQL 4.1:

- `ALTER TABLE` ignores the `DATA DIRECTORY` and `INDEX DIRECTORY` table options.
- `BACKUP TABLE` and `RESTORE TABLE` do not respect symbolic links.
- The `.frm` file must *never* be a symbolic link (as indicated previously, only the data and index files can be symbolic links). Attempting to do this (for example, to make synonyms) produces incorrect results. Suppose that you have a database `db1` under the MySQL data directory, a table `tbl1` in this database, and in the `db1` directory you make a symlink `tbl2` that points to `tbl1`:

```
shell> cd /path/to/datadir/db1
shell> ln -s tbl1.frm tbl2.frm
shell> ln -s tbl1.MYD tbl2.MYD
shell> ln -s tbl1.MYI tbl2.MYI
```

Problems result if one thread reads `db1.tbl1` and another thread updates `db1.tbl2`:

- The query cache is “fooled” (it has no way of knowing that `tbl1` has not been updated, so it returns outdated results).



- `ALTER` statements on `tbl2` fail.

### 7.6.1.3. Using Symbolic Links for Databases on Windows

Beginning with MySQL 3.23.16, the `mysqld-max` and `mysql-max-nt` servers for Windows are compiled with the `-DUSE_SYMDIR` option. This enables you to put a database directory on a different disk by setting up a symbolic link to it. This is similar to the way that symbolic links work on Unix, although the procedure for setting up the link is different.

It is necessary to define `USE_SYMDIR` explicitly only before MySQL 4.0; for `mysqld-max` and `mysql-max-nt`, you can enable symbolic links by using the `--symbolic-links` option. As of MySQL 4.0, symbolic links are enabled by default for all Windows servers. If you do not need them, you can disable them with the `--skip-symbolic-links` option.

On Windows, create a symbolic link to a MySQL database by creating a file in the data directory that contains the path to the destination directory. The file should be named `db_name.sym`, where `db_name` is the database name.

Suppose that the MySQL data directory is `C:\mysql\data` and you want to have database `foo` located at `D:\data\foo`. Set up a symlink using this procedure:

1. Make sure that the `D:\data\foo` directory exists by creating it if necessary. If you already have a database directory named `foo` in the data directory, you should move it to `D:\data`. Otherwise, the symbolic link will be ineffective. To avoid problems, make sure that the server is not running when you move the database directory.
2. Create a text file `C:\mysql\data\foo.sym` that contains the pathname `D:\data\foo\`.

After this, all tables created in the database `foo` are created in `D:\data\foo`. *Note that the symbolic link is not used if a directory with the same name as the database exists in the MySQL data directory.*

---

# Chapter 8. Client and Utility Programs

There are many different MySQL client programs that connect to the server to access databases or perform administrative tasks. Other utilities are available as well. These do not establish a client connection with the server but perform MySQL-related operations.

This chapter provides a brief overview of these programs and then a more detailed description of each one. Each program's description indicates its invocation syntax and the options that it understands. See [Chapter 4, \*Using MySQL Programs\*](#), for general information on invoking programs and specifying program options.

## 8.1. Overview of Client and Utility Programs

The following list briefly describes the MySQL client programs and utilities:

- `myisam_ftdump`

A utility that displays information about full-text indexes in [MyISAM](#) tables. See [Section 8.2, “myisam\\_ftdump — Display Full-Text Index information”](#).

- `myisamchk`, `isamchk`

A utility to describe, check, optimize, and repair [MyISAM](#) tables. `isamchk` is a similar program for [ISAM](#) tables. See [Section 8.3, “myisamchk — MyISAM Table-Maintenance Utility”](#).

- `myisamlog`, `isamlog`

Utilities that process the contents of a [MyISAM](#) or [ISAM](#) log file. See [Section 8.4, “myisamlog — Display MyISAM Log File Contents”](#).

- `myisampack`, `pack_isam`

Utilities that compress [MyISAM](#) or [ISAM](#) tables to produce smaller read-only tables. See [Section 8.5, “myisampack — Generate Compressed, Read-Only MyISAM Tables”](#).

- `mysql`

The command-line tool for interactively entering SQL statements or executing them from a file in batch mode. See [Section 8.6, “mysql — The MySQL Command-Line Tool”](#).

- `mysql_explain_log`

A utility that analyzes queries in the MySQL query log using [EXPLAIN](#). See [Section 8.7, “mysql\\_explain\\_log — Use EXPLAIN on Statements in Query Log”](#).

- `mysqlaccess`

A script that checks the access privileges for a hostname, username, and database combination. See [Section 8.8, “mysqlaccess — Client for Checking Access Privileges”](#).

- `mysqladmin`

A client that performs administrative operations, such as creating or dropping databases, reloading the grant tables, flushing tables to disk, and reopening log files. `mysqladmin` can also be used to retrieve version, process, and status information from the server. See [Section 8.9, “mysqladmin — Client for Administering a MySQL Server”](#).

- `mysqlbinlog`  
A utility for reading statements from a binary log. The log of executed statements contained in the binary log files can be used to help recover from a crash. See [Section 8.10, “mysqlbinlog — Utility for Processing Binary Log Files”](#).
- `mysqlcheck`  
A table-maintenance client that checks, repairs, analyzes, and optimizes tables. See [Section 8.11, “mysqlcheck — A Table Maintenance and Repair Program”](#).
- `mysqldump`  
A client that dumps a MySQL database into a file as SQL statements or as tab-separated text files. See [Section 8.12, “mysqldump — A Database Backup Program”](#).
- `mysqlhotcopy`  
A utility that quickly makes backups of `MyISAM` or `ISAM` tables while the server is running. See [Section 8.13, “mysqlhotcopy — A Database Backup Program”](#).
- `mysqlimport`  
A client that imports text files into their respective tables using `LOAD DATA INFILE`. See [Section 8.14, “mysqlimport — A Data Import Program”](#).
- `mysqlshow`  
A client that displays information about databases, tables, columns, and indexes. See [Section 8.15, “mysqlshow — Display Database, Table, and Column Information”](#).
- `mysql_zap`  
A utility that kills processes that match a pattern. [Section 8.16, “mysql\\_zap — Kill Processes That Match a Pattern”](#).
- `perror`  
A utility that displays the meaning of system or MySQL error codes. See [Section 8.17, “perror — Explain Error Codes”](#).
- `replace`  
A utility program that performs string replacement in the input text. See [Section 8.18, “replace — A String-Replacement Utility”](#).

MySQL AB also provides a number of GUI tools for administering and otherwise working with MySQL servers. For basic information about these, see [Chapter 4, \*Using MySQL Programs\*](#).

Each MySQL program takes many different options. Most programs provide a `--help` option that you can use to get a full description of the program's different options. For example, try `mysql --help`.

MySQL client programs that communicate with the server using the MySQL client/server library use the following environment variables:

<code>MYSQL_UNIX_PORT</code>	The default Unix socket file; used for connections to <code>localhost</code>
<code>MYSQL_TCP_PORT</code>	The default port number; used for TCP/IP connections
<code>MYSQL_PWD</code>	The default password

<code>MYSQL_DEBUG</code>	Debug trace options when debugging
<code>TMPDIR</code>	The directory where temporary tables and files are created

Use of `MYSQL_PWD` is insecure. See [Section 5.8.6, “Keeping Your Password Secure”](#).

You can override the default option values or values specified in environment variables for all standard programs by specifying options in an option file or on the command line. See [Section 4.3, “Specifying Program Options”](#).

## 8.2. `myisam_ftdump` — Display Full-Text Index information

`myisam_ftdump` displays information about `FULLTEXT` indexes in `MyISAM` tables. It reads the `MyISAM` index file directly, so it must be run on the server host where the table is located

Invoke `myisam_ftdump` like this:

```
shell> myisam_ftdump [options] tbl_name index_num
```

The `tbl_name` argument should be the name of a `MyISAM` table. You can also specify a table by naming its index file (the file with the `.MYI` suffix). If you do not invoke `myisam_ftdump` in the directory where the table files are located, the table or index file name must be preceded by the pathname to the table's database directory. Index numbers begin with 0.

Example: Suppose that the `test` database contains a table named `mytexttable1` that has the following definition:

```
CREATE TABLE mytexttable
(
  id   INT NOT NULL,
  txt  TEXT NOT NULL,
  PRIMARY KEY (id),
  FULLTEXT (txt)
);
```

The index on `id` is index 0 and the `FULLTEXT` index on `txt` is index 1. If your working directory is the `test` database directory, invoke `myisam_ftdump` as follows:

```
shell> myisam_ftdump mytexttable 1
```

If the pathname to the `test` database directory is `/usr/local/mysql/data/test`, you can also specify the table name argument using that pathname. This is useful if you do not invoke `myisam_ftdump` in the database directory:

```
shell> myisam_ftdump /usr/local/mysql/data/test/mytexttable 1
```

`myisam_ftdump` understands the following options:

- `--help, -h -?`  
Display a help message and exit.
- `--count, -c`  
Calculate per-word statistics (counts and global weights).

- `--dump, -d`  
Dump the index, including data offsets and word weights.
- `--length, -l`  
Report the length distribution.
- `--stats, -s`  
Report global index statistics. This is the default operation if no other operation is specified.
- `--verbose, -v`  
Verbose mode. Print more output about what the program does.

## 8.3. `myisamchk` — MyISAM Table-Maintenance Utility

The `myisamchk` utility gets information about your database tables or checks, repairs, or optimizes them. `myisamchk` works with MyISAM tables (tables that have `.MYD` and `.MYI` files for storing data and indexes). A related utility, `isamchk`, works with ISAM tables (tables that have `.ISD` and `.ISM` files for storing data and indexes).

Invoke `myisamchk` like this:

```
shell> myisamchk [options] tbl_name ...
```

The *options* specify what you want `myisamchk` to do. They are described in the following sections. You can also get a list of options by invoking `myisamchk --help`.

With no options, `myisamchk` simply checks your table as the default operation. To get more information or to tell `myisamchk` to take corrective action, specify options as described in the following discussion.

*tbl\_name* is the database table you want to check or repair. If you run `myisamchk` somewhere other than in the database directory, you must specify the path to the database directory, because `myisamchk` has no idea where the database is located. In fact, `myisamchk` does not actually care whether the files you are working on are located in a database directory. You can copy the files that correspond to a database table into some other location and perform recovery operations on them there.

You can name several tables on the `myisamchk` command line if you wish. You can also specify a table by naming its index file (the file with the `.MYI` suffix). This allows you to specify all tables in a directory by using the pattern `*.MYI`. For example, if you are in a database directory, you can check all the MyISAM tables in that directory like this:

```
shell> myisamchk *.MYI
```

If you are not in the database directory, you can check all the tables there by specifying the path to the directory:

```
shell> myisamchk /path/to/database_dir/*.MYI
```

You can even check all tables in all databases by specifying a wildcard with the path to the MySQL data directory:

```
shell> myisamchk /path/to/datadir/*/*.MYI
```

The recommended way to quickly check all MyISAM and ISAM tables is:

```
shell> myisamchk --silent --fast /path/to/datadir/*/*.MYI
shell> isamchk --silent /path/to/datadir/*/*.ISM
```

If you want to check all MyISAM and ISAM tables and repair any that are corrupted, you can use the following commands:

```
shell> myisamchk --silent --force --fast --update-state \
--key_buffer_size=64M --sort_buffer_size=64M \
--read_buffer_size=1M --write_buffer_size=1M \
/path/to/datadir/*/*.MYI
shell> isamchk --silent --force --key_buffer_size=64M \
--sort_buffer_size=64M --read_buffer_size=1M --write_buffer_size=1M \
/path/to/datadir/*/*.ISM
```

These commands assume that you have more than 64MB free. For more information about memory allocation with `myisamchk`, see [Section 8.3.5, “myisamchk Memory Usage”](#).

You must ensure that no other program is using the tables while you are running `myisamchk`. Otherwise, when you run `myisamchk`, it may display the following error message:

```
warning: clients are using or haven't closed the table properly
```

This means that you are trying to check a table that has been updated by another program (such as the `mysqld` server) that hasn't yet closed the file or that has died without closing the file properly.

If `mysqld` is running, you must force it to flush any table modifications that are still buffered in memory by using `FLUSH TABLES`. You should then ensure that no one is using the tables while you are running `myisamchk`. The easiest way to avoid this problem is to use `CHECK TABLE` instead of `myisamchk` to check tables.

### 8.3.1. `myisamchk` General Options

The options described in this section can be used for any type of table maintenance operation performed by `myisamchk`. The sections following this one describe options that pertain only to specific operations, such as table checking or repairing.

- `--help, -?`

Display a help message and exit.

- `--debug=debug_options, -# debug_options`

Write a debugging log. The `debug_options` string often is `'d:t:o,file_name'`.

- `--silent, -s`

Silent mode. Write output only when errors occur. You can use `-s` twice (`-ss`) to make `myisamchk` very silent.

- `--verbose, -v`

Verbose mode. Print more information about what the program does. This can be used with `-d` and `-e`. Use `-v` multiple times (`-vv, -vvv`) for even more output.

- `--version, -V`

Display version information and exit.

- `--wait, -w`

Instead of terminating with an error if the table is locked, wait until the table is unlocked before continuing. Note that if you are running `mysqld` with external locking disabled, the table can be locked only by another `myisamchk` command.

You can also set the following variables by using `--var_name=value` syntax:

Variable	Default Value
<code>decode_bits</code>	9
<code>ft_max_word_len</code>	version-dependent
<code>ft_min_word_len</code>	4
<code>ft_stopword_file</code>	built-in list
<code>key_buffer_size</code>	523264
<code>myisam_block_size</code>	1024
<code>read_buffer_size</code>	262136
<code>sort_buffer_size</code>	2097144
<code>sort_key_blocks</code>	16
<code>stats_method</code>	nulls_unequal
<code>write_buffer_size</code>	262136

It is also possible to set variables by using `--set-variable=var_name=value` or `-O var_name=value` syntax. However, this syntax is deprecated as of MySQL 4.0.

The possible `myisamchk` variables and their default values can be examined with `myisamchk -help`:

`sort_buffer_size` is used when the keys are repaired by sorting keys, which is the normal case when you use `--recover`.

`key_buffer_size` is used when you are checking the table with `--extend-check` or when the keys are repaired by inserting keys row by row into the table (like when doing normal inserts). Repairing through the key buffer is used in the following cases:

- You use `--safe-recover`.
- The temporary files needed to sort the keys would be more than twice as big as when creating the key file directly. This is often the case when you have large key values for `CHAR`, `VARCHAR`, or `TEXT` columns, because the sort operation needs to store the complete key values as it proceeds. If you have lots of temporary space and you can force `myisamchk` to repair by sorting, you can use the `--sort-recover` option.

Repairing through the key buffer takes much less disk space than using sorting, but is also much slower.

If you want a faster repair, set the `key_buffer_size` and `sort_buffer_size` variables to about 25% of your available memory. You can set both variables to large values, because only one of them is used at a time.

`myisam_block_size` is the size used for index blocks. It is available as of MySQL 4.0.0.

`stats_method` influences how `NULL` values are treated for index statistics collection when the `--analyze` option is given. It acts like the `myisam_stats_method` system variable. For more information, see the description of `myisam_stats_method` in [Section 5.2.2, “System Variables”](#), and [Section 7.4.7, “MyISAM Index Statistics Collection”](#). `stats_method` was added in MySQL 4.1.15/5.0.14. For older versions, the statistics collection method is equivalent to `nulls_equal`.

The `ft_min_word_len` and `ft_max_word_len` variables are available as of MySQL 4.0.0. `ft_stopword_file` is available as of MySQL 4.0.19.

`ft_min_word_len` and `ft_max_word_len` indicate the minimum and maximum word length for `FULLTEXT` indexes. `ft_stopword_file` names the stopword file. These need to be set under the following circumstances.

If you use `myisamchk` to perform an operation that modifies table indexes (such as repair or analyze), the `FULLTEXT` indexes are rebuilt using the default full-text parameter values for minimum and maximum word length and the stopword file unless you specify otherwise. This can result in queries failing.

The problem occurs because these parameters are known only by the server. They are not stored in `MyISAM` index files. To avoid the problem if you have modified the minimum or maximum word length or the stopword file in the server, specify the same `ft_min_word_len`, `ft_max_word_len`, and `ft_stopword_file` values to `myisamchk` that you use for `mysqld`. For example, if you have set the minimum word length to 3, you can repair a table with `myisamchk` like this:

```
shell> myisamchk --recover --ft_min_word_len=3 tbl_name.MYI
```

To ensure that `myisamchk` and the server use the same values for full-text parameters, you can place each one in both the `[mysqld]` and `[myisamchk]` sections of an option file:

```
[mysqld]
ft_min_word_len=3

[myisamchk]
ft_min_word_len=3
```

An alternative to using `myisamchk` is to use the `REPAIR TABLE`, `ANALYZE TABLE`, `OPTIMIZE TABLE`, or `ALTER TABLE`. These statements are performed by the server, which knows the proper full-text parameter values to use.

## 8.3.2. `myisamchk` Check Options

`myisamchk` supports the following options for table checking operations:

- `--check, -c`

Check the table for errors. This is the default operation if you specify no option that selects an operation type explicitly.

- `--check-only-changed, -C`

Check only tables that have changed since the last check.

- `--extend-check, -e`

Check the table very thoroughly. This is quite slow if the table has many indexes. This option should only be used in extreme cases. Normally, `myisamchk` or `myisamchk --medium-check` should be able to determine whether there are any errors in the table.



If you are using `--extend-check` and have plenty of memory, setting the `key_buffer_size` variable to a large value helps the repair operation run faster.

- `--fast, -F`

Check only tables that haven't been closed properly.

- `--force, -f`

Do a repair operation automatically if `myisamchk` finds any errors in the table. The repair type is the same as that specified with the `--recover` or `-r` option.

- `--information, -i`

Print informational statistics about the table that is checked.

- `--medium-check, -m`

Do a check that is faster than an `--extend-check` operation. This finds only 99.99% of all errors, which should be good enough in most cases.

- `--read-only, -T`

Don't mark the table as checked. This is useful if you use `myisamchk` to check a table that is in use by some other application that does not use locking, such as `mysqld` when run with external locking disabled.

- `--update-state, -U`

Store information in the `.MYI` file to indicate when the table was checked and whether the table crashed. This should be used to get full benefit of the `--check-only-changed` option, but you shouldn't use this option if the `mysqld` server is using the table and you are running it with external locking disabled.

### 8.3.3. `myisamchk` Repair Options

`myisamchk` supports the following options for table repair operations:

- `--backup, -B`

Make a backup of the `.MYD` file as `file_name-time.BAK`

- `--character-sets-dir=path`

The directory where character sets are installed. See [Section 5.10.1, “The Character Set Used for Data and Sorting”](#).

- `--correct-checksum`

Correct the checksum information for the table.

- `--data-file-length=len, -D len`

Maximum length of the data file (when re-creating data file when it is “full”).

- `--extend-check, -e`

Do a repair that tries to recover every possible row from the data file. Normally, this also finds a lot of garbage rows. Don't use this option unless you are totally desperate.

- `--force, -f`

Overwrite old intermediate files (files with names like `tbl_name.TMD`) instead of aborting.

- `--keys-used=val, -k val`

For `myisamchk`, the option value is a bit-value that indicates which indexes to update. Each binary bit of the option value corresponds to a table index, where the first index is bit 0. For `isamchk`, the option value indicates that only the first `val` of the table indexes should be updated. In either case, an option value of 0 disables updates to all indexes, which can be used to get faster inserts. Deactivated indexes can be reactivated by using `myisamchk -r` or `(isamchk -r)`.

- `--no-symlinks, -l`

Do not follow symbolic links. Normally `myisamchk` repairs the table that a symlink points to. This option does not exist as of MySQL 4.0 because versions from 4.0 on do not remove symlinks during repair operations.

- `--max-record-length=len`

Skip rows larger than the given length if `myisamchk` cannot allocate memory to hold them. This option was added in MySQL 4.1.1.

- `--parallel-recover, -p`

Uses the same technique as `-r` and `-n`, but creates all the keys in parallel, using different threads. This option was added in MySQL 4.0.2. *This is beta-quality code; use at your own risk!*

- `--quick, -q`

Achieve a faster repair by not modifying the data file. You can specify this option twice to force `myisamchk` to modify the original data file in case of duplicate keys.

- `--recover, -r`

Do a repair that can fix almost any problem except unique keys that are not unique (which is an extremely unlikely error with `ISAM/MyISAM` tables). If you want to recover a table, this is the option to try first. You should try `--safe-recover` only if `myisamchk` reports that the table cannot be recovered by `--recover`. (In the unlikely case that `--recover` fails, the data file remains intact.)

If you have lots of memory, you should increase the value of `sort_buffer_size`.

- `--safe-recover, -o`

Do a repair using an old recovery method that reads through all rows in order and updates all index trees based on the rows found. This is an order of magnitude slower than `--recover`, but can handle a couple of very unlikely cases that `--recover` cannot. This recovery method also uses much less disk space than `--recover`. Normally, you should repair first using `--recover`, and then with `--safe-recover` only if `--recover` fails.

If you have lots of memory, you should increase the value of `key_buffer_size`.

- `--set-character-set=name`

Change the character set used by the table indexes. This option was replaced by `--set-collation` in MySQL 4.1.1.

- `--set-collation=name`

Specify the collation to use for sorting table indexes. The character set name is implied by the first part of the collation name. This option was added in MySQL 4.1.11.

- `--sort-recover, -n`

Force `myisamchk` to use sorting to resolve the keys even if the temporary files should be very big.

- `--tmpdir=path, -t path`

Path of the directory to be used for storing temporary files. If this is not set, `myisamchk` uses the value of the `TMPDIR` environment variable. Starting from MySQL 4.1, `tmpdir` can be set to a list of directory paths that are used successively in round-robin fashion for creating temporary files. The separator character between directory names should be colon (':') on Unix and semicolon (';') on Windows, NetWare, and OS/2.

- `--unpack, -u`

Unpack a table that was packed with `myisampack`.

### 8.3.4. Other `myisamchk` Options

`myisamchk` supports the following options for actions other than table checks and repairs:

- `--analyze, -a`

Analyze the distribution of key values. This improves join performance by enabling the join optimizer to better choose the order in which to join the tables and which indexes it should use. To obtain information about the key distribution, use a `myisamchk --description --verbose tbl_name` command or the `SHOW INDEX FROM tbl_name` statement.

- `--block-search=offset, -b offset`

Find the record that a block at the given offset belongs to.

- `--description, -d`

Print some descriptive information about the table.

- `--set-auto-increment[=value], -A[value]`

Force `AUTO_INCREMENT` numbering for new records to start at the given value (or higher, if there are existing records with `AUTO_INCREMENT` values this large). If `value` is not specified, `AUTO_INCREMENT` numbers for new records begin with the largest value currently in the table, plus one.

- `--sort-index, -S`

Sort the index tree blocks in high-low order. This optimizes seeks and makes table scans that use indexes faster.

- `--sort-records=N, -R N`

Sort records according to a particular index. This makes your data much more localized and may speed up range-based `SELECT` and `ORDER BY` operations that use this index. (The first time you use this option to sort a table, it may be very slow.) To determine a table's index numbers, use `SHOW INDEX`, which displays a table's indexes in the same order that `myisamchk` sees them. Indexes are numbered beginning with 1.

If keys are not packed (`PACK_KEYS=0`), they have the same length, so when `myisamchk` sorts and moves records, it just overwrites record offsets in the index. If keys are packed (`PACK_KEYS=1`), `myisamchk` must unpack key blocks first, then re-create indexes and pack the key blocks again. (In this case, re-creating indexes is faster than updating offsets for each index.)

### 8.3.5. `myisamchk` Memory Usage

Memory allocation is important when you run `myisamchk`. `myisamchk` uses no more memory than its memory-related variables are set to. If you are going to use `myisamchk` on very large tables, you should first decide how much memory you want it to use. The default is to use only about 3MB to perform repairs. By using larger values, you can get `myisamchk` to operate faster. For example, if you have more than 32MB RAM, you could use options such as these (in addition to any other options you might specify):

```
shell> myisamchk --sort_buffer_size=16M --key_buffer_size=16M \
           --read_buffer_size=1M --write_buffer_size=1M ...
```

Using `--sort_buffer_size=16M` should probably be enough for most cases.

Be aware that `myisamchk` uses temporary files in `TMPDIR`. If `TMPDIR` points to a memory filesystem, you may easily get out of memory errors. If this happens, run `myisamchk` with the `--tmpdir=path` option to specify some directory located on a filesystem that has more space.

When repairing, `myisamchk` also needs a lot of disk space:

- Double the size of the data file (the original file and a copy). This space is not needed if you do a repair with `--quick`; in this case, only the index file is re-created. *This space must be available on the same filesystem as the original data file*, as the copy is created in the same directory as the original.
- Space for the new index file that replaces the old one. The old index file is truncated at the start of the repair operation, so you usually ignore this space. This space must be available on the same filesystem as the original data file.
- When using `--recover` or `--sort-recover` (but not when using `--safe-recover`), you need space for a sort buffer. The following formula yields the amount of space required:

$$(\textit{largest\_key} + \textit{row\_pointer\_length}) \times \textit{number\_of\_rows} \times 2$$

You can check the length of the keys and the `row_pointer_length` with `myisamchk -dv tbl_name`. This space is allocated in the temporary directory (specified by `TMPDIR` or `--tmpdir=path`).

If you have a problem with disk space during repair, you can try `--safe-recover` instead of `--recover`.

## 8.4. `myisamlog` — Display MyISAM Log File Contents

`myisamlog` processes the contents of a MyISAM log file. `isamlog` is similar, but is used with ISAM log files.

Invoke `myisamlog` or `isamlog` like this:

```
shell> myisamlog [options] [log_file [tbl_name] ...]
shell> isamlog [options] [log_file [tbl_name] ...]
```

The default operation is update (`-u`). If a recovery is done (`-r`), all writes and possibly updates and deletes are done and errors are only counted. The default log file name is `myisam.log` for `myisamlog` and `isam.log` for `isamlog` if no `log_file` argument is given. If tables are named on the command line, only those tables are updated.

`myisamlog` and `isamlog` understand the following options:

- `-?, -I`  
Display a help message and exit.
- `-c N`  
Execute only `N` commands.
- `-f N`  
Specify the maximum number of open files.
- `-i`  
Display extra information before exiting.
- `-o offset`  
Specify the starting offset.
- `-p N`  
Remove `N` components from path.
- `-r`  
Perform a recovery operation.
- `-R record_pos_file record_pos`  
Specify record position file and record position.
- `-u`  
Perform an update operation.
- `-v`  
Verbose mode. Print more output about what the program does. This option can be given multiple times to produce more and more output.
- `-w write_file`  
Specify the write file.

- `-V`

Display version information.

## 8.5. `myisampack` — Generate Compressed, Read-Only MyISAM Tables

The `myisampack` utility compresses MyISAM tables. `myisampack` works by compressing each column in the table separately. Usually, `myisampack` packs the data file 40%-70%.

When the table is used later, the server reads into memory the information needed to decompress columns. This results in much better performance when accessing individual rows, because you only have to uncompress exactly one row.

MySQL uses `mmap()` when possible to perform memory mapping on compressed tables. If `mmap()` does not work, MySQL falls back to normal read/write file operations.

A similar utility, `pack_isam`, compresses ISAM tables. Because ISAM tables are deprecated, this section discusses only `myisampack`, but the general procedures for using `myisampack` are also true for `pack_isam` unless otherwise specified. References to `myisamchk` should be read as references to `isamchk` if you are using `pack_isam`.

Please note the following:

- If the `mysqld` server was invoked with external locking disabled, it is not a good idea to invoke `myisampack` if the table might be updated by the server during the packing process. It is safest to compress tables with the server stopped.
- After packing a table, it becomes read-only. This is generally intended (such as when accessing packed tables on a CD). Allowing writes to a packed table is on our TODO list, but with low priority.
- `myisampack` can pack BLOB or TEXT columns. (The older `pack_isam` program for ISAM tables does not have this capability.)

Invoke `myisampack` like this:

```
shell> myisampack [options] file_name ...
```

Each filename argument should be the name of an index (`.MYI`) file. If you are not in the database directory, you should specify the pathname to the file. It is permissible to omit the `.MYI` extension.

After you compress a table with `myisampack`, you should use `myisamchk -rq` to rebuild its indexes. [Section 8.3, “myisamchk — MyISAM Table-Maintenance Utility”](#).

`myisampack` supports the following options:

- `--help, -?`

Display a help message and exit.

- `--backup, -b`

Make a backup of each table's data file using the name `tbl_name.OLD`.

- `--character-sets-dir=path`

The directory where character sets are installed. See [Section 5.10.1](#), “The Character Set Used for Data and Sorting”.

- `--debug[=debug_options], -# [debug_options]`

Write a debugging log. The `debug_options` string often is `'d:t:o,file_name'`.

- `--force, -f`

Produce a packed table even if it becomes larger than the original or if the intermediate file from an earlier invocation of `mysampack` exists. (`mysampack` creates an intermediate file named `tbl_name.TMD` in the database directory while it compresses the table. If you kill `mysampack`, the `.TMD` file might not be deleted.) Normally, `mysampack` exits with an error if it finds that `tbl_name.TMD` exists. With `--force`, `mysampack` packs the table anyway.

- `--join=big_tbl_name, -j big_tbl_name`

Join all tables named on the command line into a single table `big_tbl_name`. All tables that are to be combined *must* have identical structure (same column names and types, same indexes, and so forth).

- `--packlength=len, -p len`

Specify the row length storage size, in bytes. The value should be 1, 2, or 3. `mysampack` stores all rows with length pointers of 1, 2, or 3 bytes. In most normal cases, `mysampack` can determine the correct length value before it begins packing the file, but it may notice during the packing process that it could have used a shorter length. In this case, `mysampack` prints a note that you could use a shorter row length the next time you pack the same file.

- `--silent, -s`

Silent mode. Write output only when errors occur.

- `--test, -t`

Do not actually pack the table, just test packing it.

- `--tmpdir=path, -T path`

Use the named directory as the location where `mysampack` creates temporary files.

- `--verbose, -v`

Verbose mode. Write information about the progress of the packing operation and its result.

- `--version, -V`

Display version information and exit.

- `--wait, -w`

Wait and retry if the table is in use. If the `mysqld` server was invoked with external locking disabled, it is not a good idea to invoke `mysampack` if the table might be updated by the server during the packing process.

The following sequence of commands illustrates a typical table compression session:

```

shell> ls -l station.*
-rw-rw-r-- 1 monty my 994128 Apr 17 19:00 station.MYD
-rw-rw-r-- 1 monty my 53248 Apr 17 19:00 station.MYI
-rw-rw-r-- 1 monty my 5767 Apr 17 19:00 station.frm

shell> myisamchk -dvv station

MyISAM file: station
Isam-version: 2
Creation time: 1996-03-13 10:08:58
Recover time: 1997-02-02 3:06:43
Data records: 1192 Deleted blocks: 0
Datafile parts: 1192 Deleted data: 0
Datafile pointer (bytes): 2 Keyfile pointer (bytes): 2
Max datafile length: 54657023 Max keyfile length: 33554431
Recordlength: 834
Record format: Fixed length

table description:
Key Start Len Index Type Root Blocksize Rec/key
1 2 4 unique unsigned long 1024 1024 1
2 32 30 multip. text 10240 1024 1

Field Start Length Type
1 1 1
2 2 4
3 6 4
4 10 1
5 11 20
6 31 1
7 32 30
8 62 35
9 97 35
10 132 35
11 167 4
12 171 16
13 187 35
14 222 4
15 226 16
16 242 20
17 262 20
18 282 20
19 302 30
20 332 4
21 336 4
22 340 1
23 341 8
24 349 8
25 357 8
26 365 2
27 367 2
28 369 4
29 373 4
30 377 1
31 378 2
32 380 8
33 388 4
34 392 4
35 396 4
36 400 4
37 404 1
38 405 4
39 409 4
40 413 4
41 417 4
42 421 4
43 425 4
44 429 20
45 449 30
46 479 1
47 480 1
48 481 79
49 560 79
50 639 79
51 718 79
52 797 8
53 805 1
54 806 1

```



```

55 807 20
56 827 4
57 831 4

shell> myisampack station.MYI
Compressing station.MYI: (1192 records)
- Calculating statistics

normal:      20  empty-space:  16  empty-zero:    12  empty-fill:  11
pre-space:   0  end-space:    12  table-lookups: 5  zero:         7
Original trees: 57 After join: 17
- Compressing file
87.14%
Remember to run myisamchk -rq on compressed tables

shell> ls -l station.*
-rw-rw-r-- 1 monty my 127874 Apr 17 19:00 station.MYD
-rw-rw-r-- 1 monty my 55296 Apr 17 19:04 station.MYI
-rw-rw-r-- 1 monty my 5767 Apr 17 19:00 station.frm

shell> myisamchk -dvv station

MyISAM file:      station
Isam-version:     2
Creation time:    1996-03-13 10:08:58
Recover time:    1997-04-17 19:04:26
Data records:    1192 Deleted blocks:      0
Datafile parts: 1192 Deleted data:         0
Datafile pointer (bytes): 3 Keyfile pointer (bytes): 1
Max datafile length: 16777215 Max keyfile length: 131071
Recordlength:   834
Record format:   Compressed

table description:
Key Start Len Index Type Root Blocksize Rec/key
1 2 4 unique unsigned long 10240 1024 1
2 32 30 multip. text 54272 1024 1

Field Start Length Type Huff tree Bits
1 1 1 constant 1 0
2 2 4 zerofill(1) 2 9
3 6 4 no zeros, zerofill(1) 2 9
4 10 1 3 9
5 11 20 table-lookup 4 0
6 31 1 3 9
7 32 30 no endspace, not_always 5 9
8 62 35 no endspace, not_always, no empty 6 9
9 97 35 no empty 7 9
10 132 35 no endspace, not_always, no empty 6 9
11 167 4 zerofill(1) 2 9
12 171 16 no endspace, not_always, no empty 5 9
13 187 35 no endspace, not_always, no empty 6 9
14 222 4 zerofill(1) 2 9
15 226 16 no endspace, not_always, no empty 5 9
16 242 20 no endspace, not_always 8 9
17 262 20 no endspace, no empty 8 9
18 282 20 no endspace, no empty 5 9
19 302 30 no endspace, no empty 6 9
20 332 4 always zero 2 9
21 336 4 always zero 2 9
22 340 1 3 9
23 341 8 table-lookup 9 0
24 349 8 table-lookup 10 0
25 357 8 always zero 2 9
26 365 2 2 9
27 367 2 no zeros, zerofill(1) 2 9
28 369 4 no zeros, zerofill(1) 2 9
29 373 4 table-lookup 11 0
30 377 1 3 9
31 378 2 no zeros, zerofill(1) 2 9
32 380 8 no zeros 2 9
33 388 4 always zero 2 9
34 392 4 table-lookup 12 0
35 396 4 no zeros, zerofill(1) 13 9
36 400 4 no zeros, zerofill(1) 2 9
37 404 1 2 9
38 405 4 no zeros 2 9
39 409 4 always zero 2 9
40 413 4 no zeros 2 9
41 417 4 always zero 2 9
42 421 4 no zeros 2 9

```

43	425	4	always zero	2	9
44	429	20	no empty	3	9
45	449	30	no empty	3	9
46	479	1		14	4
47	480	1		14	4
48	481	79	no endspace, no empty	15	9
49	560	79	no empty	2	9
50	639	79	no empty	2	9
51	718	79	no endspace	16	9
52	797	8	no empty	2	9
53	805	1		17	1
54	806	1		3	9
55	807	20	no empty	3	9
56	827	4	no zeros, zerofill(2)	2	9
57	831	4	no zeros, zerofill(1)	2	9

`myisampack` displays the following kinds of information:

- `normal`

The number of columns for which no extra packing is used.

- `empty-space`

The number of columns containing values that are only spaces. These occupy one bit.

- `empty-zero`

The number of columns containing values that are only binary zeros. These occupy one bit.

- `empty-fill`

The number of integer columns that do not occupy the full byte range of their type. These are changed to a smaller type. For example, a `BIGINT` column (eight bytes) can be stored as a `TINY-INT` column (one byte) if all its values are in the range from `-128` to `127`.

- `pre-space`

The number of decimal columns that are stored with leading spaces. In this case, each value contains a count for the number of leading spaces.

- `end-space`

The number of columns that have a lot of trailing spaces. In this case, each value contains a count for the number of trailing spaces.

- `table-lookup`

The column had only a small number of different values, which were converted to an `ENUM` before Huffman compression.

- `zero`

The number of columns for which all values are zero.

- `Original trees`

The initial number of Huffman trees.

- `After join`

The number of distinct Huffman trees left after joining trees to save some header space.

After a table has been compressed, `myisamchk -dvv` prints additional information about each column:

- `Type`

The data type. The value may contain any of the following descriptors:

- `constant`

All rows have the same value.

- `no endspace`

Do not store endspace.

- `no endspace, not_always`

Do not store endspace and do not do endspace compression for all values.

- `no endspace, no empty`

Do not store endspace. Do not store empty values.

- `table-lookup`

The column was converted to an `ENUM`.

- `zerofill(N)`

The most significant `N` bytes in the value are always 0 and are not stored.

- `no zeros`

Do not store zeros.

- `always zero`

Zero values are stored using one bit.

- `Huff tree`

The number of the Huffman tree associated with the column.

- `Bits`

The number of bits used in the Huffman tree.

After you run `myisampack`, you must run `myisamchk` to re-create any indexes. At this time, you can also sort the index blocks and create statistics needed for the MySQL optimizer to work more efficiently:

```
shell> myisamchk -rq --sort-index --analyze tbl_name.MYI
```

A similar procedure applies for `ISAM` tables. After using `pack_isam`, use `isamchk` to re-create the indexes:

```
shell> isamchk -rq --sort-index --analyze tbl_name.ISM
```

After you have installed the packed table into the MySQL database directory, you should execute `mysqladmin flush-tables` to force `mysqld` to start using the new table.

To unpack a packed table, use the `--unpack` option to `myisamchk` or `isamchk`.

## 8.6. `mysql` — The MySQL Command-Line Tool

`mysql` is a simple SQL shell (with GNU `readline` capabilities). It supports interactive and non-interactive use. When used interactively, query results are presented in an ASCII-table format. When used non-interactively (for example, as a filter), the result is presented in tab-separated format. The output format can be changed using command options.

If you have problems due to insufficient memory for large result sets, use the `--quick` option. This forces `mysql` to retrieve results from the server a row at a time rather than retrieving the entire result set and buffering it in memory before displaying it. This is done by returning the result set using the `mysql_use_result()` C API function in the client/server library rather than `mysql_store_result()`.

Using `mysql` is very easy. Invoke it from the prompt of your command interpreter as follows:

```
shell> mysql db_name
```

Or:

```
shell> mysql --user=user_name --password=your_password db_name
```

Then type an SQL statement, end it with `;`, `\g`, or `\G` and press Enter.

You can execute SQL statements in a script file (batch file) like this:

```
shell> mysql db_name < script.sql > output.tab
```

### 8.6.1. `mysql` Options

`mysql` supports the following options:

- `--help, -?`

Display a help message and exit.

- `--auto-rehash`

Enable automatic rehashing. This option is on by default, which enables table and column name completion. Use `--skip-auto-rehash` to disable rehashing. That causes `mysql` to start faster, but you must issue the `rehash` command if you want to use table and column name completion.

- `--batch, -B`

Print results using tab as the column separator, with each row on a new line. With this option, `mysql` does not use the history file.

- `--character-sets-dir=path`

The directory where character sets are installed. See [Section 5.10.1, “The Character Set Used for Data and Sorting”](#).

- `--column-names`  
Write column names in results.
- `--compress, -C`  
Compress all information sent between the client and the server if both support compression.
- `--database=db_name, -D db_name`  
The database to use. This is useful primarily in an option file.
- `--debug[=debug_options], -# [debug_options]`  
Write a debugging log. The *debug\_options* string often is `'d:t:o,file_name'`. The default is `'d:t:o,/tmp/mysql.trace'`.
- `--debug-info, -T`  
Print some debugging information when the program exits.
- `--default-character-set=charset_name`  
Use *charset\_name* as the default character set. See [Section 5.10.1, “The Character Set Used for Data and Sorting”](#).
- `--delimiter=str`  
Set the statement delimiter. The default is the semicolon character (`;`).
- `--execute=statement, -e statement`  
Execute the statement and quit. The default output format is like that produced with `--batch`. See [Section 4.3.1, “Using Options on the Command Line”](#), for some examples.
- `--force, -f`  
Continue even if an SQL error occurs.
- `--host=host_name, -h host_name`  
Connect to the MySQL server on the given host.
- `--html, -H`  
Produce HTML output.
- `--ignore-spaces, -i`  
Ignore spaces after function names. The effect of this is described in the discussion for the `IGNORE_SPACE` SQL mode (see [Section 5.2.5, “SQL Modes”](#)).
- `--line-numbers`  
Write line numbers for errors. Disable this with `--skip-line-numbers`.
- `--local-infile[={0|1}]`  
Enable or disable `LOCAL` capability for `LOAD DATA INFILE`. With no value, the option enables `LOCAL`. The option may be given as `--local-infile=0` or `--local-infile=1` to explicitly

disable or enable `LOCAL`. Enabling `LOCAL` has no effect if the server does not also support it.

- `--named-commands, -G`

Enable named `mysql` commands. Long-format commands are allowed, not just short-format commands. For example, `quit` and `\q` both are recognized. Use `--skip-named-commands` to disable named commands. See [Section 8.6.2, “mysql Commands”](#).

- `--no-auto-rehash, -A`

Deprecated form of `-skip-auto-rehash`. See the description for `--auto-rehash`.

- `--no-beep, -b`

Do not beep when errors occur.

- `--no-named-commands, -g`

Disable named commands. Use the `\*` form only, or use named commands only at the beginning of a line ending with a semicolon (`;`). As of MySQL 3.23.22, `mysql` starts with this option *enabled* by default. However, even with this option, long-format commands still work from the first line. See [Section 8.6.2, “mysql Commands”](#).

- `--no-pager`

Deprecated form of `--skip-pager`. See the `--pager` option.

- `--no-tee`

Do not copy output to a file. [Section 8.6.2, “mysql Commands”](#), discusses tee files further.

- `--one-database, -o`

Ignore statements except those for the default database named on the command line. This is useful for skipping updates to other databases in the binary log.

- `--pager[=command]`

Use the given command for paging query output. If the command is omitted, the default pager is the value of your `PAGER` environment variable. Valid pagers are `less`, `more`, `cat [> filename]`, and so forth. This option works only on Unix. It does not work in batch mode. To disable paging, use `--skip-pager`. [Section 8.6.2, “mysql Commands”](#), discusses output paging further.

- `--password[=password], -p[password]`

The password to use when connecting to the server. If you use the short option form (`-p`), you *cannot* have a space between the option and the password. If you omit the `password` value following the `--password` or `-p` option on the command line, you are prompted for one.

Specifying a password on the command line should be considered insecure. See [Section 5.8.6, “Keeping Your Password Secure”](#).

- `--port=port_num, -P port_num`

The TCP/IP port number to use for the connection.

- `--prompt=format_str`

Set the prompt to the specified format. The default is `mysql>`. The special sequences that the prompt can contain are described in [Section 8.6.2, “mysql Commands”](#).

- `--protocol={TCP|SOCKET|PIPE|MEMORY}`

The connection protocol to use. Added in MySQL 4.1.

- `--quick, -q`

Do not cache each query result, print each row as it is received. This may slow down the server if the output is suspended. With this option, `mysql` does not use the history file.

- `--raw, -r`

Write column values without escape conversion. Often used with the `--batch` option.

- `--reconnect`

If the connection to the server is lost, automatically try to reconnect. A single reconnect attempt is made each time the connection is lost. To suppress reconnection behavior, use `--skip-reconnect`. Added in MySQL 4.1.0.

- `--safe-updates, --i-am-a-dummy, -U`

Allow only those `UPDATE` and `DELETE` statements that specify which rows to modify by using key values. If you have set this option in an option file, you can override it by using `--safe-updates` on the command line. See [Section 8.6.5, “mysql Tips”](#), for more information about this option.

- `--secure-auth`

Do not send passwords to the server in old (pre-4.1.1) format. This prevents connections except for servers that use the newer password format. This option was added in MySQL 4.1.1.

- `--sigint-ignore`

Ignore `SIGINT` signals (typically the result of typing Control-C). This option was added in MySQL 4.1.6.

- `--silent, -s`

Silent mode. Produce less output. This option can be given multiple times to produce less and less output.

- `--skip-column-names, -N`

Do not write column names in results.

- `--skip-line-numbers, -L`

Do not write line numbers for errors. Useful when you want to compare result files that include error messages.

- `--socket=path, -S path`

For connections to `localhost`, the Unix socket file to use, or, on Windows, the name of the named pipe to use.

- `--ssl*`

Options that begin with `--ssl` specify whether to connect to the server via SSL and indicate where to find SSL keys and certificates. See [Section 5.8.7.3, “SSL Command Options”](#).

- `--table, -t`

Display output in table format. This is the default for interactive use, but can be used to produce table output in batch mode.

- `--tee=file_name`

Append a copy of output to the given file. This option does not work in batch mode. in [Section 8.6.2, “mysql Commands”](#), discusses tee files further.

- `--unbuffered, -n`

Flush the buffer after each query.

- `--user=user_name, -u user_name`

The MySQL username to use when connecting to the server.

- `--verbose, -v`

Verbose mode. Produce more output about what the program does. This option can be given multiple times to produce more and more output. (For example, `-v -v -v` produces table output format even in batch mode.)

- `--version, -V`

Display version information and exit.

- `--vertical, -E`

Print query output rows vertically (one line per column value). Without this option, you can specify vertical output for individual statements by terminating them with `\G`.

- `--wait, -w`

If the connection cannot be established, wait and retry instead of aborting.

- `--xml, -X`

Produce XML output.

You can also set the following variables by using `--var_name=value` syntax:

- `connect_timeout`

The number of seconds before connection timeout. (Default value is 0.)

- `max_allowed_packet`

The maximum packet length to send to or receive from the server. (Default value is 16MB.)

- `max_join_size`

The automatic limit for rows in a join when using `--safe-updates`. (Default value is 1,000,000.)

- `net_buffer_length`

The buffer size for TCP/IP and socket communication. (Default value is 16KB.)



- `select_limit`

The automatic limit for `SELECT` statements when using `--safe-updates`. (Default value is 1,000.)

It is also possible to set variables by using `--set-variable=var_name=value` or `-O var_name=value` syntax. In MySQL 4.1, this syntax is deprecated.

On Unix, the `mysql` client writes a record of executed statements to a history file. By default, the history file is named `.mysql_history` and is created in your home directory. To specify a different file, set the value of the `MYSQL_HISTFILE` environment variable.

If you do not want to maintain a history file, first remove `.mysql_history` if it exists, and then use either of the following techniques:

- Set the `MYSQL_HISTFILE` variable to `/dev/null`. To cause this setting to take effect each time you log in, put the setting in one of your shell's startup files.
- Create `.mysql_history` as a symbolic link to `/dev/null`:

```
shell> ln -s /dev/null $HOME/.mysql_history
```

You need do this only once.

## 8.6.2. `mysql` Commands

`mysql` sends each SQL statement that you issue to the server to be executed. There is also a set of commands that `mysql` itself interprets. For a list of these commands, type `help` or `\h` at the `mysql>` prompt:

```
mysql> help

MySQL commands:
?          (? )  Synonym for `help`.
clear      (c)   Clear command.
connect    (r)   Reconnect to the server. Optional arguments are db and host.
delimiter (d)   Set statement delimiter. NOTE: Takes the rest of the line as
          new delimiter.
edit       (e)   Edit command with $EDITOR.
ego        (G)   Send command to mysql server, display result vertically.
exit       (q)   Exit mysql. Same as quit.
go         (g)   Send command to mysql server.
help       (h)   Display this help.
nopager    (n)   Disable pager, print to stdout.
notee      (t)   Don't write into outfile.
pager      (P)   Set PAGER [to_pager]. Print the query results via PAGER.
print      (p)   Print current command.
prompt     (R)   Change your mysql prompt.
quit       (q)   Quit mysql.
rehash     (#)  Rebuild completion hash.
source     (.)  Execute an SQL script file. Takes a file name as an argument.
status     (s)  Get status information from the server.
system     (!)  Execute a system shell command.
tee        (T)  Set outfile [to_outfile]. Append everything into given
          outfile.
use        (u)  Use another database. Takes database name as argument.
charset    (C)  Switch to another charset. Might be needed for processing binlog with multi-byte charset.
warnings   (W)  Show warnings after every statement.
nowarning  (w)  Don't show warnings after every statement.

For server side help, type 'help contents'
```

Each command has both a long and short form. The long form is not case sensitive; the short form is.

The long form can be followed by an optional semicolon terminator, but the short form should not.

If you provide an argument to the `help` command, `mysql` uses it as a search string to access server-side help from the contents of the MySQL Reference Manual. For more information, see [Section 8.6.3](#), “`mysql Server-Side Help`”.

In the `delimiter` command, you should avoid the use of the backslash (`\`) character because that is the escape character for MySQL.

The `edit`, `nopager`, `pager`, and `system` commands work only in Unix.

The `status` command provides some information about the connection and the server you are using. If you are running in `--safe-updates` mode, `status` also prints the values for the `mysql` variables that affect your queries.

To log queries and their output, use the `tee` command. All the data displayed on the screen is appended into a given file. This can be very useful for debugging purposes also. You can enable this feature on the command line with the `--tee` option, or interactively with the `tee` command. The `tee` file can be disabled interactively with the `notee` command. Executing `tee` again re-enables logging. Without a parameter, the previous file is used. Note that `tee` flushes query results to the file after each statement, just before `mysql` prints its next prompt.

By using the `--pager` option, it is possible to browse or search query results in interactive mode with Unix programs such as `less`, `more`, or any other similar program. If you specify no value for the option, `mysql` checks the value of the `PAGER` environment variable and sets the pager to that. Output paging can be enabled interactively with the `pager` command and disabled with `nopager`. The command takes an optional argument; if given, the paging program is set to that. With no argument, the pager is set to the pager that was set on the command line, or `stdout` if no pager was specified.

Output paging works only in Unix because it uses the `popen()` function, which does not exist on Windows. For Windows, the `tee` option can be used instead to save query output, although this is not as convenient as `pager` for browsing output in some situations.

Here are a few tips about the `pager` command:

- You can use it to write to a file and the results go only to the file:

```
mysql> pager cat > /tmp/log.txt
```

You can also pass any options for the program that you want to use as your pager:

```
mysql> pager less -n -i -S
```

- In the preceding example, note the `-S` option. You may find it very useful for browsing wide query results. Sometimes a very wide result set is difficult to read on the screen. The `-S` option to `less` can make the result set much more readable because you can scroll it horizontally using the left-arrow and right-arrow keys. You can also use `-S` interactively within `less` to switch the horizontal-browse mode on and off. For more information, read the `less` manual page:

```
shell> man less
```

- You can specify very complex pager commands for handling query output:

```
mysql> pager cat | tee /dr1/tmp/res.txt \  
          | tee /dr2/tmp/res2.txt | less -n -i -S
```

In this example, the command would send query results to two files in two different directories on

two different filesystems mounted on `/dr1` and `/dr2`, yet still display the results onscreen via `less`.

You can also combine the `tee` and `pager` functions. Have a `tee` file enabled and `pager` set to `less`, and you are able to browse the results using the `less` program and still have everything appended into a file the same time. The difference between the Unix `tee` used with the `pager` command and the `mysql` built-in `tee` command is that the built-in `tee` works even if you do not have the Unix `tee` available. The built-in `tee` also logs everything that is printed on the screen, whereas the Unix `tee` used with `pager` does not log quite that much. Additionally, `tee` file logging can be turned on and off interactively from within `mysql`. This is useful when you want to log some queries to a file, but not others.

From MySQL 4.0.2 on, the default `mysql>` prompt can be reconfigured. The string for defining the prompt can contain the following special sequences:

Option	Description
<code>\v</code>	The server version
<code>\d</code>	The default database
<code>\h</code>	The server host
<code>\p</code>	The current TCP/IP port or socket file
<code>\u</code>	Your username
<code>\U</code>	Your full <code>user_name@host_name</code> account name
<code>\\</code>	A literal <code>'\'</code> backslash character
<code>\n</code>	A newline character
<code>\t</code>	A tab character
<code>\ </code>	A space (a space follows the backslash)
<code>\_</code>	A space
<code>\R</code>	The current time, in 24-hour military time (0-23)
<code>\r</code>	The current time, standard 12-hour time (1-12)
<code>\m</code>	Minutes of the current time
<code>\y</code>	The current year, two digits
<code>\Y</code>	The current year, four digits
<code>\D</code>	The full current date
<code>\s</code>	Seconds of the current time
<code>\w</code>	The current day of the week in three-letter format (Mon, Tue, ...)
<code>\P</code>	am/pm
<code>\o</code>	The current month in numeric format
<code>\O</code>	The current month in three-letter format (Jan, Feb, ...)
<code>\c</code>	A counter that increments for each statement you issue
<code>\S</code>	Semicolon
<code>\'</code>	Single quote
<code>\"</code>	Double quote

`'\'` followed by any other letter just becomes that letter.

If you specify the `prompt` command with no argument, `mysql` resets the prompt to the default of `mysql>`.

You can set the prompt in several ways:

- *Use an environment variable.* You can set the `MYSQL_PS1` environment variable to a prompt string. For example:

```
shell> export MYSQL_PS1="(\\u@\\h) [\\d]> "
```

- *Use a command-line option.* You can set the `--prompt` option on the command line to `mysql`. For example:

```
shell> mysql --prompt="(\\u@\\h) [\\d]> "
(user@host) [database]>
```

- *Use an option file.* You can set the `prompt` option in the `[mysql]` group of any MySQL option file, such as `/etc/my.cnf` or the `.my.cnf` file in your home directory. For example:

```
[mysql]
prompt=(\\u@\\h) [\\d]>\\_
```

In this example, note that the backslashes are doubled. If you set the prompt using the `prompt` option in an option file, it is advisable to double the backslashes when using the special prompt options. There is some overlap in the set of allowable prompt options and the set of special escape sequences that are recognized in option files. (These sequences are listed in [Section 4.3.2, “Using Option Files”](#).) The overlap may cause you problems if you use single backslashes. For example, `\s` is interpreted as a space rather than as the current seconds value. The following example shows how to define a prompt within an option file to include the current time in `HH:MM:SS>` format:

```
[mysql]
prompt="\\r:\\m:\\s> "
```

- *Set the prompt interactively.* You can change your prompt interactively by using the `prompt` (or `\R`) command. For example:

```
mysql> prompt (\\u@\\h) [\\d]>\\_
PROMPT set to '(\\u@\\h) [\\d]>\\_'
(user@host) [database]>
(user@host) [database]> prompt
Returning to default PROMPT of mysql>
mysql>
```

•

### 8.6.3. `mysql` Server-Side Help

```
mysql> help search_string
```

As of MySQL 4.1, if you provide an argument to the `help` command, `mysql` uses it as a search string to access server-side help from the contents of the MySQL Reference Manual. The proper operation of this command requires that the help tables in the `mysql` database be initialized with help topic information (see [Section 5.2.7, “Server-Side Help”](#)).

If there is no match for the search string, the search fails:

```
mysql> help me
```

```
Nothing found
Please try to run 'help contents' for a list of all accessible topics
```

Use `help contents` to see a list of the help categories:

```
mysql> help contents
You asked for help about help category: "Contents"
For more information, type 'help <item>', where <item> is one of the
following categories:
  Account Management
  Administration
  Data Definition
  Data Manipulation
  Data Types
  Functions
  Functions and Modifiers for Use with GROUP BY
  Geographic Features
  Language Structure
  Storage Engines
  Table Maintenance
  Transactions
```

If the search string matches multiple items, `mysql` shows a list of matching topics:

```
mysql> help logs
Many help items for your request exist.
To make a more specific request, please type 'help <item>',
where <item> is one of the following topics:
  SHOW
  SHOW BINARY LOGS
  SHOW ENGINE
  SHOW LOGS
```

Use a topic as the search string to see the help entry for that topic:

```
mysql> help show binary logs
Name: 'SHOW BINARY LOGS'
Description:
Syntax:
SHOW BINARY LOGS
SHOW MASTER LOGS

Lists the binary log files on the server. This statement is used as
part of the procedure described in [purge-master-logs], that shows how
to determine which logs can be purged.

mysql> SHOW BINARY LOGS;
+-----+-----+
| Log_name          | File_size |
+-----+-----+
| binlog.000015     | 724935   |
| binlog.000016     | 733481   |
+-----+-----+
```

## 8.6.4. Executing SQL Statements from a Text File

The `mysql` client typically is used interactively, like this:

```
shell> mysql db_name
```

However, it is also possible to put your SQL statements in a file and then tell `mysql` to read its input from that file. To do so, create a text file `text_file` that contains the statements you wish to execute. Then invoke `mysql` as shown here:

```
shell> mysql db_name < text_file
```

If you place a `USE db_name` statement as the first statement in the file, it is unnecessary to specify the

database name on the command line:

```
shell> mysql < text_file
```

If you are already running `mysql`, you can execute an SQL script file using the `source` or `\.` command:

```
mysql> source file_name
mysql> \. file_name
```

Sometimes you may want your script to display progress information to the user. For this you can insert statements like this:

```
SELECT '<info_to_display>' AS ' ';
```

The statement shown outputs `<info_to_display>`.

For more information about batch mode, see [Section 3.5, “Using `mysql` in Batch Mode”](#).

## 8.6.5. `mysql` Tips

This section describes some techniques that can help you use `mysql` more effectively.

### 8.6.5.1. Displaying Query Results Vertically

Some query results are much more readable when displayed vertically, instead of in the usual horizontal table format. Queries can be displayed vertically by terminating the query with `\G` instead of a semicolon. For example, longer text values that include newlines often are much easier to read with vertical output:

```
mysql> SELECT * FROM mails WHERE LENGTH(txt) < 300 LIMIT 300,1\G
***** 1. row *****
  msg_nro: 3068
    date: 2000-03-01 23:29:50
time_zone: +0200
mail_from: Monty
      reply: monty@no.spam.com
  mail_to: "Thimble Smith" <tim@no.spam.com>
      sbj: UTF-8
      txt: >>>> "Thimble" == Thimble Smith writes:

Thimble> Hi. I think this is a good idea. Is anyone familiar
Thimble> with UTF-8 or Unicode? Otherwise, I'll put this on my
Thimble> TODO list and see what happens.

Yes, please do that.

Regards,
Monty
  file: inbox-jani-1
  hash: 190402944
1 row in set (0.09 sec)
```

### 8.6.5.2. Using the `--safe-updates` Option

For beginners, a useful startup option is `--safe-updates` (or `--i-am-a-dummy`, which has the same effect). This option was introduced in MySQL 3.23.11. It is helpful for cases when you might have issued a `DELETE FROM tbl_name` statement but forgotten the `WHERE` clause. Normally, such a statement deletes all rows from the table. With `--safe-updates`, you can delete rows only by specifying the key values that identify them. This helps prevent accidents.

When you use the `--safe-updates` option, `mysql` issues the following statement when it connects to the MySQL server:

```
SET SQL_SAFE_UPDATES=1,SQL_SELECT_LIMIT=1000, SQL_MAX_JOIN_SIZE=1000000;
```

See [Section 13.5.3](#), “[SET Syntax](#)”.

The `SET` statement has the following effects:

- You are not allowed to execute an `UPDATE` or `DELETE` statement unless you specify a key constraint in the `WHERE` clause or provide a `LIMIT` clause (or both). For example:

```
UPDATE tbl_name SET not_key_column=val WHERE key_column=val;
UPDATE tbl_name SET not_key_column=val LIMIT 1;
```

- The server limits all large `SELECT` results to 1,000 rows unless the statement includes a `LIMIT` clause.
- The server aborts multiple-table `SELECT` statements that probably need to examine more than 1,000,000 row combinations.

To specify limits different from 1,000 and 1,000,000, you can override the defaults by using the `--select_limit` and `--max_join_size` options:

```
shell> mysql --safe-updates --select_limit=500 --max_join_size=10000
```

### 8.6.5.3. Disabling `mysql` Auto-Reconnect

If the `mysql` client loses its connection to the server while sending a query, it immediately and automatically tries to reconnect once to the server and send the query again. However, even if `mysql` succeeds in reconnecting, your first connection has ended and all your previous session objects and settings are lost: temporary tables, the autocommit mode, and user-defined and session variables. Also, any current transaction rolls back. This behavior may be dangerous for you, as in the following example where the server was shut down and restarted without you knowing it:

```
mysql> SET @a=1;
Query OK, 0 rows affected (0.05 sec)

mysql> INSERT INTO t VALUES(@a);
ERROR 2006: MySQL server has gone away
No connection. Trying to reconnect...
Connection id: 1
Current database: test

Query OK, 1 row affected (1.30 sec)

mysql> SELECT * FROM t;
+-----+
| a     |
+-----+
| NULL |
+-----+
1 row in set (0.05 sec)
```

The `@a` user variable has been lost with the connection, and after the reconnection it is undefined. If it is important to have `mysql` terminate with an error if the connection has been lost, you can start the `mysql` client with the `--skip-reconnect` option.

## 8.7. `mysql_explain_log` — Use `EXPLAIN` on Statements in Query Log

`mysql_explain_log` reads its standard input for query log contents. It uses `EXPLAIN` to analyze `SELECT` statements found in the input. `UPDATE` statements are rewritten to `SELECT` statements and also analyzed with `EXPLAIN`. `mysql_explain_log` then displays a summary of its results.

The results may assist you in determining which queries result in table scans and where it would be beneficial to add indexes to your tables.

Invoke `mysql_explain_log` like this, where `log_file` contains all or part of a MySQL query log:

```
shell> mysql_explain_log [options] < log_file
```

`mysql_explain_log` understands the following options:

- `--date=YYMMDD, -d YYMMDD`  
Select entries from the log only for the given date.
- `--host=host_name, -h host_name`  
Connect to the MySQL server on the given host.
- `--password=password, -ppassword`  
The password to use when connecting to the server.  
Specifying a password on the command line should be considered insecure. See [Section 5.8.6, “Keeping Your Password Secure”](#).
- `--printerror=1, -e 1`  
Enable error output.
- `--socket=path, -S path`  
For connections to `localhost`, the Unix socket file to use, or, on Windows, the name of the named pipe to use.
- `--user=user_name, -u user_name`  
The MySQL username to use when connecting to the server.

## 8.8. `mysqlaccess` — Client for Checking Access Privileges

`mysqlaccess` is a diagnostic tool that Yves Carlier has provided for the MySQL distribution. It checks the access privileges for a hostname, username, and database combination. Note that `mysqlaccess` checks access using only the `user`, `db`, and `host` tables. It does not check table, column, or routine privileges specified in the `tables_priv`, `columns_priv`, or `procs_priv` tables.

Invoke `mysqlaccess` like this:

```
shell> mysqlaccess [host_name [user_name [db_name]]] [options]
```

`mysqlaccess` understands the following options:



- `--help, -?`  
Display a help message and exit.
- `--brief, -b`  
Generate reports in single-line tabular format.
- `--commit`  
Copy the new access privileges from the temporary tables to the original grant tables. The grant tables must be flushed for the new privileges to take effect. (For example, execute a [mysqladmin reload](#) command.)
- `--copy`  
Reload the temporary grant tables from original ones.
- `--db=db_name, -d db_name`  
Specify the database name.
- `--debug=N`  
Specify the debug level. *N* can be an integer from 0 to 3.
- `--host=host_name, -h host_name`  
The hostname to use in the access privileges.
- `--howto`  
Display some examples that show how to use [mysqlaccess](#).
- `--old_server`  
Assume that the server is an old MySQL server (before MySQL 3.21) that does not yet know how to handle full [WHERE](#) clauses.
- `--password[=password], -p[password]`  
The password to use when connecting to the server. If you omit the *password* value following the `--password` or `-p` option on the command line, you are prompted for one.  
  
Specifying a password on the command line should be considered insecure. See [Section 5.8.6, “Keeping Your Password Secure”](#).
- `--plan`  
Display suggestions and ideas for future releases.
- `--preview`  
Show the privilege differences after making changes to the temporary grant tables.
- `--relnotes`  
Display the release notes.
- `--rhost=host_name, -H host_name`

Connect to the MySQL server on the given host.

- `--rollback`

Undo the most recent changes to the temporary grant tables.

- `--spassword[=password], -P[password]`

The password to use when connecting to the server as the superuser. If you omit the *password* value following the `--password` or `-p` option on the command line, you are prompted for one.

Specifying a password on the command line should be considered insecure. See [Section 5.8.6, “Keeping Your Password Secure”](#).

- `--superuser=user_name, -U user_name`

Specify the username for connecting as the superuser.

- `--table, -t`

Generate reports in table format.

- `--user=user_name, -u user_name`

The username to use in the access privileges.

- `--version, -v`

Display version information and exit.

If your MySQL distribution is installed in some non-standard location, you must change the location where `mysqlaccess` expects to find the `mysql` client. Edit the `mysqlaccess` script at approximately line 18. Search for a line that looks like this:

```
$MYSQL = '/usr/local/bin/mysql'; # path to mysql executable
```

Change the path to reflect the location where `mysql` actually is stored on your system. If you do not do this, a [Broken pipe](#) error will occur when you run `mysqlaccess`.

## 8.9. `mysqladmin` — Client for Administering a MySQL Server

`mysqladmin` is a client for performing administrative operations. You can use it to check the server's configuration and current status, to create and drop databases, and more.

Invoke `mysqladmin` like this:

```
shell> mysqladmin [options] command [command-arg] [command [command-arg]] ...
```

`mysqladmin` supports the commands described in the following list. Some of the commands take an argument following the command name.

- `create db_name`

Create a new database named *db\_name*.

- `debug`

Tell the server to write debug information to the error log.

- `drop db_name`

Delete the database named *db\_name* and all its tables.

- `extended-status`

Display the server status variables and their values.

- `flush-hosts`

Flush all information in the host cache.

- `flush-logs`

Flush all logs.

- `flush-privileges`

Reload the grant tables (same as `reload`).

- `flush-status`

Clear status variables.

- `flush-tables`

Flush all tables.

- `flush-threads`

Flush the thread cache. (Added in MySQL 3.23.16.)

- `kill id,id,...`

Kill server threads. If multiple thread ID values are given, there must be no spaces in the list.

- `old-password new-password`

This is like the `password` command but stores the password using the old (pre-4.1) password-hashing format. This command was added in MySQL 4.1.0. (See [Section 5.7.9, “Password Hashing as of MySQL 4.1”](#).)

- `password new-password`

Set a new password. This changes the password to *new-password* for the account that you use with `mysqladmin` for connecting to the server. Thus, the next time you invoke `mysqladmin` (or any other client program) using the same account, you will need to specify the new password.

If the *new-password* value contains spaces or other characters that are special to your command interpreter, you need to enclose it within quotes. On Windows, be sure to use double quotes rather than single quotes; single quotes are not stripped from the password, but rather are interpreted as part of the password. For example:

```
shell> mysqladmin password "my new password"
```

- `ping`

Check whether the server is alive. The return status from `mysqladmin` is 0 if the server is running, 1 if it is not. Beginning with MySQL 4.0.22, the status is 0 even in case of an error such as `Access denied`, because that means the server is running but refused the connection, which is different from the server not running.

- `processlist`

Show a list of active server threads. This is like the output of the `SHOW PROCESSLIST` statement. If the `--verbose` option is given, the output is like that of `SHOW FULL PROCESSLIST`.

- `reload`

Reload the grant tables.

- `refresh`

Flush all tables and close and open log files.

- `shutdown`

Stop the server.

- `start-slave`

Start replication on a slave server. (Added in MySQL 3.23.16.)

- `status`

Display a short server status message.

- `stop-slave`

Stop replication on a slave server. (Added in MySQL 3.23.16.)

- `variables`

Display the server system variables and their values.

- `version`

Display version information from the server.

All commands can be shortened to any unique prefix. For example:

```
shell> mysqladmin proc stat
+-----+-----+-----+-----+-----+-----+-----+-----+
| Id | User | Host | db | Command | Time | State | Info |
+-----+-----+-----+-----+-----+-----+-----+-----+
| 51 | monty | localhost | | Query | 0 | | show processlist |
+-----+-----+-----+-----+-----+-----+-----+-----+
Uptime: 1473624 Threads: 1 Questions: 39487
Slow queries: 0 Opens: 541 Flush tables: 1
Open tables: 19 Queries per second avg: 0.0268
```

The `mysqladmin status` command result displays the following values:

- `Uptime`  
The number of seconds the MySQL server has been running.
- `Threads`  
The number of active threads (clients).
- `Questions`  
The number of questions (queries) from clients since the server was started.
- `Slow queries`  
The number of queries that have taken more than `long_query_time` seconds. See [Section 5.11.5, “The Slow Query Log”](#).
- `Opens`  
The number of tables the server has opened.
- `Flush tables`  
The number of `flush-*`, `refresh`, and `reload` commands the server has executed.
- `Open tables`  
The number of tables that currently are open.
- `Memory in use`  
The amount of memory allocated directly by `mysqld`. This value is displayed only when MySQL has been compiled with `--with-debug=full`.
- `Maximum memory used`  
The maximum amount of memory allocated directly by `mysqld`. This value is displayed only when MySQL has been compiled with `--with-debug=full`.

If you execute `mysqladmin shutdown` when connecting to a local server using a Unix socket file, `mysqladmin` waits until the server's process ID file has been removed, to ensure that the server has stopped properly.

`mysqladmin` supports the following options:

- `--help, -?`  
Display a help message and exit.
- `--character-sets-dir=path`  
The directory where character sets are installed. See [Section 5.10.1, “The Character Set Used for Data and Sorting”](#).
- `--compress, -C`  
Compress all information sent between the client and the server if both support compression.

- `--count=N, -c N`

The number of iterations to make for repeated command execution. This works only with the `--sleep` option.

- `--debug[=debug_options], -# [debug_options]`

Write a debugging log. The `debug_options` string often is `'d:t:o,file_name'`. The default is `'d:t:o,/tmp/mysqladmin.trace'`.

- `--default-character-set=charset_name`

Use `charset_name` as the default character set. See [Section 5.10.1, “The Character Set Used for Data and Sorting”](#). Added in MySQL 4.1.9.

- `--force, -f`

Do not ask for confirmation for the `drop db_name` command. With multiple commands, continue even if an error occurs.

- `--host=host_name, -h host_name`

Connect to the MySQL server on the given host.

- `--password[=password], -p[password]`

The password to use when connecting to the server. If you use the short option form (`-p`), you *cannot* have a space between the option and the password. If you omit the `password` value following the `--password` or `-p` option on the command line, you are prompted for one.

Specifying a password on the command line should be considered insecure. See [Section 5.8.6, “Keeping Your Password Secure”](#).

- `--port=port_num, -P port_num`

The TCP/IP port number to use for the connection.

- `--protocol={TCP|SOCKET|PIPE|MEMORY}`

The connection protocol to use. Added in MySQL 4.1.

- `--relative, -r`

Show the difference between the current and previous values when used with the `--sleep` option. Currently, this option works only with the `extended-status` command.

- `--silent, -s`

Exit silently if a connection to the server cannot be established.

- `--sleep=delay, -i delay`

Execute commands repeatedly, sleeping for `delay` seconds in between. The `--count` option determines the number of iterations.

- `--socket=path, -S path`

For connections to `localhost`, the Unix socket file to use, or, on Windows, the name of the named pipe to use.

- `--ssl*`

Options that begin with `--ssl` specify whether to connect to the server via SSL and indicate where to find SSL keys and certificates. See [Section 5.8.7.3, “SSL Command Options”](#).

- `--user=user_name, -u user_name`

The MySQL username to use when connecting to the server.

- `--verbose, -v`

Verbose mode. Print more information about what the program does.

- `--version, -V`

Display version information and exit.

- `--vertical, -E`

Print output vertically. This is similar to `--relative`, but prints output vertically.

- `--wait[=count], -w[count]`

If the connection cannot be established, wait and retry instead of aborting. If a *count* value is given, it indicates the number of times to retry. The default is one time.

You can also set the following variables by using `--var_name=value` syntax:

- `connect_timeout`

The maximum number of seconds before connection timeout. The default value is 43200 (12 hours).

- `shutdown_timeout`

The maximum number of seconds to wait for server shutdown. The default value is 3600 (1 hour).

It is also possible to set variables by using `--set-variable=var_name=value` or `-O var_name=value` syntax. However, this syntax is deprecated as of MySQL 4.0.

## 8.10. `mysqlbinlog` — Utility for Processing Binary Log Files

The binary log files that the server generates are written in binary format. To examine these files in text format, use the `mysqlbinlog` utility. It is available as of MySQL 3.23.14. You can also use `mysqlbinlog` to read relay log files written by a slave server in a replication setup. Relay logs have the same format as binary log files.

Invoke `mysqlbinlog` like this:

```
shell> mysqlbinlog [options] log_file ...
```

For example, to display the contents of the binary log file named `binlog.000003`, use this command:

```
shell> mysqlbinlog binlog.000003
```

The output includes all events contained in `binlog.000003`. Event information includes the statement executed, the time the statement took, the thread ID of the client that issued it, the timestamp when it was executed, and so forth.

The output from `mysqlbinlog` can be re-executed (for example, by using it as input to `mysql`) to re-apply the statements in the log. This is useful for recovery operations after a server crash. For other usage examples, see the discussion later in this section.

Normally, you use `mysqlbinlog` to read binary log files directly and apply them to the local MySQL server. It is also possible to read binary logs from a remote server by using the `--read-from-remote-server` option. When you read remote binary logs, the connection parameter options can be given to indicate how to connect to the server. These options are `--host`, `--password`, `--port`, `--protocol`, `--socket`, and `--user`; they are ignored except when you also use the `--read-from-remote-server` option.

Binary logs and relay logs are discussed further in [Section 5.11.4, “The Binary Log”](#), and [Section 6.3.4, “Replication Relay and Status Files”](#).

`mysqlbinlog` supports the following options:

- `--help, -?`  
Display a help message and exit.
- `--character-sets-dir=path`  
The directory where character sets are installed. See [Section 5.10.1, “The Character Set Used for Data and Sorting”](#).
- `--database=db_name, -d db_name`  
List entries for just this database (local log only). You can only specify one database with this option - if you specify multiple `--database` options, only the last one is used. This option forces `mysqlbinlog` to output entries from the binary log where the default database (that is, the one selected by `USE`) is `db_name`. Note that this does not replicate cross-database statements such as `UPDATE some_db.some_table SET foo='bar'` while having selected a different database or no database.
- `--debug[=debug_options], -# [debug_options]`  
Write a debugging log. A typical `debug_options` string is often `'d:t:o,file_name'`.
- `--disable-log-bin, -D`  
Disable binary logging. This is useful for avoiding an endless loop if you use the `--to-last-log` option and are sending the output to the same MySQL server. This option also is useful when restoring after a crash to avoid duplication of the statements you have logged. This option is available as of MySQL 4.1.8.  
  
This option requires that you have the `SUPER` privilege. It causes `mysqlbinlog` to include a `SET SQL_LOG_BIN=0` statement in its output to disable binary logging of the remaining output. The `SET` statement is ineffective unless you have the `SUPER` privilege.
- `--force-read, -f`  
With this option, if `mysqlbinlog` reads a binary log event that it does not recognize, it prints a warning, ignores the event, and continues. Without this option, `mysqlbinlog` stops if it reads such an event.



- `--host=host_name, -h host_name`

Get the binary log from the MySQL server on the given host.

- `--local-load=path, -l path`

Prepare local temporary files for `LOAD DATA INFILE` in the specified directory.

- `--offset=N, -o N`

Skip the first `N` entries in the log.

- `--password[=password], -p[password]`

The password to use when connecting to the server. If you use the short option form (`-p`), you *cannot* have a space between the option and the password. If you omit the `password` value following the `--password` or `-p` option on the command line, you are prompted for one.

Specifying a password on the command line should be considered insecure. See [Section 5.8.6, “Keeping Your Password Secure”](#).

- `--port=port_num, -P port_num`

The TCP/IP port number to use for connecting to a remote server.

- `--position=N, -j N`

Deprecated. Use `--start-position` instead (starting from MySQL 4.1.4).

- `--protocol={TCP|SOCKET|PIPE|MEMORY}`

The connection protocol to use. Added in MySQL 4.1.

- `--read-from-remote-server, -R`

Read the binary log from a MySQL server rather than reading a local log file. Any connection parameter options are ignored unless this option is given as well. These options are `--host`, `--password`, `--port`, `--protocol`, `--socket`, and `--user`.

- `--result-file=name, -r name`

Direct output to the given file.

- `--set-charset=charset_name`

Add a `SET NAMES charset_name` statement to the output to specify the character set to be used for processing log files. This option was added in MySQL 4.1.21.

- `--short-form, -s`

Display only the statements contained in the log, without any extra information.

- `--socket=path, -S path`

For connections to `localhost`, the Unix socket file to use, or, on Windows, the name of the named pipe to use.

- `--start-datetime=datetime`

Start reading the binary log at the first event having a timestamp equal to or later than the `date-`

*time* argument. The *datetime* value is relative to the local time zone on the machine where you run `mysqlbinlog`. The value should be in a format accepted for the `DATETIME` or `TIMESTAMP` data types. For example:

```
shell> mysqlbinlog --start-datetime="2005-12-25 11:25:56" binlog.000003
```

This option is available as of MySQL 4.1.4. It is useful for point-in-time recovery. See [Section 5.9.2, “Example Backup and Recovery Strategy”](#).

- `--stop-datetime=datetime`

Stop reading the binary log at the first event having a timestamp equal or posterior to the *datetime* argument. This option is useful for point-in-time recovery. See the description of the `--start-datetime` option for information about the *datetime* value. This option is available as of MySQL 4.1.4.

- `--start-position=N`

Start reading the binary log at the first event having a position equal to the *N* argument. This option applies to the first log file named on the command line. Available as of MySQL 4.1.4 (previously named `--position`).

- `--stop-position=N`

Stop reading the binary log at the first event having a position equal or greater than the *N* argument. This option applies to the last log file named on the command line. Available as of MySQL 4.1.4.

- `--to-last-log, -t`

Do not stop at the end of the requested binary log from a MySQL server, but rather continue printing until the end of the last binary log. If you send the output to the same MySQL server, this may lead to an endless loop. This option requires `--read-from-remote-server`. Available as of MySQL 4.1.2.

- `--user=user_name, -u user_name`

The MySQL username to use when connecting to a remote server.

- `--version, -V`

Display version information and exit.

You can also set the following variable by using `--var_name=value` syntax:

- `open_files_limit`

Specify the number of open file descriptors to reserve.

It is also possible to set variables by using `--set-variable=var_name=value` or `-O var_name=value` syntax. *This syntax is deprecated.*

You can pipe the output of `mysqlbinlog` into the `mysql` client to execute the statements contained in the binary log. This is used to recover from a crash when you have an old backup (see [Section 5.9.1, “Database Backups”](#)). For example:

```
shell> mysqlbinlog binlog.000001 | mysql
```

Or:

```
shell> mysqlbinlog binlog.[0-9]* | mysql
```

You can also redirect the output of `mysqlbinlog` to a text file instead, if you need to modify the statement log first (for example, to remove statements that you do not want to execute for some reason). After editing the file, execute the statements that it contains by using it as input to the `mysql` program.

`mysqlbinlog` has the `--start-position` option, which prints only those statements with an offset in the binary log greater than or equal to a given position (the given position must match the start of one event). It also has options to stop and start when it sees an event with a given date and time. This enables you to perform point-in-time recovery using the `--stop-datetime` option (to be able to say, for example, “roll forward my databases to how they were today at 10:30 a.m.”).

If you have more than one binary log to execute on the MySQL server, the safe method is to process them all using a single connection to the server. Here is an example that demonstrates what may be *unsafe*:

```
shell> mysqlbinlog binlog.000001 | mysql # DANGER!!
shell> mysqlbinlog binlog.000002 | mysql # DANGER!!
```

Processing binary logs this way using different connections to the server causes problems if the first log file contains a `CREATE TEMPORARY TABLE` statement and the second log contains a statement that uses the temporary table. When the first `mysql` process terminates, the server drops the temporary table. When the second `mysql` process attempts to use the table, the server reports “unknown table.”

To avoid problems like this, use a *single* connection to execute the contents of all binary logs that you want to process. Here is one way to do so:

```
shell> mysqlbinlog binlog.000001 binlog.000002 | mysql
```

Another approach is to write all the logs to a single file and then process the file:

```
shell> mysqlbinlog binlog.000001 > /tmp/statements.sql
shell> mysqlbinlog binlog.000002 >> /tmp/statements.sql
shell> mysql -e "source /tmp/statements.sql"
```

In MySQL 3.23, the binary log did not contain the data to load for `LOAD DATA INFILE` statements. To execute such a statement from a binary log file, the original data file was needed. Starting from MySQL 4.0.14, the binary log does contain the data, so `mysqlbinlog` can produce output that reproduces the `LOAD DATA INFILE` operation without the original data file. `mysqlbinlog` copies the data to a temporary file and writes a `LOAD DATA LOCAL INFILE` statement that refers to the file. The default location of the directory where these files are written is system-specific. To specify a directory explicitly, use the `--local-load` option.

Because `mysqlbinlog` converts `LOAD DATA INFILE` statements to `LOAD DATA LOCAL INFILE` statements (that is, it adds `LOCAL`), both the client and the server that you use to process the statements must be configured to allow `LOCAL` capability. See [Section 5.6.4, “Security Issues with `LOAD DATA LOCAL`”](#).

**Warning:** The temporary files created for `LOAD DATA LOCAL` statements are *not* automatically deleted because they are needed until you actually execute those statements. You should delete the temporary files yourself after you no longer need the statement log. The files can be found in the temporary file directory and have names like `original_file_name-#-#`.

Before MySQL 4.1, `mysqlbinlog` could not prepare output suitable for `mysql` if the binary log contained interlaced statements originating from different clients that used temporary tables of the same name. This is fixed in MySQL 4.1. However, the problem still existed for `LOAD DATA INFILE` state-

ments until it was fixed in MySQL 4.1.8.

## 8.11. `mysqlcheck` — A Table Maintenance and Repair Program

The `mysqlcheck` client checks, repairs, optimizes, and analyzes tables. `mysqlcheck` is available as of MySQL 3.23.38.

`mysqlcheck` is similar in function to `myisamchk`, but works differently. The main operational difference is that `mysqlcheck` must be used when the `mysqld` server is running, whereas `myisamchk` should be used when it is not. The benefit of using `mysqlcheck` is that you do not have to stop the server to check or repair your tables.

`mysqlcheck` uses the SQL statements `CHECK TABLE`, `REPAIR TABLE`, `ANALYZE TABLE`, and `OPTIMIZE TABLE` in a convenient way for the user. It determines which statements to use for the operation you want to perform, and then sends the statements to the server to be executed. For details about which storage engines each statement works with, see the descriptions for those statements in [Chapter 13, SQL Statement Syntax](#).

The `MyISAM` storage engine supports all four statements, so `mysqlcheck` can be used to perform all four operations on `MyISAM` tables. Other storage engines do not necessarily support all operations. In such cases, an error message is displayed. For example, if `test.t` is a `MEMORY` table, an attempt to check it produces this result:

```
shell> mysqlcheck test t
test.t
note      : The storage engine for the table doesn't support check
```

There are three general ways to invoke `mysqlcheck`:

```
shell> mysqlcheck [options] db_name [tables]
shell> mysqlcheck [options] --databases db_name1 [db_name2 db_name3...]
shell> mysqlcheck [options] --all-databases
```

If you do not name any tables following `db_name` or if you use the `--databases` or `--all-databases` option, entire databases are checked.

`mysqlcheck` has a special feature compared to other client programs. The default behavior of checking tables (`--check`) can be changed by renaming the binary. If you want to have a tool that repairs tables by default, you should just make a copy of `mysqlcheck` named `mysqlrepair`, or make a symbolic link to `mysqlcheck` named `mysqlrepair`. If you invoke `mysqlrepair`, it repairs tables.

The following names can be used to change `mysqlcheck` default behavior:

<code>mysqlrepair</code>	The default option is <code>--repair</code>
<code>mysqlanalyze</code>	The default option is <code>--analyze</code>
<code>mysqloptimize</code>	The default option is <code>--optimize</code>

`mysqlcheck` supports the following options:

- `--help, -?`

Display a help message and exit.

- `--all-databases, -A`

Check all tables in all databases. This is the same as using the `--databases` option and naming all the databases on the command line.
- `--all-in-1, -1`

Instead of issuing a statement for each table, execute a single statement for each database that names all the tables from that database to be processed.
- `--analyze, -a`

Analyze the tables.
- `--auto-repair`

If a checked table is corrupted, automatically fix it. Any necessary repairs are done after all tables have been checked.
- `--character-sets-dir=path`

The directory where character sets are installed. See [Section 5.10.1, “The Character Set Used for Data and Sorting”](#).
- `--check, -c`

Check the tables for errors. This is the default operation.
- `--check-only-changed, -C`

Check only tables that have changed since the last check or that have not been closed properly.
- `--compress`

Compress all information sent between the client and the server if both support compression.
- `--databases, -B`

Process all tables in the named databases. Normally, `mysqlcheck` treats the first name argument on the command line as a database name and following names as table names. With this option, it treats all name arguments as database names.
- `--debug[=debug_options], -# [debug_options]`

Write a debugging log. A typical `debug_options` string is often `'d:t:o,file_name'`.
- `--default-character-set=charset_name`

Use `charset_name` as the default character set. See [Section 5.10.1, “The Character Set Used for Data and Sorting”](#).
- `--extended, -e`

If you are using this option to check tables, it ensures that they are 100% consistent but takes a long time.

If you are using this option to repair tables, it runs an extended repair that may not only take a long time to execute, but may produce a lot of garbage rows also!
- `--fast, -F`

Check only tables that have not been closed properly.

- `--force, -f`

Continue even if an SQL error occurs.

- `--host=host_name, -h host_name`

Connect to the MySQL server on the given host.

- `--medium-check, -m`

Do a check that is faster than an `--extended` operation. This finds only 99.99% of all errors, which should be good enough in most cases.

- `--optimize, -o`

Optimize the tables.

- `--password[=password], -p[password]`

The password to use when connecting to the server. If you use the short option form (`-p`), you *cannot* have a space between the option and the password. If you omit the `password` value following the `--password` or `-p` option on the command line, you are prompted for one.

Specifying a password on the command line should be considered insecure. See [Section 5.8.6, “Keeping Your Password Secure”](#).

- `--port=port_num, -P port_num`

The TCP/IP port number to use for the connection.

- `--protocol={TCP|SOCKET|PIPE|MEMORY}`

The connection protocol to use. This option was added in MySQL 4.1.

- `--quick, -q`

If you are using this option to check tables, it prevents the check from scanning the rows to check for incorrect links. This is the fastest check method.

If you are using this option to repair tables, it tries to repair only the index tree. This is the fastest repair method.

- `--repair, -r`

Perform a repair that can fix almost anything except unique keys that are not unique.

- `--silent, -s`

Silent mode. Print only error messages.

- `--socket=path, -S path`

For connections to `localhost`, the Unix socket file to use, or, on Windows, the name of the named pipe to use.

- `--ssl*`

Options that begin with `--ssl` specify whether to connect to the server via SSL and indicate where to find SSL keys and certificates. See [Section 5.8.7.3, “SSL Command Options”](#).

- `--tables`  
Overrides the `--databases` or `-B` option. All name arguments following the option are regarded as table names.
- `--use-frm`  
For repair operations on [MyISAM](#) tables, get the table structure from the `.frm` file so that the table can be repaired even if the `.MYI` header is corrupted. This option was added in MySQL 4.0.5.
- `--user=user_name, -u user_name`  
The MySQL username to use when connecting to the server.
- `--verbose, -v`  
Verbose mode. Print information about the various stages of program operation.
- `--version, -V`  
Display version information and exit.

## 8.12. `mysqldump` — A Database Backup Program

The `mysqldump` client is a backup program originally written by Igor Romanenko. It can be used to dump a database or a collection of databases for backup or for transferring the data to another SQL server (not necessarily a MySQL server). The dump contains SQL statements to create the table or populate it, or both.

If you are doing a backup on the server, and your tables all are [MyISAM](#) tables, consider using the `mysqlhotcopy` instead because it can accomplish faster backups and faster restores. See [Section 8.13, “mysqlhotcopy — A Database Backup Program”](#).

There are three general ways to invoke `mysqldump`:

```
shell> mysqldump [options] db_name [tables]
shell> mysqldump [options] --databases db_name1 [db_name2 db_name3...]
shell> mysqldump [options] --all-databases
```

If you do not name any tables following `db_name` or if you use the `--databases` or `--all-databases` option, entire databases are dumped.

To get a list of the options your version of `mysqldump` supports, execute `mysqldump --help`.

If you run `mysqldump` without the `--quick` or `--opt` option, `mysqldump` loads the whole result set into memory before dumping the result. This can be a problem if you are dumping a big database. As of MySQL 4.1, `--opt` is enabled by default, but can be disabled with `--skip-opt`.

If you are using a recent copy of the `mysqldump` program to generate a dump to be reloaded into a very old MySQL server, you should not use the `--opt` or `--extended-insert` option. Use `--skip-opt` instead.

Before MySQL 4.1.2, out-of-range numeric values such as `-inf` and `inf`, as well as `NaN` (not-a-number) values are dumped by `mysqldump` as `NULL`. You can see this using the following

sample table:

```
mysql> CREATE TABLE t (f DOUBLE);
mysql> INSERT INTO t VALUES(1e+1111111111111111111);
mysql> INSERT INTO t VALUES(-1e1111111111111111111);
mysql> SELECT f FROM t;
+-----+
| f      |
+-----+
| inf    |
| -inf   |
+-----+
```

For this table, `mysqldump` produces the following data output:

```
--
-- Dumping data for table `t`
--
INSERT INTO t VALUES (NULL);
INSERT INTO t VALUES (NULL);
```

The significance of this behavior is that if you dump and restore the table, the new table has contents that differ from the original contents. This problem is fixed as of MySQL 4.1.2; you cannot insert `inf` in the table, so this `mysqldump` behavior is only relevant when you deal with old servers.

`mysqldump` supports the following options:

- `--help, -?`  
Display a help message and exit.
- `--add-drop-database`  
Add a `DROP DATABASE` statement before each `CREATE DATABASE` statement. Added in MySQL 4.1.13.
- `--add-drop-table`  
Add a `DROP TABLE` statement before each `CREATE TABLE` statement.
- `--add-locks`  
Surround each table dump with `LOCK TABLES` and `UNLOCK TABLES` statements. This results in faster inserts when the dump file is reloaded. See [Section 7.2.13, “Speed of INSERT Statements”](#).
- `--all-databases, -A`  
Dump all tables in all databases. This is the same as using the `--databases` option and naming all the databases on the command line.
- `--allow-keywords`  
Allow creation of column names that are keywords. This works by prefixing each column name with the table name.
- `--character-sets-dir=path`  
The directory where character sets are installed. See [Section 5.10.1, “The Character Set Used for Data and Sorting”](#).
- `--comments, -i`



Write additional information in the dump file such as program version, server version, and host. . This option is enabled by default. To suppress additional, use `--skip-comments`. This option was added in MySQL 4.0.17.

- `--compact`

Produce less verbose output. This option suppresses comments and enables the `-skip-add-drop-table`, `--no-set-names`, `--skip-disable-keys`, and `-skip-add-locks` options. Added in MySQL 4.1.2.

- `--compatible=name`

Produce output that is more compatible with other database systems or with older MySQL servers. The value of `name` can be `ansi`, `mysql323`, `mysql40`, `postgresql`, `oracle`, `mssql`, `db2`, `maxdb`, `no_key_options`, `no_table_options`, or `no_field_options`. To use several values, separate them by commas. These values have the same meaning as the corresponding options for setting the server SQL mode. See [Section 5.2.5, “SQL Modes”](#).

This option does not guarantee compatibility with other servers. It only enables those SQL mode values that are currently available for making dump output more compatible. For example, `--compatible=oracle` does not map data types to Oracle types or use Oracle comment syntax.

*This option requires a server version of 4.1.0 or higher. With older servers, it does nothing.*

- `--complete-insert, -c`

Use complete `INSERT` statements that include column names.

- `--compress, -C`

Compress all information sent between the client and the server if both support compression.

- `--create-options`

Include all MySQL-specific table options in the `CREATE TABLE` statements. Before MySQL 4.1.2, use `--all` instead.

- `--databases, -B`

Dump several databases. Normally, `mysqldump` treats the first name argument on the command line as a database name and following names as table names. With this option, it treats all name arguments as database names. `CREATE DATABASE` and `USE` statements are included in the output before each new database.

- `--debug[=debug_options], -# [debug_options]`

Write a debugging log. The `debug_options` string is often `'d:t:o,file_name'`. The default is `'d:t:o,/tmp/mysqldump.trace'`.

- `--default-character-set=charset_name`

Use `charset_name` as the default character set. See [Section 5.10.1, “The Character Set Used for Data and Sorting”](#). If not specified, `mysqldump` from MySQL 4.1.2 or later uses `utf8`, and earlier versions use `latin1`.

- `--delayed-insert`

Write `INSERT DELAYED` statements rather than `INSERT` statements.

- `--delete-master-logs`

On a master replication server, delete the binary logs after performing the dump operation. This option automatically enables `--first-slave` before MySQL 4.1.8 and enables `--master-data` thereafter. It was added in MySQL 3.23.57 (for MySQL 3.23) and MySQL 4.0.13 (for MySQL 4.0).

- `--disable-keys, -K`

For each table, surround the `INSERT` statements with `/*!40000 ALTER TABLE tbl_name DISABLE KEYS */;` and `/*!40000 ALTER TABLE tbl_name ENABLE KEYS */;` statements. This makes loading the dump file into a MySQL 4.0 or newer server faster because the indexes are created after all rows are inserted. This option is effective for `MyISAM` tables only.

- `--extended-insert, -e`

Use multiple-row `INSERT` syntax that include several `VALUES` lists. This results in a smaller dump file and speeds up inserts when the file is reloaded.

- `--fields-terminated-by=...`, `--fields-enclosed-by=...`,  
`--fields-optionally-enclosed-by=...`, `--fields-escaped-by=...`,  
`--lines-terminated-by=...`

These options are used with the `-T` option and have the same meaning as the corresponding clauses for `LOAD DATA INFILE`. See [Section 13.2.5, “LOAD DATA INFILE Syntax”](#).

- `--first-slave, -x`

Deprecated. Renamed to `--lock-all-tables` in MySQL 4.1.8.

- `--flush-logs, -F`

Flush the MySQL server log files before starting the dump. This option requires the `RELOAD` privilege. Note that if you use this option in combination with the `--all-databases` (or `-A`) option, the logs are flushed *for each database dumped*. The exception is when using `--lock-all-tables` or `--master-data`: In this case, the logs are flushed only once, corresponding to the moment that all tables are locked. If you want your dump and the log flush to happen at exactly the same moment, you should use `--flush-logs` together with either `--lock-all-tables` or `--master-data`.

- `--force, -f`

Continue even if an SQL error occurs during a table dump.

- `--host=host_name, -h host_name`

Dump data from the MySQL server on the given host. The default host is `localhost`.

- `--hex-blob`

Dump binary columns using hexadecimal notation (for example, `'abc'` becomes `0x616263`). The affected data types are `BINARY`, `VARBINARY`, and `BLOB` in MySQL 4.1 and up, and `CHAR BINARY`, `VARCHAR BINARY`, and `BLOB` in MySQL 4.0. This option was added in MySQL 4.0.23 and 4.1.8.

- `--ignore-table=db_name.tbl_name`

Do not dump the given table, which must be specified using both the database and table names. To ignore multiple tables, use this option multiple times. This option was added in MySQL 4.1.9.

- `--insert-ignore`

Write `INSERT` statements with the `IGNORE` option. This option was added in MySQL 4.1.12.

- `--lock-all-tables, -x`

Lock all tables across all databases. This is achieved by acquiring a global read lock for the duration of the whole dump. This option automatically turns off `--single-transaction` and `--lock-tables`. Added in MySQL 4.1.8.

- `--lock-tables, -l`

Lock all tables before starting the dump. The tables are locked with `READ LOCAL` to allow concurrent inserts in the case of `MyISAM` tables. For transactional tables such as `InnoDB` and `BDB`, `--single-transaction` is a much better option, because it does not need to lock the tables at all.

Please note that when dumping multiple databases, `--lock-tables` locks tables for each database separately. So, this option does not guarantee that the tables in the dump file are logically consistent between databases. Tables in different databases may be dumped in completely different states.

- `--master-data[=value]`

Write the binary log filename and position to the output. This option requires the `RELOAD` privilege and the binary log must be enabled. If the option value is equal to 1, the position and filename are written to the dump output in the form of a `CHANGE MASTER` statement that makes a slave server start from the correct position in the master's binary logs if you use this SQL dump of the master to set up a slave. If the option value is equal to 2, the `CHANGE MASTER` statement is written as an SQL comment. This is the default action if *value* is omitted. *value* may be given as of MySQL 4.1.8; before that, do not specify an option value.

The `--master-data` option turns on `--lock-all-tables`, unless `--single-transaction` also is specified (in which case, a global read lock is only acquired a short time at the beginning of the dump. See also the description for `--single-transaction`. In all cases, any action on logs happens at the exact moment of the dump. This option automatically turns off `--lock-tables`.

- `--no-autocommit`

Enclose the `INSERT` statements for each dumped table within `SET AUTOCOMMIT=0` and `COMMIT` statements.

- `--no-create-db, -n`

This option suppresses the `CREATE DATABASE` statements that are otherwise included in the output if the `--databases` or `--all-databases` option is given.

- `--no-create-info, -t`

Do not write `CREATE TABLE` statements that re-create each dumped table.

- `--no-data, -d`

Do not write any row information for the table. This is very useful if you want to dump only the `CREATE TABLE` statement for the table.

- `--opt`

This option is shorthand; it is the same as specifying `--add-drop-table --add-locks --create-options --disable-keys --extended-insert --lock-tables --quick --set-charset`. It should give you a fast dump operation and produce a dump file that can be reloaded into a MySQL server quickly.

As of MySQL 4.1, `--opt` is enabled by default. To disable the options that it enables, use `--skip-opt`. To disable only certain of the options enabled by `--opt`, use their `--skip` forms; for example, `--skip-add-drop-table` or `--skip-quick`. Alternatively, use `--skip-opt` to disable the options enabled by `--opt`, followed by options to enable the features that you want. Options are processed in order, so the options to enable features must follow `--skip-opt`. For example, `--skip-opt --extended-insert` enables extended inserts, but `--extended-insert --skip-opt` does not.

- `--order-by-primary`

Sorts each table's rows by its primary key, or its first unique index, if such an index exists. This is useful when dumping a `MyISAM` table to be loaded into an `InnoDB` table, but will make the dump itself take considerably longer. This option was added in MySQL 4.1.8.

- `--password[=password], -p[password]`

The password to use when connecting to the server. If you use the short option form (`-p`), you *cannot* have a space between the option and the password. If you omit the `password` value following the `--password` or `-p` option on the command line, you are prompted for one.

Specifying a password on the command line should be considered insecure. See [Section 5.8.6, “Keeping Your Password Secure”](#).

- `--port=port_num, -P port_num`

The TCP/IP port number to use for the connection.

- `--protocol={TCP | SOCKET | PIPE | MEMORY}`

The connection protocol to use. Added in MySQL 4.1.

- `--quick, -q`

This option is useful for dumping large tables. It forces `mysqldump` to retrieve rows for a table from the server a row at a time rather than retrieving the entire row set and buffering it in memory before writing it out.

- `--quote-names, -Q`

Quote database, table, and column names within ‘`’ characters. If the `ANSI_QUOTES` SQL mode is enabled, names are quoted within ‘”’ characters. As of MySQL 4.1.1, `--quote-names` is enabled by default. It can be disabled with `--skip-quote-names`, but this option should be given after any option such as `--compatible` that may enable `--quote-names`.

- `--result-file=file, -r file`

Direct output to a given file. This option should be used on Windows to prevent newline ‘\n’ characters from being converted to ‘\r\n’ carriage return/newline sequences. The result file is created and its contents overwritten, even if an error occurs while generating the dump. The previous contents are lost.

- `--set-charset`

Add `SET NAMES default_character_set` to the output. This option is enabled by default. To suppress the `SET NAMES` statement, use `--skip-set-charset`. This option was added in MySQL 4.1.2.

- `--single-transaction`

This option issues a `BEGIN` SQL statement before dumping data from the server. It is useful only with transactional tables such as `InnoDB` and `BDB`, because then it dumps the consistent state of the database at the time when `BEGIN` was issued without blocking any applications.

When using this option, you should keep in mind that only `InnoDB` tables are dumped in a consistent state. For example, any `MyISAM` or `MEMORY` tables dumped while using this option may still change state.

The `--single-transaction` option was added in MySQL 4.0.2. This option is mutually exclusive with the `--lock-tables` option, because `LOCK TABLES` causes any pending transactions to be committed implicitly.

This option is not supported for MySQL Cluster tables; the results cannot be guaranteed to be consistent due to the fact that the `NDBCluster` storage engine supports only the `READ_COMMITTED` transaction isolation level. You should always use `NDB` backup and restore instead.

To dump big tables, you should combine this option with `--quick`.

- `--skip-opt`

See the description for the `--opt` option.

- `--socket=path, -S path`

For connections to `localhost`, the Unix socket file to use, or, on Windows, the name of the named pipe to use.

- `--skip-comments`

See the description for the `--comments` option.

- `--ssl*`

Options that begin with `--ssl` specify whether to connect to the server via SSL and indicate where to find SSL keys and certificates. See [Section 5.8.7.3, “SSL Command Options”](#).

- `--tab=path, -T path`

Produce tab-separated data files. For each dumped table, `mysqldump` creates a `tbl_name.sql` file that contains the `CREATE TABLE` statement that creates the table, and a `tbl_name.txt` file that contains its data. The option value is the directory in which to write the files.

By default, the `.txt` data files are formatted using tab characters between column values and a newline at the end of each line. The format can be specified explicitly using the `--fields-xxx` and `--lines-xxx` options.

**Note:** This option should be used only when `mysqldump` is run on the same machine as the `mysqld` server. You must have the `FILE` privilege, and the server must have permission to write files in the directory that you specify.

- `--tables`

Override the `--databases` or `-B` option. All name arguments following the option are regarded as table names.

- `--user=user_name, -u user_name`

The MySQL username to use when connecting to the server.

- `--verbose, -v`

Verbose mode. Print more information about what the program does.

- `--version, -V`

Display version information and exit.

- `--where='where_condition', -w 'where_condition'`

Dump only rows selected by the given `WHERE` condition. Note that quotes around the condition are mandatory if it contains spaces or other characters that are special to your command interpreter.

Examples:

```
--where="user='jimf' "  
-w"userid>1"  
-w"userid<1"
```

- `--xml, -X`

Write dump output as well-formed XML.

You can also set the following variables by using `--var_name=value` syntax:

- `max_allowed_packet`

The maximum size of the buffer for client/server communication. The value of the variable can be up to 16MB before MySQL 4.0, and up to 1GB from MySQL 4.0 on.

- `net_buffer_length`

The initial size of the buffer for client/server communication. When creating multiple-row-insert statements (as with option `--extended-insert` or `--opt`), `mysqldump` creates rows up to `net_buffer_length` length. If you increase this variable, you should also ensure that the `net_buffer_length` variable in the MySQL server is at least this large.

It is also possible to set variables by using `--set-variable=var_name=value` or `-O var_name=value` syntax. However, this syntax is deprecated as of MySQL 4.0.

The most common use of `mysqldump` is probably for making a backup of an entire database:

```
shell> mysqldump --opt db_name > backup-file.sql
```

You can read the dump file back into the server like this:

```
shell> mysql db_name < backup-file.sql
```

Or like this:

```
shell> mysql -e "source /path-to-backup/backup-file.sql" db_name
```

`mysqldump` is also very useful for populating databases by copying data from one MySQL server to another:

```
shell> mysqldump --opt db_name | mysql --host=remote_host -C db_name
```

It is possible to dump several databases with one command:

```
shell> mysqldump --databases db_name1 [db_name2 ...] > my_databases.sql
```

To dump all databases, use the `--all-databases` option:

```
shell> mysqldump --all-databases > all_databases.sql
```

For InnoDB tables, `mysqldump` provides a way of making an online backup:

```
shell> mysqldump --all-databases --single-transaction > all_databases.sql
```

This backup just needs to acquire a global read lock on all tables (using `FLUSH TABLES WITH READ LOCK`) at the beginning of the dump. As soon as this lock has been acquired, the binary log coordinates are read and the lock is released. If and only if one long updating statement is running when the `FLUSH` statement is issued, the MySQL server may get stalled until that long statement finishes, and then the dump becomes lock-free. If the update statements that the MySQL server receives are short (in terms of execution time), the initial lock period should not be noticeable, even with many updates.

For point-in-time recovery (also known as “roll-forward,” when you need to restore an old backup and replay the changes that happened since that backup), it is often useful to rotate the binary log (see [Section 5.11.4, “The Binary Log”](#)) or at least know the binary log coordinates to which the dump corresponds:

```
shell> mysqldump --all-databases --master-data=2 > all_databases.sql
```

Or:

```
shell> mysqldump --all-databases --flush-logs --master-data=2  
> all_databases.sql
```

The simultaneous use of `--master-data` and `--single-transaction` works as of MySQL 4.1.8. It provides a convenient way to make an online backup suitable for point-in-time recovery if tables are stored in the InnoDB storage engine.

For more information on making backups, see [Section 5.9.1, “Database Backups”](#), and [Section 5.9.2, “Example Backup and Recovery Strategy”](#).

## 8.13. `mysqlhotcopy` — A Database Backup Program

`mysqlhotcopy` is a Perl script that was originally written and contributed by Tim Bunce. It uses `LOCK TABLES`, `FLUSH TABLES`, and `cp` or `scp` to make a database backup quickly. It is the fastest way to make a backup of the database or single tables, but it can be run only on the same machine where the database directories are located. `mysqlhotcopy` works only for backing up MyISAM and ISAM tables, and ARCHIVE tables as of MySQL 4.1. `mysqlhotcopy` runs on Unix, and as of MySQL 4.0.18 also on NetWare.

```
shell> mysqlhotcopy db_name [/path/to/new_directory]
```

```
shell> mysqlhotcopy db_name_1 ... db_name_n /path/to/new_directory
```

Back up tables in the given database that match a regular expression:

```
shell> mysqlhotcopy db_name ./regex/
```

The regular expression for the table name can be negated by prefixing it with a tilde ('~'):

```
shell> mysqlhotcopy db_name ./~regex/
```

`mysqlhotcopy` supports the following options:

- `--help, -?`  
Display a help message and exit.
- `--addtodest`  
Do not rename target directory (if it exists); merely add files to it. This option was added in MySQL 4.0.13.
- `--allowold`  
Do not abort if a target exists; rename it by adding an `_old` suffix.
- `--checkpoint=db_name.tbl_name`  
Insert checkpoint entries into the specified database `db_name` and table `tbl_name`.
- `--chroot=path`  
Base directory of the `chroot` jail in which `mysqld` operates. The `path` value should match that of the `--chroot` option given to `mysqld`. This option was added in MySQL 4.0.19.
- `--debug`  
Enable debug output.
- `--dryrun, -n`  
Report actions without performing them.
- `--flushlog`  
Flush logs after all tables are locked.
- `--host=host_name, -h host_name`  
The hostname of the local host to use for making a TCP/IP connection to the local server. By default, the connection is made to `localhost` using a Unix socket file.
- `--keepold`  
Do not delete previous (renamed) target when done.



- `--method=command`

The method for copying files (`cp` or `scp`).

- `--noindices`

Do not include full index files in the backup. This makes the backup smaller and faster. The indexes for reloaded tables can be reconstructed later with `myisamchk -rq` for MyISAM tables or `isamchk -rq` for ISAM tables.

- `--password=password, -ppassword`

The password to use when connecting to the server. Note that the password value is not optional for this option, unlike for other MySQL programs. You can use an option file to avoid giving the password on the command line.

Specifying a password on the command line should be considered insecure. See [Section 5.8.6, “Keeping Your Password Secure”](#).

- `--port=port_num, -P port_num`

The TCP/IP port number to use when connecting to the local server.

- `--quiet, -q`

Be silent except for errors.

- `--record_log_pos=db_name.tbl_name`

Record master and slave status in the specified database `db_name` and table `tbl_name`.

- `--regexp=expr`

Copy all databases with names that match the given regular expression.

- `--resetmaster`

Reset the binary log after locking all the tables.

- `--resetslave`

Reset the `master.info` file after locking all the tables.

- `--socket=path, -S path`

The Unix socket file to use for the connection.

- `--suffix=str`

The suffix for names of copied databases.

- `--tmpdir=path`

The temporary directory. The default is `/tmp`.

- `--user=user_name, -u user_name`

The MySQL username to use when connecting to the server.

`mysqlhotcopy` reads the `[client]` and `[mysqlhotcopy]` option groups from option files.

To execute `mysqlhotcopy`, you must have access to the files for the tables that you are backing up, the `SELECT` privilege for those tables, the `RELOAD` privilege (to be able to execute `FLUSH TABLES`), and the `LOCK TABLES` privilege (to be able to lock the tables).

Use `perldoc` for additional `mysqlhotcopy` documentation, including information about the structure of the tables needed for the `--checkpoint` and `--record_log_pos` options:

```
shell> perldoc mysqlhotcopy
```

## 8.14. `mysqlimport` — A Data Import Program

The `mysqlimport` client provides a command-line interface to the `LOAD DATA INFILE` SQL statement. Most options to `mysqlimport` correspond directly to clauses of `LOAD DATA INFILE` syntax. See [Section 13.2.5, “LOAD DATA INFILE Syntax”](#).

Invoke `mysqlimport` like this:

```
shell> mysqlimport [options] db_name textfile1 [textfile2 ...]
```

For each text file named on the command line, `mysqlimport` strips any extension from the filename and uses the result to determine the name of the table into which to import the file's contents. For example, files named `patient.txt`, `patient.text`, and `patient` all would be imported into a table named `patient`.

`mysqlimport` supports the following options:

- `--help, -?`  
Display a help message and exit.
- `--character-sets-dir=path`  
The directory where character sets are installed. See [Section 5.10.1, “The Character Set Used for Data and Sorting”](#).
- `--columns=column_list, -c column_list`  
This option takes a comma-separated list of column names as its value. The order of the column names indicates how to match data file columns with table columns.
- `--compress, -C`  
Compress all information sent between the client and the server if both support compression.
- `--debug[=debug_options], -# [debug_options]`  
Write a debugging log. The `debug_options` string often is `'d:t:o,file_name'`.
- `--default-character-set=charset_name`  
Use `charset_name` as the default character set. See [Section 5.10.1, “The Character Set Used for Data and Sorting”](#).
- `--delete, -D`

Empty the table before importing the text file.

- `--fields-terminated-by=...`, `--fields-enclosed-by=...`, `--fields-optionally-enclosed-by=...`, `--fields-escaped-by=...`, `--lines-terminated-by=...`

These options have the same meaning as the corresponding clauses for `LOAD DATA INFILE`. For example, to import Windows files that have lines terminated with carriage return/linefeed pairs, use `--lines-terminated-by="\r\n"`. (You might have to double the backslashes, depending on the escaping conventions of your command interpreter.) See [Section 13.2.5, “LOAD DATA INFILE Syntax”](#).

- `--force, -f`

Ignore errors. For example, if a table for a text file does not exist, continue processing any remaining files. Without `--force`, `mysqlimport` exits if a table does not exist.

- `--host=host_name, -h host_name`

Import data to the MySQL server on the given host. The default host is `localhost`.

- `--ignore, -i`

See the description for the `--replace` option.

- `--ignore-lines=N`

Ignore the first *N* lines of the data file.

- `--local, -L`

Read input files locally from the client host.

- `--lock-tables, -l`

Lock *all* tables for writing before processing any text files. This ensures that all tables are synchronized on the server.

- `--low-priority`

Use `LOW_PRIORITY` when loading the table.

- `--password[=password], -p[password]`

The password to use when connecting to the server. If you use the short option form (`-p`), you *cannot* have a space between the option and the password. If you omit the `password` value following the `--password` or `-p` option on the command line, you are prompted for one.

Specifying a password on the command line should be considered insecure. See [Section 5.8.6, “Keeping Your Password Secure”](#).

- `--port=port_num, -P port_num`

The TCP/IP port number to use for the connection.

- `--protocol={TCP|SOCKET|PIPE|MEMORY}`

The connection protocol to use. Added in MySQL 4.1.

- `--replace, -r`

The `--replace` and `--ignore` options control handling of input rows that duplicate existing rows on unique key values. If you specify `--replace`, new rows replace existing rows that have the same unique key value. If you specify `--ignore`, input rows that duplicate an existing row on a unique key value are skipped. If you do not specify either option, an error occurs when a duplicate key value is found, and the rest of the text file is ignored.

- `--silent, -s`

Silent mode. Produce output only when errors occur.

- `--socket=path, -S path`

For connections to `localhost`, the Unix socket file to use, or, on Windows, the name of the named pipe to use.

- `--ssl*`

Options that begin with `--ssl` specify whether to connect to the server via SSL and indicate where to find SSL keys and certificates. See [Section 5.8.7.3, “SSL Command Options”](#).

- `--user=user_name, -u user_name`

The MySQL username to use when connecting to the server.

- `--verbose, -v`

Verbose mode. Print more information about what the program does.

- `--version, -V`

Display version information and exit.

Here is a sample session that demonstrates use of `mysqlimport`:

```
shell> mysql -e 'CREATE TABLE impptest(id INT, n VARCHAR(30))' test
shell> ed
a
100      Max Sydow
101      Count Dracula
.
w impptest.txt
32
q
shell> od -c impptest.txt
0000000  1  0  0  \t  M  a  x           S  y  d  o  w  \n  1  0
0000020  1  \t  C  o  u  n  t       D  r  a  c  u  l  a  \n
0000040
shell> mysqlimport --local test impptest.txt
test.impptest: Records: 2 Deleted: 0 Skipped: 0 Warnings: 0
shell> mysql -e 'SELECT * FROM impptest' test
+-----+-----+
| id  | n          |
+-----+-----+
| 100 | Max Sydow  |
| 101 | Count Dracula |
+-----+-----+
```

## 8.15. `mysqlshow` — Display Database, Table, and Column Information

The `mysqlshow` client can be used to quickly see which databases exist, their tables, or a table's columns or indexes.

`mysqlshow` provides a command-line interface to several SQL `SHOW` statements. See [Section 13.5.4, “SHOW Syntax”](#). The same information can be obtained by using those statements directly. For example, you can issue them from the `mysql` client program.

Invoke `mysqlshow` like this:

```
shell> mysqlshow [options] [db_name [tbl_name [col_name]]]
```

- If no database is given, a list of database names is shown.
- If no table is given, all matching tables in the database are shown.
- If no column is given, all matching columns and column types in the table are shown.

The output displays only the names of those databases, tables, or columns for which you have some privileges.

If the last argument contains shell or SQL wildcard characters (`*`, `?`, `%`, or `_`), only those names that are matched by the wildcard are shown. If a database name contains any underscores, those should be escaped with a backslash (some Unix shells require two) to get a list of the proper tables or columns. `*` and `?` characters are converted into SQL `%` and `_` wildcard characters. This might cause some confusion when you try to display the columns for a table with a `_` in the name, because in this case, `mysqlshow` shows you only the table names that match the pattern. This is easily fixed by adding an extra `%` last on the command line as a separate argument.

`mysqlshow` supports the following options:

- `--help, -?`  
Display a help message and exit.
- `--character-sets-dir=path`  
The directory where character sets are installed. See [Section 5.10.1, “The Character Set Used for Data and Sorting”](#).
- `--compress, -C`  
Compress all information sent between the client and the server if both support compression.
- `--debug[=debug_options], -# [debug_options]`  
Write a debugging log. The `debug_options` string often is `'d:t:o,file_name'`.
- `--default-character-set=charset_name`  
Use `charset_name` as the default character set. See [Section 5.10.1, “The Character Set Used for Data and Sorting”](#).
- `--host=host_name, -h host_name`  
Connect to the MySQL server on the given host.

- `--keys, -k`

Show table indexes.

- `--password[=password], -p[password]`

The password to use when connecting to the server. If you use the short option form (`-p`), you *cannot* have a space between the option and the password. If you omit the *password* value following the `--password` or `-p` option on the command line, you are prompted for one.

Specifying a password on the command line should be considered insecure. See [Section 5.8.6, “Keeping Your Password Secure”](#).

Specifying a password on the command line should be considered insecure. See [Section 5.8.6, “Keeping Your Password Secure”](#).

- `--port=port_num, -P port_num`

The TCP/IP port number to use for the connection.

- `--protocol={TCP|SOCKET|PIPE|MEMORY}`

The connection protocol to use. Added in MySQL 4.1.

- `--socket=path, -S path`

For connections to `localhost`, the Unix socket file to use, or, on Windows, the name of the named pipe to use.

- `--ssl*`

Options that begin with `--ssl` specify whether to connect to the server via SSL and indicate where to find SSL keys and certificates. See [Section 5.8.7.3, “SSL Command Options”](#).

- `--status, -i`

Display extra information about each table.

- `--user=user_name, -u user_name`

The MySQL username to use when connecting to the server.

- `--verbose, -v`

Verbose mode. Print more information about what the program does. This option can be used multiple times to increase the amount of information.

- `--version, -V`

Display version information and exit.

## 8.16. `mysql_zap` — Kill Processes That Match a Pattern

`mysql_zap` kills processes that match a pattern. It uses the `ps` command and Unix signals, so it runs on Unix and Unix-like systems.

Invoke `mysql_zap` like this:

```
shell> mysql_zap [-signal] [-?Ift] pattern
```

A process matches if its output line from the `ps` command contains the pattern. By default, `mysql_zap` asks for confirmation for each process. Respond `y` to kill the process, or `q` to exit `mysql_zap`. For any other response, `mysql_zap` does not attempt to kill the process.

If the `-signal` option is given, it specifies the name or number of the signal to send to each process. Otherwise, `mysql_zap` tries first with `TERM` (signal 15) and then with `KILL` (signal 9).

`mysql_zap` understands the following additional options:

- `--help, -?, -I`

Display a help message and exit.

- `-f`

Force mode. `mysql_zap` attempts to kill each process without confirmation.

- `-t`

Test mode. Display information about each process but do not kill it.

## 8.17. `pererror` — Explain Error Codes

For most system errors, MySQL displays, in addition to an internal text message, the system error code in one of the following styles:

```
message ... (errno: #)
message ... (Errcode: #)
```

You can find out what the error code means by examining the documentation for your system or by using the `pererror` utility.

`pererror` prints a description for a system error code or for a storage engine (table handler) error code.

Invoke `pererror` like this:

```
shell> pererror [options] errorcode ...
```

Example:

```
shell> pererror 13 64
Error code 13: Permission denied
Error code 64: Machine is not on the network
```

To obtain the error message for a MySQL Cluster error code, invoke `pererror` with the `--ndb` option:

```
shell> pererror --ndb errorcode
```

Note that the meaning of system error messages may be dependent on your operating system. A given error code may mean different things on different operating systems.

`perror` supports the following options:

- `--help`, `--info`, `-I`, `-?`  
Display a help message and exit.
- `--ndb`  
Print the error message for a MySQL Cluster error code.
- `--silent`, `-s`  
Silent mode. Print only the error message.
- `--verbose`, `-v`  
Verbose mode. Print error code and message. This is the default behavior.
- `--version`, `-V`  
Display version information and exit.

## 8.18. `replace` — A String-Replacement Utility

The `replace` utility program changes strings in place in files or on the standard input.

Invoke `replace` in one of the following ways:

```
shell> replace from to [from to] ... -- file [file] ...  
shell> replace from to [from to] ... < file
```

`from` represents a string to look for and `to` represents its replacement. There can be one or more pairs of strings.

Use the `--` option to indicate where the string-replacement list ends and the filenames begin. In this case, any file named on the command line is modified in place, so you may want to make a copy of the original before converting it. `replace` prints a message indicating which of the input files it actually modifies.

If the `--` option is not given, `replace` reads the standard input and writes to the standard output.

`replace` uses a finite state machine to match longer strings first. It can be used to swap strings. For example, the following command swaps `a` and `b` in the given files, `file1` and `file2`:

```
shell> replace a b b a -- file1 file2 ...
```

The `replace` program is used by `mysql2mysql`. See [Section 17.9.1, “mysql2mysql — Convert mSQL Programs for Use with MySQL”](#).

`replace` supports the following options:

- `-?`, `-I`  
Display a help message and exit.



- `-# debug_options`

Write a debugging log. The `debug_options` string often is `'d:t:o,file_name'`.

- `-s`

Silent mode. Print less information what the program does.

- `-v`

Verbose mode. Print more information about what the program does.

- `-V`

Display version information and exit.

---

# Chapter 9. Language Structure

This chapter discusses the rules for writing the following elements of SQL statements when using MySQL:

- Literal values such as strings and numbers
- Identifiers such as database, table, and column names
- User-defined and system variables
- Comments
- Reserved words

## 9.1. Literal Values

This section describes how to write literal values in MySQL. These include strings, numbers, hexadecimal values, boolean values, and `NULL`. The section also covers the various nuances and “gotchas” that you may run into when dealing with these basic types in MySQL.

### 9.1.1. Strings

A string is a sequence of bytes or characters, enclosed within either single quote (`'`) or double quote (`"`) characters. Examples:

```
'a string'  
"another string"
```

If the `ANSI_QUOTES` SQL mode is enabled, string literals can be quoted only within single quotes because a string quoted within double quotes is interpreted as an identifier.

As of MySQL 4.1.1, a *binary string* is a string of bytes that has no character set or collation. A *non-binary string* is a string of characters that has a character set and collation. For both types of strings, comparisons are based on the numeric values of the string unit. For binary strings, the unit is the byte. For non-binary strings the unit is the character and some character sets allow multi-byte characters. Character value ordering is a function of the string collation.

Also as of MySQL 4.1.1, string literals may have an optional character set introducer and `COLLATE` clause:

```
[_charset_name]'string' [COLLATE collation_name]
```

Examples:

```
SELECT _latin1'string';  
SELECT _latin1'string' COLLATE latin1_danish_ci;
```

For more information about these forms of string syntax, see [Section 10.3.5, “Character String Literal Character Set and Collation”](#).

Within a string, certain sequences have special meaning. Each of these sequences begins with a backslash (`\`), known as the *escape character*. MySQL recognizes the following escape sequences:

<code>\0</code>	An ASCII 0 (NUL) character.
<code>\'</code>	A single quote ( <code>'</code> ) character.
<code>\"</code>	A double quote ( <code>"</code> ) character.
<code>\b</code>	A backspace character.
<code>\n</code>	A newline (linefeed) character.
<code>\r</code>	A carriage return character.
<code>\t</code>	A tab character.
<code>\z</code>	ASCII 26 (Control-Z). See note following the table.
<code>\\</code>	A backslash ( <code>\</code> ) character.
<code>\%</code>	A <code>'%</code> character. See note following the table.
<code>\_</code>	A <code>'_'</code> character. See note following the table.

For all other escape sequences, backslash is ignored. That is, the escaped character is interpreted as if it was not escaped. For example, `\x` is just `x`.

These sequences are case sensitive. For example, `\b` is interpreted as a backspace, but `\B` is interpreted as `B`.

The ASCII 26 character can be encoded as `\z` to enable you to work around the problem that ASCII 26 stands for END-OF-FILE on Windows. ASCII 26 within a file causes problems if you try to use `mysql db_name < file_name`.

The `\%` and `\_` sequences are used to search for literal instances of `'%'` and `'_'` in pattern-matching contexts where they would otherwise be interpreted as wildcard characters. See the description of the `LIKE` operator in [Section 12.3.1, “String Comparison Functions”](#). If you use `\%` or `\_` in non-pattern-matching contexts, they evaluate to the strings `\%` and `\_`, not to `'%'` and `'_'`.

There are several ways to include quote characters within a string:

- A `'` inside a string quoted with `'` may be written as `''`.
- A `"` inside a string quoted with `"` may be written as `""`.
- Precede the quote character by an escape character (`\`).
- A `'` inside a string quoted with `"` needs no special treatment and need not be doubled or escaped. In the same way, `"` inside a string quoted with `'` needs no special treatment.

The following `SELECT` statements demonstrate how quoting and escaping work:

```
mysql> SELECT 'hello', "hello", "'hello'", 'hel'lo', '\hello';
+-----+-----+-----+-----+-----+
| hello | "hello" | "'hello'" | hel'lo | '\hello' |
+-----+-----+-----+-----+-----+

mysql> SELECT "hello", "'hello'", "'hello'", "hel"lo", "\"hello";
+-----+-----+-----+-----+-----+
| hello | 'hello' | "'hello'" | hel"lo | "\"hello" |
+-----+-----+-----+-----+-----+

mysql> SELECT 'This\nIs\nFour\nLines';
+-----+
| This
| Is
| Four
| Lines |
```

```

+-----+
mysql> SELECT 'disappearing\ backslash';
+-----+
| disappearing backslash |
+-----+

```

If you want to insert binary data into a string column (such as a [BLOB](#) column), the following characters must be represented by escape sequences:

<code>NUL</code>	<code>NUL</code> byte (ASCII 0). Represent this character by <code>'\0'</code> (a backslash followed by an ASCII <code>'0'</code> character).
<code>\</code>	Backslash (ASCII 92). Represent this character by <code>'\\'</code> .
<code>'</code>	Single quote (ASCII 39). Represent this character by <code>'\''</code> .
<code>"</code>	Double quote (ASCII 34). Represent this character by <code>'\"'</code> .

When writing application programs, any string that might contain any of these special characters must be properly escaped before the string is used as a data value in an SQL statement that is sent to the MySQL server. You can do this in two ways:

- Process the string with a function that escapes the special characters. In a C program, you can use the `mysql_real_escape_string()` C API function to escape characters. See [Section 17.2.3.51](#), “`mysql_real_escape_string()`”. The Perl DBI interface provides a `quote` method to convert special characters to the proper escape sequences. See [Section 17.4](#), “MySQL Perl API”. Other language interfaces may provide a similar capability.
- As an alternative to explicitly escaping special characters, many MySQL APIs provide a placeholder capability that enables you to insert special markers into a statement string, and then bind data values to them when you issue the statement. In this case, the API takes care of escaping special characters in the values for you.

## 9.1.2. Numbers

Integers are represented as a sequence of digits. Floats use `'.'` as a decimal separator. Either type of number may be preceded by `'-'` or `'+'` to indicate a negative or positive value, respectively

Examples of valid integers:

```

1221
0
-32

```

Examples of valid floating-point numbers:

```

294.42
-32032.6809e+10
148.00

```

An integer may be used in a floating-point context; it is interpreted as the equivalent floating-point number.

## 9.1.3. Hexadecimal Values

MySQL supports hexadecimal values. In numeric contexts, these act like integers (64-bit precision). In string contexts, these act like binary strings, where each pair of hex digits is converted to a character:

```
mysql> SELECT x'4D7953514C';
-> 'MySQL'
mysql> SELECT 0xa+0;
-> 10
mysql> SELECT 0x5061756c;
-> 'Paul'
```

In MySQL 4.1 (and in MySQL 4.0 when using the `--new` option), the default type of a hexadecimal value is a string. If you want to ensure that the value is treated as a number, you can use `CAST(... AS UNSIGNED)`:

```
mysql> SELECT 0x41, CAST(0x41 AS UNSIGNED);
-> 'A', 65
```

The `x'hexstring'` syntax is new in 4.0 and is based on standard SQL. The `0x` syntax is based on ODBC. Hexadecimal strings are often used by ODBC to supply values for `BLOB` columns.

Beginning with MySQL 4.0.1, you can convert a string or a number to a string in hexadecimal format with the `HEX()` function:

```
mysql> SELECT HEX('cat');
-> '636174'
mysql> SELECT 0x636174;
-> 'cat'
```

## 9.1.4. Boolean Values

Beginning with MySQL 4.1, The constants `TRUE` and `FALSE` evaluate to `1` and `0`, respectively. The constant names can be written in any lettercase.

```
mysql> SELECT TRUE, true, FALSE, false;
-> 1, 1, 0, 0
```

## 9.1.5. NULL Values

The `NULL` value means “no data.” `NULL` can be written in any lettercase.

Be aware that the `NULL` value is different from values such as `0` for numeric types or the empty string for string types. See [Section A.5.3, “Problems with NULL Values”](#).

For text file import or export operations performed with `LOAD DATA INFILE` or `SELECT ... INTO OUTFILE`, `NULL` is represented by the `\N` sequence. See [Section 13.2.5, “LOAD DATA INFILE Syntax”](#).

## 9.2. Database, Table, Index, Column, and Alias Names

Database, table, index, column, and alias names are identifiers. This section describes the allowable syntax for identifiers in MySQL.

The following table describes the maximum length for each type of identifier.

Identifier	Maximum Length
Database	64
Table	64
Column	64
Index	64

Alias	255
-------	-----

There are some restrictions on the characters that may appear in identifiers:

- No identifier can contain ASCII 0 (0x00) or a byte with a value of 255.
- Before MySQL 4.1, identifier quote characters should not be used in identifiers. As of 4.1, the use of identifier quote characters in identifiers is permitted, although it is best to avoid doing so if possible.
- Database, table, and column names should not end with space characters.
- Database names cannot contain '/', '\', '.', or characters that are not allowed in a directory name.
- Table names cannot contain '/', '\', '.', or characters that are not allowed in a filename.

Beginning with MySQL 4.1, identifiers are stored using Unicode (UTF-8). This applies to identifiers in table definitions that stored in `.frm` files and to identifiers stored in the grant tables in the `mysql` database. Although Unicode identifiers can include multi-byte characters, note that the maximum lengths shown in the table are byte counts until MySQL 4.1.5; until that version, if an identifier does contain multi-byte characters, the number of *characters* allowed in the identifier is less than the value shown in the table.

An identifier may be quoted or unquoted. If an identifier is a reserved word or contains special characters, you *must* quote it whenever you refer to it. (Exception: A word that follows a period in a qualified name must be an identifier, so it is not necessary to quote it, even if it is a reserved word.) For a list of reserved words, see [Section 9.5, “Treatment of Reserved Words in MySQL”](#). Special characters are those outside the set of alphanumeric characters from the current character set, '\_', and '\$'.

The identifier quote character is the backtick (`):

```
mysql> SELECT * FROM `select` WHERE `select`.id > 100;
```

If the `ANSI_QUOTES` SQL mode is enabled, it is also allowable to quote identifiers within double quotes:

```
mysql> CREATE TABLE "test" (col INT);
ERROR 1064: You have an error in your SQL syntax. (...)
mysql> SET sql_mode='ANSI_QUOTES';
mysql> CREATE TABLE "test" (col INT);
Query OK, 0 rows affected (0.00 sec)
```

Note: Because the `ANSI_QUOTES` mode causes the server to interpret double-quoted strings as identifiers, string literals must be enclosed within single quotes when this mode is enabled. They cannot be enclosed within double quotes.

The server SQL mode is controlled as described in [Section 5.2.5, “SQL Modes”](#).

As of MySQL 4.1, identifier quote characters can be included within an identifier if you quote the identifier. If the character to be included within the identifier is the same as that used to quote the identifier itself, double the character. The following statement creates a table named `a`b` that contains a column named `c`d`:

```
mysql> CREATE TABLE `a``b` (`c`d` INT);
```

Identifier quoting was introduced in MySQL 3.23.6 to allow use of identifiers that are reserved words or that contain special characters. Before 3.23.6, you cannot use identifiers that require quotes, so the rules

for legal identifiers are more restrictive:

- A name may consist of alphanumeric characters from the current character set, ‘\_’, and ‘\$’. The default character set is cp1252 (Latin1). This may be changed with the `-default-character-set` option to `mysqld`. See [Section 5.10.1, “The Character Set Used for Data and Sorting”](#).
- A name may start with any character that is legal in a name. In particular, a name may start with a digit; this differs from many other database systems! However, an unquoted name cannot consist *only* of digits.
- You cannot use the ‘.’ character in names because it is used to extend the format by which you can refer to columns (see [Section 9.2.1, “Identifier Qualifiers”](#)).

It is recommended that you do not use names of the form *Me* or *MeN*, where *M* and *N* are integers. For example, avoid using `1e` or `2e2` as identifiers, because an expression such as `1e+3` is ambiguous. Depending on context, it might be interpreted as the expression `1e + 3` or as the number `1e+3`.

Be careful when using `MD5()` to produce table names because it can produce names in illegal or ambiguous formats such as those just described.

## 9.2.1. Identifier Qualifiers

MySQL allows names that consist of a single identifier or multiple identifiers. The components of a multiple-part name should be separated by period (‘.’) characters. The initial parts of a multiple-part name act as qualifiers that affect the context within which the final identifier is interpreted.

In MySQL you can refer to a column using any of the following forms:

Column Reference	Meaning
<code>col_name</code>	The column <code>col_name</code> from whichever table used in the statement contains a column of that name.
<code>tbl_name.col_name</code>	The column <code>col_name</code> from table <code>tbl_name</code> of the default database.
<code>db_name.tbl_name.col_name</code>	The column <code>col_name</code> from table <code>tbl_name</code> of the database <code>db_name</code> . This syntax is unavailable before MySQL 3.22.

If any components of a multiple-part name require quoting, quote them individually rather than quoting the name as a whole. For example, write ``my-table`.`my-column``, not ``my-table.my-column``.

You need not specify a `tbl_name` or `db_name.tbl_name` prefix for a column reference in a statement unless the reference would be ambiguous. Suppose that tables `t1` and `t2` each contain a column `c`, and you retrieve `c` in a `SELECT` statement that uses both `t1` and `t2`. In this case, `c` is ambiguous because it is not unique among the tables used in the statement. You must qualify it with a table name as `t1.c` or `t2.c` to indicate which table you mean. Similarly, to retrieve from a table `t` in database `db1` and from a table `t` in database `db2` in the same statement, you must refer to columns in those tables as `db1.t.col_name` and `db2.t.col_name`.

A word that follows a period in a qualified name must be an identifier, so it is not necessary to quote it, even if it is a reserved word.

The syntax `.tbl_name` means the table `tbl_name` in the default database. This syntax is accepted for ODBC compatibility because some ODBC programs prefix table names with a ‘.’ character.

## 9.2.2. Identifier Case Sensitivity

In MySQL, databases correspond to directories within the data directory. Each table within a database corresponds to at least one file within the database directory (and possibly more, depending on the storage engine). Consequently, the case sensitivity of the underlying operating system determines the case sensitivity of database and table names. This means database and table names are not case sensitive in Windows, and case sensitive in most varieties of Unix. One notable exception is Mac OS X, which is Unix-based but uses a default filesystem type (HFS+) that is not case sensitive. However, Mac OS X also supports UFS volumes, which are case sensitive just as on any Unix. See [Section 1.9.4, “MySQL Extensions to Standard SQL”](#). The `lower_case_table_names` system variable also affects how the server handles identifier case sensitivity, as described later in this section.

**Note:** Although database and table names are not case sensitive on some platforms, you should not refer to a given database or table using different cases within the same statement. The following statement would not work because it refers to a table both as `my_table` and as `MY_TABLE`:

```
mysql> SELECT * FROM my_table WHERE MY_TABLE.col=1;
```

Column, index, stored routine, and trigger names are not case sensitive on any platform, nor are column aliases.

Table aliases are case sensitive before MySQL 4.1.1. The following query would not work because it refers to the alias both as `a` and as `A`:

```
mysql> SELECT col_name FROM tbl_name AS a
-> WHERE a.col_name = 1 OR A.col_name = 2;
```

If you have trouble remembering the allowable lettercase for database and table names, it is best to adopt a consistent convention, such as always creating and referring to databases and tables using lowercase names. This convention is recommended for maximum portability and ease of use.

How table and database names are stored on disk and used in MySQL is affected by the `lower_case_table_names` system variable, which you can set when starting `mysqld`. `lower_case_table_names` can take the values shown in the following table. On Unix, the default value of `lower_case_table_names` is 0. On Windows, the default value is 1. On Mac OS X, the default is 1 before MySQL 4.0.18 and 2 as of 4.0.18.

Value	Meaning
0	Table and database names are stored on disk using the lettercase specified in the <code>CREATE TABLE</code> or <code>CREATE DATABASE</code> statement. Name comparisons are case sensitive. Note that if you force this variable to 0 with <code>--lower-case-table-names=0</code> on a case-insensitive filesystem and access <code>MyISAM</code> tablenamees using different lettercases, this may lead to index corruption.
1	Table names are stored in lowercase on disk and name comparisons are not case sensitive. MySQL converts all table names to lowercase on storage and lookup. This behavior also applies to database names as of MySQL 4.0.2, and to table aliases as of 4.1.1.
2	Table and database names are stored on disk using the lettercase specified in the <code>CREATE TABLE</code> or <code>CREATE DATABASE</code> statement, but MySQL converts them to lowercase on lookup. Name comparisons are not case sensitive. <b>Note:</b> This works <i>only</i> on filesystems that are not case sensitive! <code>InnoDB</code> table names are stored in lowercase, as for <code>lower_case_table_names=1</code> . Setting <code>lower_case_table_names</code> to 2 can be done as of MySQL 4.0.18.

If you are using MySQL on only one platform, you don't normally have to change the `lower_case_table_names` variable. However, you may encounter difficulties if you want to trans-



fer tables between platforms that differ in filesystem case sensitivity. For example, on Unix, you can have two different tables named `my_table` and `MY_TABLE`, but on Windows those names are considered identical. To avoid data transfer problems stemming from lettercase of database or table names, you have two options:

- Use `lower_case_table_names=1` on all systems. The main disadvantage with this is that when you use `SHOW TABLES` or `SHOW DATABASES`, you don't see the names in their original lettercase.
- Use `lower_case_table_names=0` on Unix and `lower_case_table_names=2` on Windows. This preserves the lettercase of database and table names. The disadvantage of this is that you must ensure that your statements always refer to your database and table names with the correct lettercase on Windows. If you transfer your statements to Unix, where lettercase is significant, they do not work if the lettercase is incorrect.

Exception: If you are using InnoDB tables, you should set `lower_case_table_names` to 1 on all platforms to force names to be converted to lowercase.

Note that if you plan to set the `lower_case_table_names` system variable to 1 on Unix, you must first convert your old database and table names to lowercase before restarting `mysqld` with the new variable setting.

## 9.3. User-Defined Variables

MySQL supports user variables as of version 3.23.6. You can store a value in a user-defined variable and then refer to it later. This enables you to pass values from one statement to another. *User-defined variables are connection-specific*. That is, a user variable defined by one client cannot be seen or used by other clients. All variables for a given client connection are automatically freed when that client exits.

User variables are written as `@var_name`, where the variable name `var_name` may consist of alphanumeric characters from the current character set, `'.`, `'_'`, and `'$'`. The default character set is `latin1` (cp1252 West European). This may be changed with the `--default-character-set` option to `mysqld`. See Section 5.10.1, “The Character Set Used for Data and Sorting”. Section 5.10.1, “The Character Set Used for Data and Sorting”. A user variable name can contain other characters if you quote it as a string or identifier (for example, `@'my-var'`, `@"my-var"`, or `@`my-var``).

Note: User variable names are case sensitive prior to MySQL 5.0.

One way to set a user variable is by issuing a `SET` statement:

```
SET @var_name = expr [, @var_name = expr] ...
```

For `SET`, either `=` or `:=` can be used as the assignment operator. The `expr` assigned to each variable can evaluate to an integer, real, string, or `NULL` value. However, if the value of the variable is selected in a result set, it is returned to the client as a string.

You can also assign a value to a user variable in statements other than `SET`. In this case, the assignment operator must be `:=` and not `=` because `=` is treated as a comparison operator in non-`SET` statements:

```
mysql> SET @t1=0, @t2=0, @t3=0;
mysql> SELECT @t1:=(@t2:=1)+@t3:=4,@t1,@t2,@t3;
+-----+-----+-----+-----+
| @t1:=(@t2:=1)+@t3:=4 | @t1 | @t2 | @t3 |
+-----+-----+-----+-----+
|                    5 |    5 |    1 |    4 |
+-----+-----+-----+-----+
```

User variables may be used in contexts where expressions are allowed. This does not currently include

contexts that explicitly require a literal value, such as in the `LIMIT` clause of a `SELECT` statement, or the `IGNORE N LINES` clause of a `LOAD DATA` statement.

Beginning with MySQL 4.1.1, if a user variable is assigned a string value, it has the same character set and collation as the string. The coercibility of user variables is “implicit” as of MySQL 4.1.11 and 5.0.3. (This is the same coercibility as table column values.)

**Note:** In a `SELECT` statement, each expression is evaluated only when sent to the client. This means that in a `HAVING`, `GROUP BY`, or `ORDER BY` clause, you cannot refer to an expression that involves variables that are set in the `SELECT` list. For example, the following statement does *not* work as expected:

```
mysql> SELECT (@aa:=id) AS a, (@aa+3) AS b FROM tbl_name HAVING b=5;
```

The reference to `b` in the `HAVING` clause refers to an alias for an expression in the `SELECT` list that uses `@aa`. This does not work as expected: `@aa` contains the value of `id` from the previous selected row, not from the current row.

The order of evaluation for user variables is undefined and may change based on the elements contained within a given query. In `SELECT @a, @a := @a+1 ...`, you might think that MySQL will evaluate `@a` first and then do an assignment second, but changing the query (for example, by adding a `GROUP BY`, `HAVING`, or `ORDER BY` clause) may change the order of evaluation.

The general rule is to never assign a value to a user variable in one part of a statement *and* use the same variable in some other part the same statement. You might get the results you expect, but this is not guaranteed.

Another issue with setting a variable and using it in the same statement is that the default result type of a variable is based on the type of the variable at the start of the statement. The following example illustrates this:

```
mysql> SET @a='test';
mysql> SELECT @a,(@a:=20) FROM tbl_name;
```

For this `SELECT` statement, MySQL reports to the client that column one is a string and converts all accesses of `@a` to strings, even though `@a` is set to a number for the second row. After the `SELECT` statement executes, `@a` is regarded as a number for the next statement.

To avoid problems with this behavior, either do not set and use the same variable within a single statement, or else set the variable to `0`, `0.0`, or `' '` to define its type before you use it.

If you refer to a variable that has not been initialized, it has a value of `NULL` and a type of string.

## 9.4. Comment Syntax

MySQL Server supports three comment styles:

- From a `#` character to the end of the line.
- From a `--` sequence to the end of the line. This style is supported as of MySQL 3.23.3. In MySQL, the `<quo--` (double-dash) comment style requires the second dash to be followed by at least one whitespace or control character (such as a space, tab, newline, and so on). This syntax differs slightly from standard SQL comment syntax, as discussed in [Section 1.9.5.7, “-- as the Start of a Comment”](#).
- From a `/*` sequence to the following `*/` sequence, as in the C programming language. This syntax allows a comment to extend over multiple lines because the beginning and closing sequences need not be on the same line.

The following example demonstrates all three comment styles:

```
mysql> SELECT 1+1;      # This comment continues to the end of line
mysql> SELECT 1+1;      -- This comment continues to the end of line
mysql> SELECT 1 /* this is an in-line comment */ + 1;
mysql> SELECT 1+
/*
this is a
multiple-line comment
*/
1;
```

MySQL Server supports some variants of C-style comments. These enable you to write code that includes MySQL extensions, but is still portable, by using comments of the following form:

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

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

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

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

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

The comment syntax just described applies to how the `mysqld` server parses SQL statements. The `mysql` client program also performs some parsing of statements before sending them to the server. (It does this to determine statement boundaries within a multiple-statement input line.) However, there are some limitations on the way that `mysql` parses `/* ... */` comments:

- A semicolon within the comment is taken to indicate the end of the current SQL statement and anything following it to indicate the beginning of the next statement. This problem was fixed in MySQL 4.0.13.
- A single quote, double quote, or backtick character is taken to indicate the beginning of a quoted string or identifier, even within a comment. If the quote is not matched by a second quote within the comment, the parser doesn't realize the comment has ended. If you are running `mysql` interactively, you can tell that it has gotten confused like this because the prompt changes from `mysql>` to `'>`, `">`, or ``>`. This problem was fixed in MySQL 4.1.1.

For affected versions of MySQL, these limitations apply both when you run `mysql` interactively and when you put commands in a file and use `mysql` in batch mode to process the file with `mysql < file_name`.

## 9.5. Treatment of Reserved Words in MySQL

A common problem stems from trying to use an identifier such as a table or column name that is a reserved word such as `SELECT` or the name of a built-in MySQL data type or function such as `TIMESTAMP` or `GROUP`.

If an identifier is a reserved word, you must quote it as described in [Section 9.2, “Database, Table, Index, Column, and Alias Names”](#). Exception: A word that follows a period in a qualified name must be an identifier, so it is not necessary to quote it, even if it is a reserved word.

You are permitted to use function names as identifiers. For example, `ABS` is acceptable as a column name. However, by default, no whitespace is allowed in function invocations between the function name and the following `'(` character. This requirement allows a function call to be distinguished from a reference to a column name.

A side effect of this behavior is that omitting a space in some contexts causes an identifier to be interpreted as a function name. For example, this statement is legal:

```
mysql> CREATE TABLE abs (val INT);
```

But omitting the space after `abs` causes a syntax error because the statement then appears to invoke the `ABS()` function:

```
mysql> CREATE TABLE abs(val INT);
ERROR 1064 (42000) at line 2: You have an error in your SQL
syntax ... near 'abs(val INT)'
```

If the `IGNORE_SPACE` SQL mode is enabled, the server allows function invocations to have whitespace between a function name and the following `'(` character. This causes function names to be treated as reserved words. As a result, identifiers that are the same as function names must be quoted as described in [Section 9.2, “Database, Table, Index, Column, and Alias Names”](#). The server SQL mode is controlled as described in [Section 5.2.5, “SQL Modes”](#).

The words in the following table are explicitly reserved in MySQL 4.1. At some point, you might update to a higher version, so it's a good idea to have a look at future reserved words, too. You can find these in the manuals that cover higher versions of MySQL. Most of the words in the table are forbidden by standard SQL as column or table names (for example, `GROUP`). A few are reserved because MySQL needs them and (currently) uses a `yacc` parser. A reserved word can be used as an identifier if you quote it.

ADD	ALL	ALTER
ANALYZE	AND	AS
ASC	BEFORE	BETWEEN
BIGINT	BINARY	BLOB
BOTH	BY	CASCADE
CASE	CHANGE	CHAR
CHARACTER	CHECK	COLLATE
COLUMN	COLUMNS	CONSTRAINT
CONVERT	CREATE	CROSS
CURRENT_DATE	CURRENT_TIME	CURRENT_TIMESTAMP
CURRENT_USER	DATABASE	DATABASES
DAY_HOUR	DAY_MICROSECOND	DAY_MINUTE
DAY_SECOND	DEC	DECIMAL
DEFAULT	DELAYED	DELETE
DESC	DESCRIBE	DISTINCT
DISTINCTROW	DIV	DOUBLE
DROP	DUAL	ELSE
ENCLOSED	ESCAPED	EXISTS
EXPLAIN	FALSE	FIELDS
FLOAT	FLOAT4	FLOAT8

FOR	FORCE	FOREIGN
FROM	FULLTEXT	GRANT
GROUP	HAVING	HIGH_PRIORITY
HOUR_MICROSECOND	HOUR_MINUTE	HOUR_SECOND
IF	IGNORE	IN
INDEX	INFILE	INNER
INSERT	INT	INT1
INT2	INT3	INT4
INT8	INTEGER	INTERVAL
INTO	IS	JOIN
KEY	KEYS	KILL
LEADING	LEFT	LIKE
LIMIT	LINES	LOAD
LOCALTIME	LOCALTIMESTAMP	LOCK
LONG	LOBLOB	LONGTEXT
LOW_PRIORITY	MATCH	MEDIUMBLOB
MEDIUMINT	MEDIUMTEXT	MIDDLEINT
MINUTE_MICROSECOND	MINUTE_SECOND	MOD
NATURAL	NOT	NO_WRITE_TO_BINLOG
NULL	NUMERIC	ON
OPTIMIZE	OPTION	OPTIONALLY
OR	ORDER	OUTER
OUTFILE	PRECISION	PRIMARY
PRIVILEGES	PROCEDURE	PURGE
RAID0	READ	REAL
REFERENCES	REGEXP	RENAME
REPLACE	REQUIRE	RESTRICT
REVOKE	RIGHT	RLIKE
SECOND_MICROSECOND	SELECT	SEPARATOR
SET	SHOW	SMALLINT
SONAME	SPATIAL	SQL_BIG_RESULT
SQL_CALC_FOUND_ROWS	SQL_SMALL_RESULT	SSL
STARTING	STRAIGHT_JOIN	TABLE
TABLES	TERMINATED	THEN
TINYBLOB	TINYINT	TINYTEXT
TO	TRAILING	TRUE
UNION	UNIQUE	UNLOCK
UNSIGNED	UPDATE	USAGE
USE	USING	UTC_DATE
UTC_TIME	UTC_TIMESTAMP	VALUES
VARBINARY	VARCHAR	VARCHARACTER
VARYING	WHEN	WHERE

WITH	WRITE	X509
XOR	YEAR_MONTH	ZEROFILL

The following are new reserved words in MySQL 4.0: `CHECK`, `FORCE`, `LOCALTIME`, `LOCALTIMESTAMP`, `REQUIRE`, `SQL_CALC_FOUND_ROWS`, `SSL`, `X509`, `XOR`.

The following are new reserved words in MySQL 4.1: `BEFORE`, `COLLATE`, `CONVERT`, `CURRENT_USER`, `DAY_MICROSECOND`, `DIV`, `DUAL`, `FALSE`, `HOURL_MICROSECOND`, `MINUTE_MICROSECOND`, `MOD`, `NO_WRITE_TO_BINLOG`, `SECOND_MICROSECOND`, `SEPARATOR`, `SPATIAL`, `TRUE`, `UTC_DATE`, `UTC_TIME`, `UTC_TIMESTAMP`, `VARCHARACTER`.

MySQL allows some keywords to be used as unquoted identifiers because many people previously used them. Examples are those in the following list:

- `ACTION`
- `BIT`
- `DATE`
- `ENUM`
- `NO`
- `TEXT`
- `TIME`
- `TIMESTAMP`

---

# Chapter 10. Character Set Support

Improved support for character set handling was added to MySQL in version 4.1. This support enables you to store data using a variety of character sets and perform comparisons according to a variety of collations. You can specify character sets at the server, database, table, and column level. MySQL supports the use of character sets for the [MyISAM](#), [MEMORY](#), and (as of MySQL 4.1.2) [InnoDB](#) storage engines. The [ISAM](#) storage engine does not include character set support; there are no plans to change this, because [ISAM](#) is deprecated.

**Note:** The [NDBCluster](#) storage engine in MySQL 4.1 (available beginning with MySQL 4.1.3-Max) provides limited character set and collation support; see [Section 15.11, “Known Limitations of MySQL Cluster”](#).

This chapter discusses the following topics:

- What are character sets and collations?
- The multiple-level default system for character set assignment
- Syntax for specifying character sets and collations
- Affected functions and operations
- Unicode support
- The character sets and collations that are available, with notes

Character set issues affect data storage, but also communication between client programs and the MySQL server. If you want the client program to communicate with the server using a character set different from the default, you'll need to indicate which one. For example, to use the [utf8](#) Unicode character set, issue this statement after connecting to the server:

```
SET NAMES 'utf8';
```

For more information about character set-related issues in client/server communication, see [Section 10.4, “Connection Character Sets and Collations”](#).

## 10.1. Character Sets and Collations in General

A *character set* is a set of symbols and encodings. A *collation* is a set of rules for comparing characters in a character set. Let's make the distinction clear with an example of an imaginary character set.

Suppose that we have an alphabet with four letters: 'A', 'B', 'a', 'b'. We give each letter a number: 'A' = 0, 'B' = 1, 'a' = 2, 'b' = 3. The letter 'A' is a symbol, the number 0 is the **encoding** for 'A', and the combination of all four letters and their encodings is a **character set**.

Suppose that we want to compare two string values, 'A' and 'B'. The simplest way to do this is to look at the encodings: 0 for 'A' and 1 for 'B'. Because 0 is less than 1, we say 'A' is less than 'B'. What we've just done is apply a collation to our character set. The collation is a set of rules (only one rule in this case): “compare the encodings.” We call this simplest of all possible collations a *binary* collation.

But what if we want to say that the lowercase and uppercase letters are equivalent? Then we would have at least two rules: (1) treat the lowercase letters 'a' and 'b' as equivalent to 'A' and 'B'; (2) then compare the encodings. We call this a *case-insensitive* collation. It's a little more complex than a binary collation.

In real life, most character sets have many characters: not just ‘A’ and ‘B’ but whole alphabets, sometimes multiple alphabets or eastern writing systems with thousands of characters, along with many special symbols and punctuation marks. Also in real life, most collations have many rules, not just for whether to distinguish lettercase, but also for whether to distinguish accents (an “accent” is a mark attached to a character as in German <code>ö</code>), and for multiple-character mappings (such as the rule that ‘ö’ = ‘OE’ in one of the two German collations).

MySQL 4.1 can do these things for you:

- Store strings using a variety of character sets
- Compare strings using a variety of collations
- Mix strings with different character sets or collations in the same server, the same database, or even the same table
- Allow specification of character set and collation at any level

In these respects, not only is MySQL 4.1 far more flexible than MySQL 4.0, it also is far ahead of most other database management systems. However, to use these features effectively, you need to know what character sets and collations are available, how to change the defaults, and how they affect the behavior of string operators and functions.

## 10.2. Character Sets and Collations in MySQL

The MySQL server can support multiple character sets. To list the available character sets, use the `SHOW CHARACTER SET` statement. A partial listing follows. For more complete information, see [Section 10.10, “Character Sets and Collations That MySQL Supports”](#).

```
mysql> SHOW CHARACTER SET;
```

Charset	Description	Default collation	Maxlen
big5	Big5 Traditional Chinese	big5_chinese_ci	2
dec8	DEC West European	dec8_swedish_ci	1
cp850	DOS West European	cp850_general_ci	1
hp8	HP West European	hp8_english_ci	1
koi8r	KOI8-R Relcom Russian	koi8r_general_ci	1
latin1	cp1252 West European	latin1_swedish_ci	1
latin2	ISO 8859-2 Central European	latin2_general_ci	1
swe7	7bit Swedish	swe7_swedish_ci	1
ascii	US ASCII	ascii_general_ci	1
ujis	EUC-JP Japanese	ujis_japanese_ci	3
sjis	Shift-JIS Japanese	sjis_japanese_ci	2
hebrew	ISO 8859-8 Hebrew	hebrew_general_ci	1
tis620	TIS620 Thai	tis620_thai_ci	1
euckr	EUC-KR Korean	euckr_korean_ci	2
koi8u	KOI8-U Ukrainian	koi8u_general_ci	1
gb2312	GB2312 Simplified Chinese	gb2312_chinese_ci	2
greek	ISO 8859-7 Greek	greek_general_ci	1
cp1250	Windows Central European	cp1250_general_ci	1
gbk	GBK Simplified Chinese	gbk_chinese_ci	2
latin5	ISO 8859-9 Turkish	latin5_turkish_ci	1
...			

Any given character set always has at least one collation. It may have several collations. To list the collations for a character set, use the `SHOW COLLATION` statement. For example, to see the collations for the `latin1` (cp1252 West European) character set, use this statement to find those collation names that begin with `latin1`:

```
mysql> SHOW COLLATION LIKE 'latin1%';
```

Collation	Charset	Id	Default	Compiled	Sortlen
latin1_swedish_ci	latin1	43	NO	NO	1
latin1_general_ci	latin1	44	NO	NO	1
latin1_general_cs	latin1	45	NO	NO	1
latin1_bin	latin1	46	NO	NO	1
latin1_danish_ci	latin1	47	NO	NO	1



latin1_german1_ci	latin1	5			0
latin1_swedish_ci	latin1	8	Yes	Yes	1
latin1_danish_ci	latin1	15			0
latin1_german2_ci	latin1	31		Yes	2
latin1_bin	latin1	47		Yes	1
latin1_general_ci	latin1	48			0
latin1_general_cs	latin1	49			0
latin1_spanish_ci	latin1	94			0

The `latin1` collations have the following meanings:

Collation	Meaning
<code>latin1_german1_ci</code>	German DIN-1
<code>latin1_swedish_ci</code>	Swedish/Finnish
<code>latin1_danish_ci</code>	Danish/Norwegian
<code>latin1_german2_ci</code>	German DIN-2
<code>latin1_bin</code>	Binary according to <code>latin1</code> encoding
<code>latin1_general_ci</code>	Multilingual (Western European)
<code>latin1_general_cs</code>	Multilingual (ISO Western European), case sensitive
<code>latin1_spanish_ci</code>	Modern Spanish

Collations have these general characteristics:

- Two different character sets cannot have the same collation.
- Each character set has one collation that is the *default collation*. For example, the default collation for `latin1` is `latin1_swedish_ci`. The output for `SHOW CHARACTER SET` indicates which collation is the default for each displayed character set.
- There is a convention for collation names: They start with the name of the character set with which they are associated, they usually include a language name, and they end with `_ci` (case insensitive), `_cs` (case sensitive), or `_bin` (binary).

## 10.3. Specifying Character Sets and Collations

There are default settings for character sets and collations at four levels: server, database, table, and column. The following description may appear complex, but it has been found in practice that multiple-level defaulting leads to natural and obvious results.

`CHARACTER SET` is used in clauses that specify a character set. `CHARSET` may be used as a synonym for `CHARACTER SET`.

### 10.3.1. Server Character Set and Collation

MySQL Server has a server character set and a server collation. These can be set at server startup and changed at runtime.

Initially, the server character set and collation depend on the options that you use when you start `mysqld`. You can use `--character-set-server` for the character set. Along with it, you can add `--collation-server` for the collation. If you don't specify a character set, that is the same as saying `--character-set-server=latin1`. If you specify only a character set (for example, `latin1`) but not a collation, that is the same as saying `--character-set-server=latin1 -`

`--collation-server=latin1_swedish_ci` because `latin1_swedish_ci` is the default collation for `latin1`. Therefore, the following three commands all have the same effect:

```
shell> mysqld
shell> mysqld --character-set-server=latin1
shell> mysqld --character-set-server=latin1 \
            --collation-server=latin1_swedish_ci
```

One way to change the settings is by recompiling. If you want to change the default server character set and collation when building from sources, use: `--with-charset` and `--with-collation` as arguments for `configure`. For example:

```
shell> ./configure --with-charset=latin1
```

Or:

```
shell> ./configure --with-charset=latin1 \
            --with-collation=latin1_german1_ci
```

Both `mysqld` and `configure` verify that the character set/collation combination is valid. If not, each program displays an error message and terminates.

The current server character set and collation can be determined from the values of the `character_set_server` and `collation_server` system variables. These variables can be changed at runtime.

## 10.3.2. Database Character Set and Collation

Every database has a database character set and a database collation. The `CREATE DATABASE` and `ALTER DATABASE` statements have optional clauses for specifying the database character set and collation:

```
CREATE DATABASE db_name
    [[DEFAULT] CHARACTER SET charset_name]
    [[DEFAULT] COLLATE collation_name]

ALTER DATABASE db_name
    [[DEFAULT] CHARACTER SET charset_name]
    [[DEFAULT] COLLATE collation_name]
```

All database options are stored in a text file named `db.opt` that can be found in the database directory.

The `CHARACTER SET` and `COLLATE` clauses make it possible to create databases with different character sets and collations on the same MySQL server.

Example:

```
CREATE DATABASE db_name CHARACTER SET latin1 COLLATE latin1_swedish_ci;
```

MySQL chooses the database character set and database collation in the following manner:

- If both `CHARACTER SET X` and `COLLATE Y` were specified, then character set `X` and collation `Y`.
- If `CHARACTER SET X` was specified without `COLLATE`, then character set `X` and its default collation.
- If `COLLATE Y` was specified without `CHARACTER SET`, then the character set associated with `Y` and collation `Y`.

- Otherwise, the server character set and server collation.

The database character set and collation are used as default values if the table character set and collation are not specified in `CREATE TABLE` statements. They have no other purpose.

The character set and collation for the default database can be determined from the values of the `character_set_database` and `collation_database` system variables. The server sets these variables whenever the default database changes. If there is no default database, the variables have the same value as the corresponding server-level system variables, `character_set_server` and `collation_server`.

### 10.3.3. Table Character Set and Collation

Every table has a table character set and a table collation. The `CREATE TABLE` and `ALTER TABLE` statements have optional clauses for specifying the table character set and collation:

```
CREATE TABLE tbl_name (column_list)
  [[DEFAULT] CHARACTER SET charset_name] [COLLATE collation_name]

ALTER TABLE tbl_name
  [[DEFAULT] CHARACTER SET charset_name] [COLLATE collation_name]
```

Example:

```
CREATE TABLE t1 ( ... ) CHARACTER SET latin1 COLLATE latin1_danish_ci;
```

MySQL chooses the table character set and collation in the following manner:

- If both `CHARACTER SET X` and `COLLATE Y` were specified, then character set `X` and collation `Y`.
- If `CHARACTER SET X` was specified without `COLLATE`, then character set `X` and its default collation.
- If `COLLATE Y` was specified without `CHARACTER SET`, then the character set associated with `Y` and collation `Y`.
- Otherwise, the database character set and collation.

The table character set and collation are used as default values if the column character set and collation are not specified in individual column definitions. The table character set and collation are MySQL extensions; there are no such things in standard SQL.

### 10.3.4. Column Character Set and Collation

Every “character” column (that is, a column of type `CHAR`, `VARCHAR`, or `TEXT`) has a column character set and a column collation. Column definition syntax has optional clauses for specifying the column character set and collation:

```
col_name {CHAR | VARCHAR | TEXT} (col_length)
  [CHARACTER SET charset_name] [COLLATE collation_name]
```

Example:

```
CREATE TABLE Table1
(
  column1 VARCHAR(5) CHARACTER SET latin1 COLLATE latin1_german1_ci
```

```
);
```

MySQL chooses the column character set and collation in the following manner:

- If both `CHARACTER SET X` and `COLLATE Y` were specified, then character set `X` and collation `Y` are used.
- If `CHARACTER SET X` was specified without `COLLATE`, then character set `X` and its default collation are used.
- If `COLLATE Y` was specified without `CHARACTER SET`, then the character set associated with `Y` and collation `Y`.
- Otherwise, the table character set and collation are used.

The `CHARACTER SET` and `COLLATE` clauses are standard SQL.

### 10.3.5. Character String Literal Character Set and Collation

Every character string literal has a character set and a collation.

A character string literal may have an optional character set introducer and `COLLATE` clause:

```
[_charset_name]'string' [COLLATE collation_name]
```

Examples:

```
SELECT 'string';
SELECT _latin1'string';
SELECT _latin1'string' COLLATE latin1_danish_ci;
```

For the simple statement `SELECT 'string'`, the string has the character set and collation defined by the `character_set_connection` and `collation_connection` system variables.

The `_charset_name` expression is formally called an *introducer*. It tells the parser, “the string that is about to follow uses character set `X`.” Because this has confused people in the past, we emphasize that an introducer does not cause any conversion; it is strictly a signal that does not change the string's value. An introducer is also legal before standard hex literal and numeric hex literal notation (`x'literal'` and `0xn`).

Examples:

```
SELECT _latin1 x'AABBCC';
SELECT _latin1 0xAABBCC;
```

MySQL determines a literal's character set and collation in the following manner:

- If both `_X` and `COLLATE Y` were specified, then character set `X` and collation `Y` are used.
- If `_X` is specified but `COLLATE` is not specified, then character set `X` and its default collation are used.
- Otherwise, the character set and collation given by the `character_set_connection` and `collation_connection` system variables are used.

Examples:

- A string with `latin1` character set and `latin1_german1_ci` collation:

```
SELECT _latin1'Müller' COLLATE latin1_german1_ci;
```

- A string with `latin1` character set and its default collation (that is, `latin1_swedish_ci`):

```
SELECT _latin1'Müller';
```

- A string with the connection default character set and collation:

```
SELECT 'Müller';
```

Character set introducers and the `COLLATE` clause are implemented according to standard SQL specifications.

### 10.3.6. National Character Set

Before MySQL 4.1, `NCHAR` and `CHAR` were synonymous. Standard SQL defines `NCHAR` or `NATIONAL CHAR` as a way to indicate that a `CHAR` column should use some predefined character set. MySQL 4.1 and up uses `utf8` as that predefined character set. For example, these data type declarations are equivalent:

```
CHAR(10) CHARACTER SET utf8
NATIONAL CHARACTER(10)
NCHAR(10)
```

As are these:

```
VARCHAR(10) CHARACTER SET utf8
NATIONAL VARCHAR(10)
NCHAR VARCHAR(10)
NATIONAL CHARACTER VARYING(10)
NATIONAL CHAR VARYING(10)
```

You can use `N'literal'` to create a string in the national character set. These two statements are equivalent:

```
SELECT N'some text';
SELECT _utf8'some text';
```

### 10.3.7. Examples of Character Set and Collation Assignment

The following examples show how MySQL determines default character set and collation values.

#### Example 1: Table and Column Definition

```
CREATE TABLE t1
(
  c1 CHAR(10) CHARACTER SET latin1 COLLATE latin1_german1_ci
) DEFAULT CHARACTER SET latin2 COLLATE latin2_bin;
```

Here we have a column with a `latin1` character set and a `latin1_german1_ci` collation. The definition is explicit, so that's straightforward. Notice that there is no problem with storing a `latin1` column in a `latin2` table.

**Example 2: Table and Column Definition**

```
CREATE TABLE t1
(
  c1 CHAR(10) CHARACTER SET latin1
) DEFAULT CHARACTER SET latin1 COLLATE latin1_danish_ci;
```

This time we have a column with a `latin1` character set and a default collation. Although it might seem natural, the default collation is not taken from the table level. Instead, because the default collation for `latin1` is always `latin1_swedish_ci`, column `c1` has a collation of `latin1_swedish_ci` (not `latin1_danish_ci`).

**Example 3: Table and Column Definition**

```
CREATE TABLE t1
(
  c1 CHAR(10)
) DEFAULT CHARACTER SET latin1 COLLATE latin1_danish_ci;
```

We have a column with a default character set and a default collation. In this circumstance, MySQL checks the table level to determine the column character set and collation. Consequently, the character set for column `c1` is `latin1` and its collation is `latin1_danish_ci`.

**Example 4: Database, Table, and Column Definition**

```
CREATE DATABASE d1
  DEFAULT CHARACTER SET latin2 COLLATE latin2_czech_ci;
USE d1;
CREATE TABLE t1
(
  c1 CHAR(10)
);
```

We create a column without specifying its character set and collation. We're also not specifying a character set and a collation at the table level. In this circumstance, MySQL checks the database level to determine the table settings, which thereafter become the column settings.) Consequently, the character set for column `c1` is `latin2` and its collation is `latin2_czech_ci`.

## 10.3.8. Compatibility with Other DBMSs

For MaxDB compatibility these two statements are the same:

```
CREATE TABLE t1 (f1 CHAR(N) UNICODE);
CREATE TABLE t1 (f1 CHAR(N) CHARACTER SET ucs2);
```

## 10.4. Connection Character Sets and Collations

Several character set and collation system variables relate to a client's interaction with the server. Some of these have been mentioned in earlier sections:

- The server character set and collation can be determined from the values of the `character_set_server` and `collation_server` system variables.
- The character set and collation of the default database can be determined from the values of the `character_set_database` and `collation_database` system variables.

Additional character set and collation system variables are involved in handling traffic for the connec-

tion between a client and the server. Every client has connection-related character set and collation system variables.

Consider what a “connection” is: It's what you make when you connect to the server. The client sends SQL statements, such as queries, over the connection to the server. The server sends responses, such as result sets, over the connection back to the client. This leads to several questions about character set and collation handling for client connections, each of which can be answered in terms of system variables:

- What character set is the statement in when it leaves the client?

The server takes the `character_set_client` system variable to be the character set in which statements are sent by the client.

- What character set should the server translate a statement to after receiving it?

For this, the server uses the `character_set_connection` and `collation_connection` system variables. It converts statements sent by the client from `character_set_client` to `character_set_connection` (except for string literals that have an introducer such as `_latin1` or `_utf8`). `collation_connection` is important for comparisons of literal strings. For comparisons of strings with column values, `collation_connection` does not matter because columns have their own collation, which has a higher collation precedence.

- What character set should the server translate to before shipping result sets or error messages back to the client?

The `character_set_results` system variable indicates the character set in which the server returns query results to the client. This includes result data such as column values, and result metadata such as column names.

You can fine-tune the settings for these variables, or you can depend on the defaults (in which case, you can skip the rest of this section).

There are two statements that affect the connection character sets:

```
SET NAMES 'charset_name'
SET CHARACTER SET 'charset_name'
```

`SET NAMES` indicates what character set the client will use to send SQL statements to the server. Thus, `SET NAMES 'cp1251'` tells the server “future incoming messages from this client are in character set `cp1251`.” It also specifies the character set that the server should use for sending results back to the client. (For example, it indicates what character set to use for column values if you use a `SELECT` statement.)

A `SET NAMES 'x'` statement is equivalent to these three statements:

```
SET character_set_client = x;
SET character_set_results = x;
SET character_set_connection = x;
```

Setting `character_set_connection` to `x` also sets `collation_connection` to the default collation for `x`. To specify one of the character set's collations explicitly, use the optional `COLLATE` clause:

```
SET NAMES 'charset_name' COLLATE 'collation_name'
```

`SET CHARACTER SET` is similar to `SET NAMES` but sets `connection_character_set` and `collation_connection` to `character_set_database` and `collation_database`. A

`SET CHARACTER SET x` statement is equivalent to these three statements:

```
SET character_set_client = x;
SET character_set_results = x;
SET collation_connection = @@collation_database;
```

Setting `collation_connection` also sets `character_set_connection` to the character set associated with the collation.

When a client connects, it sends to the server the name of the character set that it wants to use. The server uses the name to set the `character_set_client`, `character_set_results`, and `character_set_connection` system variables. In effect, the server performs a `SET NAMES` operation using the character set name.

With the `mysql` client, it is not necessary to execute `SET NAMES` every time you start up if you want to use a character set different from the default. You can add the `--default-character-set` option setting to your `mysql` statement line, or in your option file. For example, the following option file setting changes the three character set variables set to `koi8r` each time you invoke `mysql`:

```
[mysql]
default-character-set=koi8r
```

Example: Suppose that `column1` is defined as `CHAR(5) CHARACTER SET latin2`. If you do not say `SET NAMES` or `SET CHARACTER SET`, then for `SELECT column1 FROM t`, the server sends back all the values for `column1` using the character set that the client specified when it connected. On the other hand, if you say `SET NAMES 'latin1'` or `SET CHARACTER SET latin1` before issuing the `SELECT` statement, the server converts the `latin2` values to `latin1` just before sending results back. Conversion may be lossy if there are characters that are not in both character sets.

If you do not want the server to perform any conversion of result sets, set `character_set_results` to `NULL`:

```
SET character_set_results = NULL;
```

**Note:** Currently, UCS-2 cannot be used as a client character set, which means that `SET NAMES 'ucs2'` does not work.

To see the values of the character set and collation system variables that apply to your connection, use these statements:

```
SHOW VARIABLES LIKE 'character_set%';
SHOW VARIABLES LIKE 'collation%';
```

## 10.5. Collation Issues

The following sections discuss various aspects of character set collations.

### 10.5.1. Using `COLLATE` in SQL Statements

With the `COLLATE` clause, you can override whatever the default collation is for a comparison. `COLLATE` may be used in various parts of SQL statements. Here are some examples:

- With `ORDER BY`:

```
SELECT k
FROM t1
ORDER BY k COLLATE latin1_german2_ci;
```



- With **AS**:

```
SELECT k COLLATE latin1_german2_ci AS k1
FROM t1
ORDER BY k1;
```

- With **GROUP BY**:

```
SELECT k
FROM t1
GROUP BY k COLLATE latin1_german2_ci;
```

- With aggregate functions:

```
SELECT MAX(k COLLATE latin1_german2_ci)
FROM t1;
```

- With **DISTINCT**:

```
SELECT DISTINCT k COLLATE latin1_german2_ci
FROM t1;
```

- With **WHERE**:

```
SELECT *
FROM t1
WHERE _latin1 'Müller' COLLATE latin1_german2_ci = k;
```

```
SELECT *
FROM t1
WHERE k LIKE _latin1 'Müller' COLLATE latin1_german2_ci;
```

- With **HAVING**:

```
SELECT k
FROM t1
GROUP BY k
HAVING k = _latin1 'Müller' COLLATE latin1_german2_ci;
```

## 10.5.2. COLLATE Clause Precedence

The **COLLATE** clause has high precedence (higher than `|`), so the following two expressions are equivalent:

```
x | | y COLLATE z
x | | (y COLLATE z)
```

## 10.5.3. BINARY Operator

The **BINARY** operator casts the string following it to a binary string. This is an easy way to force a comparison to be done byte by byte rather than character by character. **BINARY** also causes trailing spaces to be significant.

```
mysql> SELECT 'a' = 'A';
-> 1
mysql> SELECT BINARY 'a' = 'A';
-> 0
mysql> SELECT 'a' = 'a ';
-> 1
```

```
mysql> SELECT BINARY 'a' = 'a ';
-> 0
```

`BINARY str` is shorthand for `CAST(str AS BINARY)`.

The `BINARY` attribute in character column definitions has a different effect. A character column defined with the `BINARY` attribute is assigned the binary collation of the column's character set. Every character set has a binary collation. For example, the binary collation for the `latin1` character set is `latin1_bin`, so if the table default character set is `latin1`, these two column definitions are equivalent:

```
CHAR(10) BINARY
CHAR(10) CHARACTER SET latin1 COLLATE latin1_bin
```

The effect of `BINARY` as a column attribute differs from its effect prior to MySQL 4.1. Formerly, `BINARY` resulted in a column that was treated as a binary string. A binary string is a string of bytes that has no character set or collation, which differs from a non-binary character string that has a binary collation. For both types of strings, comparisons are based on the numeric values of the string unit, but for non-binary strings the unit is the character and some character sets allow multi-byte characters. [Section 11.4.2, “The BINARY and VARBINARY Types”](#).

The use of `CHARACTER SET binary` in the definition of a `CHAR`, `VARCHAR`, or `TEXT` column causes the column to be treated as a binary data type. For example, the following pairs of definitions are equivalent:

```
CHAR(10) CHARACTER SET binary
BINARY(10)

VARCHAR(10) CHARACTER SET binary
VARBINARY(10)

TEXT CHARACTER SET binary
BLOB
```

## 10.5.4. Some Special Cases Where the Collation Determination Is Tricky

In the great majority of statements, it is obvious what collation MySQL uses to resolve a comparison operation. For example, in the following cases, it should be clear that the collation is the collation of column `x`:

```
SELECT x FROM T ORDER BY x;
SELECT x FROM T WHERE x = x;
SELECT DISTINCT x FROM T;
```

However, when multiple operands are involved, there can be ambiguity. For example:

```
SELECT x FROM T WHERE x = 'Y';
```

Should this query use the collation of the column `x`, or of the string literal `'Y'`?

Standard SQL resolves such questions using what used to be called “coercibility” rules. Basically, this means: Both `x` and `'Y'` have collations, so which collation takes precedence? This can be difficult to resolve, but the following rules cover most situations:

- An explicit `COLLATE` clause has a coercibility of 0. (Not coercible at all.)
- The concatenation of two strings with different collations has a coercibility of 1.

- A column's collation has a coercibility of 2.
- A “system constant” (the string returned by functions such as `USER()` or `VERSION()`) has a coercibility of 3.
- A literal's collation has a coercibility of 4.
- `NULL` or an expression that is derived from `NULL` has a coercibility of 5.

The preceding coercibility values are current as of MySQL 4.1.11. See the note later in this section for additional version-related information.

Those rules resolve ambiguities in the following manner:

- Use the collation with the lowest coercibility value.
- If both sides have the same coercibility, then it is an error if the collations aren't the same.

Examples:

<code>column1 = 'A'</code>	Use collation of <code>column1</code>
<code>column1 = 'A' COLLATE x</code>	Use collation of 'A'
<code>column1 COLLATE x = 'A' COLLATE y</code>	Error

The `COERCIBILITY()` function can be used to determine the coercibility of a string expression:

```
mysql> SELECT COERCIBILITY('A' COLLATE latin1_swedish_ci);
-> 0
mysql> SELECT COERCIBILITY(VERSION());
-> 3
mysql> SELECT COERCIBILITY('A');
-> 4
```

See [Section 12.9.3, “Information Functions”](#).

Before MySQL 4.1.11, there is no system constant or ignorable coercibility. Functions such as `USER()` have a coercibility of 2 rather than 3, and literals have a coercibility of 3 rather than 4.

## 10.5.5. Collations Must Be for the Right Character Set

Each character set has one or more collations, but each collation is associated with one and only one character set. Therefore, the following statement causes an error message because the `latin2_bin` collation is not legal with the `latin1` character set:

```
mysql> SELECT _latin1 'x' COLLATE latin2_bin;
ERROR 1253 (42000): COLLATION 'latin2_bin' is not valid
for CHARACTER SET 'latin1'
```

In some cases, expressions that worked before MySQL 4.1 fail in early versions of MySQL 4.1 if you do not take character set and collation into account. For example, before 4.1, this statement works as is:

```
mysql> SELECT SUBSTRING_INDEX(USER(), '@', 1);
+-----+
| SUBSTRING_INDEX(USER(), '@', 1) |
+-----+
| root                             |
+-----+
```

```
+-----+
```

The statement also works as is in MySQL 4.1 as of 4.1.8: In MySQL 4.1, usernames are stored using the `utf8` character set (see [Section 10.8, “UTF-8 for Metadata”](#)). The literal string '@' has the server character set (`latin1` by default). Although the character sets are different, MySQL can coerce the `latin1` string to the character set (and collation) of `USER()` without data loss. It does so, performs the substring operation, and returns a result that has a character set of `utf8`.

However, in versions of MySQL 4.1 before 4.1.8, the statement fails:

```
mysql> SELECT SUBSTRING_INDEX(USER(),'@',1);
ERROR 1267 (HY000): Illegal mix of collations
(utf8_general_ci,IMPLICIT) and (latin1_swedish_ci,COERCIBLE)
for operation 'substr_index'
```

This happens because the automatic character set conversion of '@' does not occur and the string operands have different character sets (and thus different collations):

```
mysql> SELECT COLLATION(USER()), COLLATION('@');
+-----+-----+
| COLLATION(USER()) | COLLATION('@') |
+-----+-----+
| utf8_general_ci   | latin1_swedish_ci |
+-----+-----+
```

One way to deal with this is to upgrade to MySQL 4.1.8 or later. If that is not possible, you can tell MySQL to interpret the literal string as `utf8`:

```
mysql> SELECT SUBSTRING_INDEX(USER(),_utf8'@',1);
+-----+
| SUBSTRING_INDEX(USER(),_utf8'@',1) |
+-----+
| root                               |
+-----+
```

Another way is to change the connection character set and collation to `utf8`. You can do that with `SET NAMES 'utf8'` or by setting the `character_set_connection` and `collation_connection` system variables directly.

## 10.5.6. An Example of the Effect of Collation

Suppose that column `X` in table `T` has these `latin1` column values:

```
Muffler
Müller
MX Systems
MySQL
```

Suppose also that the column values are retrieved using the following statement:

```
SELECT X FROM T ORDER BY X COLLATE collation_name;
```

The following table shows the resulting order of the values if we use `ORDER BY` with different collations:

<code>latin1_swedish_ci</code>	<code>latin1_german1_ci</code>	<code>latin1_german2_ci</code>
Muffler	Muffler	Müller
MX Systems	Müller	Muffler
Müller	MX Systems	MX Systems

MySQL	MySQL	MySQL
-------	-------	-------

The character that causes the different sort orders in this example is the U with two dots over it (ü), which the Germans call “U-umlaut.”

- The first column shows the result of the `SELECT` using the Swedish/Finnish collating rule, which says that U-umlaut sorts with Y.
- The second column shows the result of the `SELECT` using the German DIN-1 rule, which says that U-umlaut sorts with U.
- The third column shows the result of the `SELECT` using the German DIN-2 rule, which says that U-umlaut sorts with UE.

## 10.6. Operations Affected by Character Set Support

This section describes operations that take character set information into account as of MySQL 4.1.

### 10.6.1. Result Strings

MySQL has many operators and functions that return a string. This section answers the question: What is the character set and collation of such a string?

For simple functions that take string input and return a string result as output, the output's character set and collation are the same as those of the principal input value. For example, `UPPER(X)` returns a string whose character string and collation are the same as that of `X`. The same applies for `INSTR()`, `LCASE()`, `LOWER()`, `LTRIM()`, `MID()`, `REPEAT()`, `REPLACE()`, `REVERSE()`, `RIGHT()`, `RPAD()`, `RTRIM()`, `SOUNDEX()`, `SUBSTRING()`, `TRIM()`, `UCASE()`, and `UPPER()`.

Note: The `REPLACE()` function, unlike all other functions, always ignores the collation of the string input and performs a case-sensitive comparison.

If a string input or function result is a binary string, the string has no character set or collation. This can be checked by using the `CHARSET()` and `COLLATION()` functions, both of which return `binary` to indicate that their argument is a binary string:

```
mysql> SELECT CHARSET(BINARY 'a'), COLLATION(BINARY 'a');
+-----+-----+
| CHARSET(BINARY 'a') | COLLATION(BINARY 'a') |
+-----+-----+
| binary              | binary                 |
+-----+-----+
```

For operations that combine multiple string inputs and return a single string output, the “aggregation rules” of standard SQL apply for determining the collation of the result:

- If an explicit `COLLATE X` occurs, use `X`.
- If explicit `COLLATE X` and `COLLATE Y` occur, raise an error.
- Otherwise, if all collations are `X`, use `X`.
- Otherwise, the result has no collation.

For example, with `CASE ... WHEN a THEN b WHEN b THEN c COLLATE X END`, the resulting collation is *X*. The same applies for `UNION`, `|`, `CONCAT()`, `ELT()`, `GREATEST()`, `IF()`, and `LEAST()`.

For operations that convert to character data, the character set and collation of the strings that result from the operations are defined by the `character_set_connection` and `collation_connection` system variables. This applies to `CAST()`, `CHAR()`, `CONV()`, `FORMAT()`, `HEX()`, and `SPACE()`.

## 10.6.2. CONVERT() and CAST()

`CONVERT()` provides a way to convert data between different character sets. The syntax is:

```
CONVERT(expr USING transcoding_name)
```

In MySQL, transcoding names are the same as the corresponding character set names.

Examples:

```
SELECT CONVERT(_latin1'Müller' USING utf8);
INSERT INTO utf8table (utf8column)
  SELECT CONVERT(latin1field USING utf8) FROM latin1table;
```

`CONVERT(... USING ...)` is implemented according to the standard SQL specification.

You may also use `CAST()` to convert a string to a different character set. The syntax is:

```
CAST(character_string AS character_data_type CHARACTER SET charset_name)
```

Example:

```
SELECT CAST(_latin1'test' AS CHAR CHARACTER SET utf8);
```

If you use `CAST()` without specifying `CHARACTER SET`, the resulting character set and collation are defined by the `character_set_connection` and `collation_connection` system variables. If you use `CAST()` with `CHARACTER SET X`, the resulting character set and collation are *X* and the default collation of *X*.

You may not use a `COLLATE` clause inside a `CAST()`, but you may use it outside. That is, `CAST(... COLLATE ...)` is illegal, but `CAST(...) COLLATE ...` is legal.

Example:

```
SELECT CAST(_latin1'test' AS CHAR CHARACTER SET utf8) COLLATE utf8_bin;
```

## 10.6.3. SHOW Statements and INFORMATION\_SCHEMA

Several `SHOW` statements are added or modified in MySQL 4.1 to provide additional character set information. `SHOW CHARACTER SET`, `SHOW COLLATION`, and `SHOW CREATE DATABASE` are new. `SHOW CREATE TABLE` and `SHOW COLUMNS` are modified. These statements are described here briefly. For more information, see [Section 13.5.4](#), “`SHOW Syntax`”.

The `SHOW CHARACTER SET` command shows all available character sets. It takes an optional `LIKE` clause that indicates which character set names to match. For example:

```
mysql> SHOW CHARACTER SET LIKE 'latin%';
+-----+-----+-----+-----+
| Charset | Description                  | Default collation | Maxlen |
+-----+-----+-----+-----+
+-----+-----+-----+-----+
```

latin1	cp1252 West European	latin1_swedish_ci	1
latin2	ISO 8859-2 Central European	latin2_general_ci	1
latin5	ISO 8859-9 Turkish	latin5_turkish_ci	1
latin7	ISO 8859-13 Baltic	latin7_general_ci	1

The output from `SHOW COLLATION` includes all available character sets. It takes an optional `LIKE` clause that indicates which collation names to match. For example:

```
mysql> SHOW COLLATION LIKE 'latin1%';
```

Collation	Charset	Id	Default	Compiled	Sortlen
latin1_german1_ci	latin1	5			0
latin1_swedish_ci	latin1	8	Yes	Yes	0
latin1_danish_ci	latin1	15			0
latin1_german2_ci	latin1	31		Yes	2
latin1_bin	latin1	47		Yes	0
latin1_general_ci	latin1	48			0
latin1_general_cs	latin1	49			0
latin1_spanish_ci	latin1	94			0

`SHOW CREATE DATABASE` displays the `CREATE DATABASE` statement that creates a given database:

```
mysql> SHOW CREATE DATABASE test;
```

Database	Create Database
test	CREATE DATABASE `test` /*!40100 DEFAULT CHARACTER SET latin1 */

If no `COLLATE` clause is shown, the default collation for the character set applies.

`SHOW CREATE TABLE` is similar, but displays the `CREATE TABLE` statement to create a given table. The column definitions indicate any character set specifications, and the table options include character set information.

The `SHOW COLUMNS` statement displays the collations of a table's columns when invoked as `SHOW FULL COLUMNS`. Columns with `CHAR`, `VARCHAR`, or `TEXT` data types have collations. Numeric and other non-character types have no collation (indicated by `NULL` as the `Collation` value). For example:

```
mysql> SHOW FULL COLUMNS FROM person\G
***** 1. row *****
Field: id
Type: smallint(5) unsigned
Collation: NULL
Null: NO
Key: PRI
Default: NULL
Extra: auto_increment
Privileges: select,insert,update,references
Comment:
***** 2. row *****
Field: name
Type: char(60)
Collation: latin1_swedish_ci
Null: NO
Key:
Default:
Extra:
Privileges: select,insert,update,references
Comment:
```

The character set is not part of the display but is implied by the collation name.

## 10.7. Unicode Support

As of MySQL version 4.1, there are two new character sets for storing Unicode data:

- `ucs2`, the UCS-2 Unicode character set.
- `utf8`, the UTF-8 encoding of the Unicode character set.

In UCS-2 (binary Unicode representation), every character is represented by a two-byte Unicode code with the most significant byte first. For example: `LATIN CAPITAL LETTER A` has the code `0x0041` and it is stored as a two-byte sequence: `0x00 0x41`. `CYRILLIC SMALL LETTER YERU` (Unicode `0x044B`) is stored as a two-byte sequence: `0x04 0x4B`. For Unicode characters and their codes, please refer to the [Unicode Home Page](http://www.unicode.org/) [http://www.unicode.org/].

Currently, UCS-2 cannot be used as a client character set. That means that `SET NAMES 'ucs2'` does not work.

The UTF-8 character set (transform Unicode representation) is an alternative way to store Unicode data. It is implemented according to RFC 3629. The idea of the UTF-8 character set is that various Unicode characters are encoded using byte sequences of different lengths:

- Basic Latin letters, digits, and punctuation signs use one byte.
- Most European and Middle East script letters fit into a two-byte sequence: extended Latin letters (with tilde, macron, acute, grave and other accents), Cyrillic, Greek, Armenian, Hebrew, Arabic, Syriac, and others.
- Korean, Chinese, and Japanese ideographs use three-byte sequences.

RFC 3629 describes encoding sequences that take from one to four bytes. Currently, MySQL support for UTF-8 does not include four-byte sequences. (An older standard for UTF-8 encoding is given by RFC 2279, which describes UTF-8 sequences that take from one to six bytes. RFC 3629 renders RFC 2279 obsolete; for this reason, sequences with five and six bytes are no longer used.)

**Tip:** To save space with UTF-8, use `VARCHAR` instead of `CHAR`. Otherwise, MySQL must reserve three bytes for each character in a `CHAR CHARACTER SET utf8` column because that is the maximum possible length. For example, MySQL must reserve 30 bytes for a `CHAR(10) CHARACTER SET utf8` column.

## 10.8. UTF-8 for Metadata

*Metadata* is “the data about the data.” Anything that *describes* the database — as opposed to being the *contents* of the database — is metadata. Thus column names, database names, usernames, version names, and most of the string results from `SHOW` are metadata.

Representation of metadata must satisfy these requirements:

- All metadata must be in the same character set. Otherwise, `SHOW` wouldn't work properly because different rows in the same column would be in different character sets.
- Metadata must include all characters in all languages. Otherwise, users would not be able to name columns and tables using their own languages.



To satisfy both requirements, MySQL stores metadata in a Unicode character set, namely UTF-8. This does not cause any disruption if you never use accented or non-Latin characters. But if you do, you should be aware that metadata is in UTF-8.

The metadata requirements mean that the return values of the `USER()`, `CURRENT_USER()`, `SESSION_USER()`, `SYSTEM_USER()`, `DATABASE()`, and `VERSION()` functions have the UTF-8 character set by default.

The server sets the `character_set_system` system variable to the name of the metadata character set:

```
mysql> SHOW VARIABLES LIKE 'character_set_system';
+-----+-----+
| Variable_name | Value |
+-----+-----+
| character_set_system | utf8 |
+-----+-----+
```

Storage of metadata using Unicode does *not* mean that the server returns headers of columns and the results of `DESCRIBE` functions in the `character_set_system` character set by default. When you use `SELECT column1 FROM t`, the name `column1` itself is returned from the server to the client in the character set determined by the value of the `character_set_results` system variable, which has a default value of `latin1`. If you want the server to pass metadata results back in a different character set, use the `SET NAMES` statement to force the server to perform character set conversion. `SET NAMES` sets the `character_set_results` and other related system variables. (See [Section 10.4, “Connection Character Sets and Collations”](#).) Alternatively, a client program can perform the conversion after receiving the result from the server. It is more efficient for the client to perform the conversion, but this option is not available for many clients until late in the MySQL 4.x product cycle.

If `character_set_results` is set to `NULL`, no conversion is performed and the server returns metadata using its original character set (the set indicated by `character_set_system`).

Beginning with MySQL 4.1.1, error messages returned from the server to the client are converted to the client character set automatically, as with metadata.

If you are using (for example) the `USER()` function for comparison or assignment within a single statement, don't worry. MySQL performs some automatic conversion for you.

```
SELECT * FROM Table1 WHERE USER() = latin1_column;
```

This works because the contents of `latin1_column` are automatically converted to UTF-8 before the comparison.

```
INSERT INTO Table1 (latin1_column) SELECT USER();
```

This works because the contents of `USER()` are automatically converted to `latin1` before the assignment. Automatic conversion is not fully implemented yet, but should work correctly in a later version.

Although automatic conversion is not in the SQL standard, the SQL standard document does say that every character set is (in terms of supported characters) a “subset” of Unicode. Because it is a well-known principle that “what applies to a superset can apply to a subset,” we believe that a collation for Unicode can apply for comparisons with non-Unicode strings.

## 10.9. Upgrading Character Sets from MySQL 4.0

What about upgrading from older versions of MySQL? MySQL 4.1 is almost upward compatible with MySQL 4.0 and earlier for the simple reason that almost all the features are new, so there's nothing in earlier versions to conflict with. However, there are some differences and a few things to be aware of.

It is important to note that the “MySQL 4.0 character set” contains both character set and collation information in one single entity. Beginning in MySQL 4.1, character sets and collations are separate entities. Though each collation corresponds to a particular character set, the two are not bundled together.

There is a special treatment of national character sets in MySQL 4.1. `NCHAR` is not the same as `CHAR`, and `N'...'` literals are not the same as `'...'` literals.

Finally, there is a different file format for storing information about character sets and collations. Make sure that you have reinstalled the `/share/mysql/charsets/` directory containing the new configuration files.

If you want to start `mysqld` from a 4.1.x distribution with data created by MySQL 4.0, you should start the server with the same character set and collation. In this case, you won't need to reindex your data.

There are two ways to do so:

```
shell> ./configure --with-charset=... --with-collation=...
shell> ./mysqld --default-character-set=... --default-collation=...
```

If you used `mysqld` with, for example, the MySQL 4.0 `danish` character set, you should use the `latin1` character set and the `latin1_danish_ci` collation:

```
shell> ./configure --with-charset=latin1 \
--with-collation=latin1_danish_ci
shell> ./mysqld --default-character-set=latin1 \
--default-collation=latin1_danish_ci
```

Use the table shown in [Section 10.9.1, “4.0 Character Sets and Corresponding 4.1 Character Set/Collation Pairs”](#), to find old 4.0 character set names and their 4.1 character set/collation pair equivalents.

If you have non-`latin1` data stored in a 4.0 `latin1` table and want to convert the table column definitions to reflect the actual character set of the data, use the instructions in [Section 10.9.2, “Converting 4.0 Character Columns to 4.1 Format”](#).

### 10.9.1. 4.0 Character Sets and Corresponding 4.1 Character Set/Collation Pairs

ID	4.0 Character Set	4.1 Character Set	4.1 Collation
1	<code>big5</code>	<code>big5</code>	<code>big5_chinese_ci</code>
2	<code>czech</code>	<code>latin2</code>	<code>latin2_czech_ci</code>
3	<code>dec8</code>	<code>dec8</code>	<code>dec8_swedish_ci</code>
4	<code>dos</code>	<code>cp850</code>	<code>cp850_general_ci</code>
5	<code>german1</code>	<code>latin1</code>	<code>latin1_german1_ci</code>
6	<code>hp8</code>	<code>hp8</code>	<code>hp8_english_ci</code>
7	<code>koi8_ru</code>	<code>koi8r</code>	<code>koi8r_general_ci</code>
8	<code>latin1</code>	<code>latin1</code>	<code>latin1_swedish_ci</code>
9	<code>latin2</code>	<code>latin2</code>	<code>latin2_general_ci</code>
10	<code>swe7</code>	<code>swe7</code>	<code>swe7_swedish_ci</code>
11	<code>usa7</code>	<code>ascii</code>	<code>ascii_general_ci</code>
12	<code>ujis</code>	<code>ujis</code>	<code>ujis_japanese_ci</code>
13	<code>sjis</code>	<code>sjis</code>	<code>sjis_japanese_ci</code>
14	<code>cp1251</code>	<code>cp1251</code>	<code>cp1251_bulgarian_ci</code>

15	danish	latin1	latin1_danish_ci
16	hebrew	hebrew	hebrew_general_ci
17	win1251	(removed)	(removed)
18	tis620	tis620	tis620_thai_ci
19	euc_kr	euckr	euckr_korean_ci
20	estonia	latin7	latin7_estonian_ci
21	hungarian	latin2	latin2_hungarian_ci
22	koi8_ukr	koi8u	koi8u_ukrainian_ci
23	win1251ukr	cp1251	cp1251_ukrainian_ci
24	gb2312	gb2312	gb2312_chinese_ci
25	greek	greek	greek_general_ci
26	win1250	cp1250	cp1250_general_ci
27	croat	latin2	latin2_croatian_ci
28	gbk	gbk	gbk_chinese_ci
29	cp1257	cp1257	cp1257_lithuanian_ci
30	latin5	latin5	latin5_turkish_ci
31	latin1_de	latin1	latin1_german2_ci

## 10.9.2. Converting 4.0 Character Columns to 4.1 Format

Normally, the server runs using the `latin1` character set by default. If you have been storing column data that actually is in some other character set that the 4.1 server supports directly, you can convert the column. However, you should avoid trying to convert directly from `latin1` to the "real" character set. This may result in data loss. Instead, convert the column to a binary data type, and then from the binary type to a non-binary type with the desired character set. Conversion to and from binary involves no attempt at character value conversion and preserves your data intact. For example, suppose that you have a 4.0 table with three columns that are used to store values represented in `latin1`, `latin2`, and `utf8`:

```
CREATE TABLE t
(
  latin1_col CHAR(50),
  latin2_col CHAR(100),
  utf8_col CHAR(150)
);
```

For MySQL 4.1, you want to convert this table to leave `latin1_col` alone but change the `latin2_col` and `utf8_col` columns to have character sets of `latin2` and `utf8`. Before upgrading to 4.1, back up your table, then convert the columns as follows:

```
ALTER TABLE t MODIFY latin2_col BINARY(100);
ALTER TABLE t MODIFY utf8_col BINARY(150);
```

Then, after upgrading to 4.1, complete the conversion by issuing these statements:

```
ALTER TABLE t MODIFY latin2_col CHAR(100) CHARACTER SET latin2;
ALTER TABLE t MODIFY utf8_col CHAR(150) CHARACTER SET utf8;
```

The first two statements "remove" the character set information from the `latin2_col` and `utf8_col` columns. The second two statements assign the proper character sets to the two columns.

If you like, you can combine the to-binary conversions and from-binary conversions into single state-

ments. In MySQL 4.0, do this:

```
ALTER TABLE t
  MODIFY latin2_col BINARY(100),
  MODIFY utf8_col BINARY(150);
```

After upgrading to 4.1, do this:

```
ALTER TABLE t
  MODIFY latin2_col CHAR(100) CHARACTER SET latin2,
  MODIFY utf8_col CHAR(150) CHARACTER SET utf8;
```

If you can ensure that the tables will not otherwise be modified before you perform the character set conversion, you can issue all of the `ALTER TABLE` statements after upgrading to MySQL 4.1.

If you specified attributes when creating a column initially, you should also specify them when altering the table with `ALTER TABLE`. For example, if you specified `NOT NULL` and an explicit `DEFAULT` value, you should also provide them in the `ALTER TABLE` statement. Otherwise, the resulting column definition will not include those attributes.

## 10.10. Character Sets and Collations That MySQL Supports

MySQL supports 70+ collations for 30+ character sets. This section indicates which character sets MySQL supports. There is one subsection for each group of related character sets. For each character set, the allowable collations are listed.

You can always list the available character sets and their default collations with the `SHOW CHARACTER SET` statement:

```
mysql> SHOW CHARACTER SET;
```

Charset	Description	Default collation
big5	Big5 Traditional Chinese	big5_chinese_ci
dec8	DEC West European	dec8_swedish_ci
cp850	DOS West European	cp850_general_ci
hp8	HP West European	hp8_english_ci
koi8r	KOI8-R Relcom Russian	koi8r_general_ci
latin1	cp1252 West European	latin1_swedish_ci
latin2	ISO 8859-2 Central European	latin2_general_ci
swe7	7bit Swedish	swe7_swedish_ci
ascii	US ASCII	ascii_general_ci
ujis	EUC-JP Japanese	ujis_japanese_ci
sjis	Shift-JIS Japanese	sjis_japanese_ci
hebrew	ISO 8859-8 Hebrew	hebrew_general_ci
tis620	TIS620 Thai	tis620_thai_ci
euckr	EUC-KR Korean	euckr_korean_ci
koi8u	KOI8-U Ukrainian	koi8u_general_ci
gb2312	GB2312 Simplified Chinese	gb2312_chinese_ci
greek	ISO 8859-7 Greek	greek_general_ci
cp1250	Windows Central European	cp1250_general_ci
gbk	GBK Simplified Chinese	gbk_chinese_ci
latin5	ISO 8859-9 Turkish	latin5_turkish_ci
armscii8	ARMScii8 Armenian	armscii8_general_ci
utf8	UTF-8 Unicode	utf8_general_ci
ucs2	UCS-2 Unicode	ucs2_general_ci
cp866	DOS Russian	cp866_general_ci
keybcs2	DOS Kamenicky Czech-Slovak	keybcs2_general_ci
macce	Mac Central European	macce_general_ci
macroman	Mac West European	macroman_general_ci
cp852	DOS Central European	cp852_general_ci
latin7	ISO 8859-13 Baltic	latin7_general_ci
cp1251	Windows Cyrillic	cp1251_general_ci
cp1256	Windows Arabic	cp1256_general_ci
cp1257	Windows Baltic	cp1257_general_ci
binary	Binary pseudo charset	binary
geostd8	GEOSTD8 Georgian	geostd8_general_ci

cp932	SJIS for Windows Japanese	cp932_japanese_ci
eucjpms	UJIS for Windows Japanese	eucjpms_japanese_ci

## 10.10.1. Unicode Character Sets

MySQL has two Unicode character sets. You can store text in about 650 languages using these character sets.

- `ucs2` (UCS-2 Unicode) collations:
  - `ucs2_bin`
  - `ucs2_czech_ci`
  - `ucs2_danish_ci`
  - `ucs2_estonian_ci`
  - `ucs2_general_ci` (default)
  - `ucs2_icelandic_ci`
  - `ucs2_latvian_ci`
  - `ucs2_lithuanian_ci`
  - `ucs2_persian_ci`
  - `ucs2_polish_ci`
  - `ucs2_roman_ci`
  - `ucs2_romanian_ci`
  - `ucs2_slovak_ci`
  - `ucs2_slovenian_ci`
  - `ucs2_spanish2_ci`
  - `ucs2_spanish_ci`
  - `ucs2_swedish_ci`
  - `ucs2_turkish_ci`
  - `ucs2_unicode_ci`
- `utf8` (UTF-8 Unicode) collations:
  - `utf8_bin`
  - `utf8_czech_ci`
  - `utf8_danish_ci`
  - `utf8_estonian_ci`

- `utf8_general_ci` (default)
- `utf8_icelandic_ci`
- `utf8_latvian_ci`
- `utf8_lithuanian_ci`
- `utf8_persian_ci`
- `utf8_polish_ci`
- `utf8_roman_ci`
- `utf8_romanian_ci`
- `utf8_slovak_ci`
- `utf8_slovenian_ci`
- `utf8_spanish2_ci`
- `utf8_spanish_ci`
- `utf8_swedish_ci`
- `utf8_turkish_ci`
- `utf8_unicode_ci`

Note that in the `ucs2_roman_ci` and `utf8_roman_ci` collations, **I** and **J** are given the same precedence, and **U** and **V** are given the same precedence.

There is a limitation in MySQL 4.1 that results in two characters not being correctly handled when a user tries to change their case using `LOWER()` or `UPPER()`:

- LATIN SMALL LETTER DOTLESS **i**
- LATIN CAPITAL LETTER **I** WITH DOT ABOVE

Here are two workarounds for MySQL 4.1:

1. Use UCS2 if you have Turkish data.
2. Use these function calls:

```
CONVERT(LOWER(CONVERT(col USING ucs2)) USING utf8)
```

MySQL implements the `utf8_unicode_ci` collation according to the Unicode Collation Algorithm (UCA) described at <http://www.unicode.org/reports/tr10/>. The collation uses the version-4.0.0 UCA weight keys: <http://www.unicode.org/Public/UCA/4.0.0/allkeys-4.0.0.txt>. The following discussion uses `utf8_unicode_ci`, but it is also true for `ucs2_unicode_ci`.

Currently, the `utf8_unicode_ci` collation has only partial support for the Unicode Collation Al-

gorithm. Some characters are not supported yet. Also, combining marks are not fully supported. This affects primarily Vietnamese and some minority languages in Russia such as Udmurt, Tatar, Bashkir, and Mari.

The most significant feature in `utf8_unicode_ci` is that it supports expansions; that is, when one character compares as equal to combinations of other characters. For example, in German and some other languages ‘ß’ is equal to ‘ss’.

`utf8_general_ci` is a legacy collation that does not support expansions. It can make only one-to-one comparisons between characters. This means that comparisons for the `utf8_general_ci` collation are faster, but slightly less correct, than comparisons for `utf8_unicode_ci`.

For example, the following equalities hold in both `utf8_general_ci` and `utf8_unicode_ci`:

```
Ä = A
Ö = O
Ü = U
```

A difference between the collations is that this is true for `utf8_general_ci`:

```
ß = s
```

Whereas this is true for `utf8_unicode_ci`:

```
ß = ss
```

MySQL implements language-specific collations for the `utf8` character set only if the ordering with `utf8_unicode_ci` does not work well for a language. For example, `utf8_unicode_ci` works fine for German and French, so there is no need to create special `utf8` collations for these two languages.

`utf8_general_ci` also is satisfactory for both German and French, except that ‘ß’ is equal to ‘s’, and not to ‘ss’. If this is acceptable for your application, then you should use `utf8_general_ci` because it is faster. Otherwise, use `utf8_unicode_ci` because it is more accurate.

`utf8_swedish_ci`, like other `utf8` language-specific collations, is derived from `utf8_unicode_ci` with additional language rules. For example, in Swedish, the following relationship holds, which is not something expected by a German or French speaker:

```
Û = Y < Ö
```

The `utf8_spanish_ci` and `utf8_spanish2_ci` collations correspond to modern Spanish and traditional Spanish, respectively. In both collations, ‘ñ’ (n-tilde) is a separate letter between ‘n’ and ‘o’. In addition, for traditional Spanish, ‘ch’ is a separate letter between ‘c’ and ‘d’, and ‘ll’ is a separate letter between ‘l’ and ‘m’.

## 10.10.2. West European Character Sets

Western European character sets cover most West European languages, such as French, Spanish, Catalan, Basque, Portuguese, Italian, Albanian, Dutch, German, Danish, Swedish, Norwegian, Finnish, Faroese, Icelandic, Irish, Scottish, and English.

- `ascii` (US ASCII) collations:
  - `ascii_bin`
  - `ascii_general_ci` (default)

- `cp850` (DOS West European) collations:
  - `cp850_bin`
  - `cp850_general_ci` (default)
- `dec8` (DEC Western European) collations:
  - `dec8_bin`
  - `dec8_swedish_ci` (default)
- `hp8` (HP Western European) collations:
  - `hp8_bin`
  - `hp8_english_ci` (default)
- `latin1` (cp1252 West European) collations:
  - `latin1_bin`
  - `latin1_danish_ci`
  - `latin1_general_ci`
  - `latin1_general_cs`
  - `latin1_german1_ci`
  - `latin1_german2_ci`
  - `latin1_spanish_ci`
  - `latin1_swedish_ci` (default)

`latin1` is the default character set. MySQL's `latin1` is the same as the Windows `cp1252` character set. This means it is the same as the official [ISO 8859-1](#) or IANA (Internet Assigned Numbers Authority) `latin1`, but IANA `latin1` treats the code points between `0x80` and `0x9f` as “undefined,” whereas `cp1252`, and therefore MySQL's `latin1`, assign characters for those positions. For example, `0x80` is the Euro sign. For the “undefined” entries in `cp1252`, MySQL translates `0x81` to Unicode `0x0081`, `0x8d` to `0x008d`, `0x8f` to `0x008f`, `0x90` to `0x0090`, and `0x9d` to `0x009d`.

The `latin1_swedish_ci` collation is the default that probably is used by the majority of MySQL customers. Although it is frequently said that it is based on the Swedish/Finnish collation rules, there are Swedes and Finns who disagree with this statement.

The `latin1_german1_ci` and `latin1_german2_ci` collations are based on the DIN-1 and DIN-2 standards, where DIN stands for *Deutsches Institut für Normung* (the German equivalent of ANSI). DIN-1 is called the “dictionary collation” and DIN-2 is called the “phone book collation.”

- `latin1_german1_ci` (dictionary) rules:

```
Ä = A
Ö = O
Ü = U
ß = s
```

- `latin1_german2_ci` (phone-book) rules:



```
Ä = AE
Ö = OE
Û = UE
ß = ss
```

In the `latin1_spanish_ci` collation, ‘ñ’ (n-tilde) is a separate letter between ‘n’ and ‘o’.

- `macroman` (Mac West European) collations:
  - `macroman_bin`
  - `macroman_general_ci` (default)
- `swe7` (7bit Swedish) collations:
  - `swe7_bin`
  - `swe7_swedish_ci` (default)

### 10.10.3. Central European Character Sets

MySQL provides some support for character sets used in the Czech Republic, Slovakia, Hungary, Romania, Slovenia, Croatia, and Poland.

- `cp1250` (Windows Central European) collations:
  - `cp1250_bin`
  - `cp1250_croatian_ci`
  - `cp1250_czech_cs`
  - `cp1250_general_ci` (default)
- `cp852` (DOS Central European) collations:
  - `cp852_bin`
  - `cp852_general_ci` (default)
- `keybcs2` (DOS Kamenicky Czech-Slovak) collations:
  - `keybcs2_bin`
  - `keybcs2_general_ci` (default)
- `latin2` (ISO 8859-2 Central European) collations:
  - `latin2_bin`
  - `latin2_croatian_ci`
  - `latin2_czech_cs`
  - `latin2_general_ci` (default)

- `latin2_hungarian_ci`
- `macce` (Mac Central European) collations:
  - `macce_bin`
  - `macce_general_ci` (default)

## 10.10.4. South European and Middle East Character Sets

South European and Middle Eastern character sets supported by MySQL include Armenian, Arabic, Georgian, Greek, Hebrew, and Turkish.

- `armscii8` (ARMScii-8 Armenian) collations:
  - `armscii8_bin`
  - `armscii8_general_ci` (default)
- `cp1256` (Windows Arabic) collations:
  - `cp1256_bin`
  - `cp1256_general_ci` (default)
- `geostd8` (GEOSTD8 Georgian) collations:
  - `geostd8_bin`
  - `geostd8_general_ci` (default)
- `greek` (ISO 8859-7 Greek) collations:
  - `greek_bin`
  - `greek_general_ci` (default)
- `hebrew` (ISO 8859-8 Hebrew) collations:
  - `hebrew_bin`
  - `hebrew_general_ci` (default)
- `latin5` (ISO 8859-9 Turkish) collations:
  - `latin5_bin`
  - `latin5_turkish_ci` (default)

## 10.10.5. Baltic Character Sets

The Baltic character sets cover Estonian, Latvian, and Lithuanian languages.

- `cp1257` (Windows Baltic) collations:

- `cp1257_bin`
- `cp1257_general_ci` (default)
- `cp1257_lithuanian_ci`
- `latin7` (ISO 8859-13 Baltic) collations:
  - `latin7_bin`
  - `latin7_estonian_cs`
  - `latin7_general_ci` (default)
  - `latin7_general_cs`

### 10.10.6. Cyrillic Character Sets

The Cyrillic character sets and collations are for use with Belarusian, Bulgarian, Russian, and Ukrainian languages.

- `cp1251` (Windows Cyrillic) collations:
  - `cp1251_bin`
  - `cp1251_bulgarian_ci`
  - `cp1251_general_ci` (default)
  - `cp1251_general_cs`
  - `cp1251_ukrainian_ci`
- `cp866` (DOS Russian) collations:
  - `cp866_bin`
  - `cp866_general_ci` (default)
- `koi8r` (KOI8-R Relcom Russian) collations:
  - `koi8r_bin`
  - `koi8r_general_ci` (default)
- `koi8u` (KOI8-U Ukrainian) collations:
  - `koi8u_bin`
  - `koi8u_general_ci` (default)

### 10.10.7. Asian Character Sets

The Asian character sets that we support include Chinese, Japanese, Korean, and Thai. These can be complicated. For example, the Chinese sets must allow for thousands of different characters. See [Sec-](#)

tion 10.10.7.1, “The `cp932` Character Set”, for additional information about the `cp932` and `sjis` character sets.

- `big5` (Big5 Traditional Chinese) collations:
  - `big5_bin`
  - `big5_chinese_ci` (default)
- `cp932` (SJIS for Windows Japanese) collations:
  - `cp932_bin`
  - `cp932_japanese_ci` (default)
- `ujis` (EUC-JP Japanese) collations:
  - `ujis_bin`
  - `ujis_japanese_ci` (default)
- `eucjpms` (UJIS for Windows Japanese) collations:
  - `eucjpms_bin`
  - `eucjpms_japanese_ci` (default)
- `euckr` (EUC-KR Korean) collations:
  - `euckr_bin`
  - `euckr_korean_ci` (default)
- `gb2312` (GB2312 Simplified Chinese) collations:
  - `gb2312_bin`
  - `gb2312_chinese_ci` (default)
- `gbk` (GBK Simplified Chinese) collations:
  - `gbk_bin`
  - `gbk_chinese_ci` (default)
- `sjis` (Shift-JIS Japanese) collations:
  - `sjis_bin`
  - `sjis_japanese_ci` (default)
- `tis620` (TIS620 Thai) collations:
  - `tis620_bin`
  - `tis620_thai_ci` (default)
- `ujis` (EUC-JP Japanese) collations:
  - `ujis_bin`
  - `ujis_japanese_ci` (default)

### 10.10.7.1. The `cp932` Character Set

### Why is `cp932` needed?

In MySQL, the `sjis` character set corresponds to the `Shift_JIS` character set defined by IANA, which supports JIS X0201 and JIS X0208 characters. (See <http://www.iana.org/assignments/character-sets>.)

However, the meaning of “SHIFT JIS” as a descriptive term has become very vague and it often includes the extensions to `Shift_JIS` that are defined by various vendors.

For example, “SHIFT JIS” used in Japanese Windows environments is a Microsoft extension of `Shift_JIS` and its exact name is `Microsoft Windows Codepage : 932` or `cp932`. In addition to the characters supported by `Shift_JIS`, `cp932` supports extension characters such as NEC special characters, NEC selected — IBM extended characters, and IBM extended characters.

Since MySQL 4.1, many Japanese users have experienced problems using these extension characters. These problems stem from the following factors:

- MySQL automatically converts character sets.
- Character sets are converted via Unicode (`ucs2`).
- The `sjis` character set does not support the conversion of these extension characters.
- There are several conversion rules from so-called “SHIFT JIS” to Unicode, and some characters are converted to Unicode differently depending on the conversion rule. MySQL supports only one of these rules (described later).

The MySQL `cp932` character set is designed to solve these problems. It is available as of MySQL 4.1.12.

Before MySQL 4.1, it was safe to use any version of “SHIFT JIS” in conjunction with the `sjis` character set. However, because MySQL supports character set conversion beginning with 4.1, it is important to separate IANA `Shift_JIS` and `cp932` into two different character sets because they provide different conversion rules.

### How does `cp932` differ from `sjis`?

The `cp932` character set differs from `sjis` in the following ways:

- `cp932` supports NEC special characters, NEC selected — IBM extended characters, and IBM selected characters.
- Some `cp932` characters have two different code points, both of which convert to the same Unicode code point. When converting from Unicode back to `cp932`, one of the code points must be selected. For this “round trip conversion,” the rule recommended by Microsoft is used. (See <http://support.microsoft.com/kb/170559/EN-US/>.)

The conversion rule works like this:

- If the character is in both JIS X 0208 and NEC special characters, use the code point of JIS X 0208.
- If the character is in both NEC special characters and IBM selected characters, use the code point of NEC special characters.
- If the character is in both IBM selected characters and NEC selected — IBM extended charac-

ters, use the code point of IBM extended characters.

The table shown at <http://www.microsoft.com/globaldev/reference/dbcs/932.htm> provides information about the Unicode values of `cp932` characters. For `cp932` table entries with characters under which a four-digit number appears, the number represents the corresponding Unicode (`ucs2`) encoding. For table entries with an underlined two-digit value appears, there is a range of `cp932` character values that begin with those two digits. Clicking such a table entry takes you to a page that displays the Unicode value for each of the `cp932` characters that begin with those digits.

The following links are of special interest. They correspond to the encodings for the following sets of characters:

- NEC special characters:

[http://www.microsoft.com/globaldev/reference/dbcs/932/932\\_87.htm](http://www.microsoft.com/globaldev/reference/dbcs/932/932_87.htm)

- NEC selected — IBM extended characters:

[http://www.microsoft.com/globaldev/reference/dbcs/932/932\\_ED.htm](http://www.microsoft.com/globaldev/reference/dbcs/932/932_ED.htm)  
[http://www.microsoft.com/globaldev/reference/dbcs/932/932\\_EE.htm](http://www.microsoft.com/globaldev/reference/dbcs/932/932_EE.htm)

- IBM selected characters:

[http://www.microsoft.com/globaldev/reference/dbcs/932/932\\_FA.htm](http://www.microsoft.com/globaldev/reference/dbcs/932/932_FA.htm)  
[http://www.microsoft.com/globaldev/reference/dbcs/932/932\\_FB.htm](http://www.microsoft.com/globaldev/reference/dbcs/932/932_FB.htm)  
[http://www.microsoft.com/globaldev/reference/dbcs/932/932\\_FC.htm](http://www.microsoft.com/globaldev/reference/dbcs/932/932_FC.htm)

For some characters, conversion to and from `ucs2` is different for `sjis` and `cp932`. The following tables illustrate these differences.

Conversion to `ucs2`:

<code>sjis/cp932</code> Value	<code>sjis</code> -> <code>ucs2</code> Conversion	<code>cp932</code> -> <code>ucs2</code> Conversion
5C	005C	005C
7E	007E	007E
815C	2015	2015
815F	005C	FF3C
8160	301C	FF5E
8161	2016	2225
817C	2212	FF0D
8191	00A2	FFE0
8192	00A3	FFE1
81CA	00AC	FFE2

Conversion from `ucs2`:

<code>ucs2</code> value	<code>ucs2</code> -> <code>sjis</code> Conversion	<code>ucs2</code> -> <code>cp932</code> Conversion
005C	815F	5C
007E	7E	7E
00A2	8191	3F

00A3	8192	3F
00AC	81CA	3F
2015	815C	815C
2016	8161	3F
2212	817C	3F
2225	3F	8161
301C	8160	3F
FF0D	3F	817C
FF3C	3F	815F
FF5E	3F	8160
FFE0	3F	8191
FFE1	3F	8192
FFE2	3F	81CA

Users of any Japanese character sets should be aware that using `-character-set-client-handshake` (or `-skip-character-set-client-handshake`) has an important effect. See [Section 5.2.1, “Command Options”](#).

---

# Chapter 11. Data Types

MySQL supports a number of data types in several categories: numeric types, date and time types, and string (character) types. This chapter first gives an overview of these data types, and then provides a more detailed description of the properties of the types in each category, and a summary of the data type storage requirements. The initial overview is intentionally brief. The more detailed descriptions later in the chapter should be consulted for additional information about particular data types, such as the allowable formats in which you can specify values.

MySQL 4.1 and up also supports extensions for handling spatial data. [Chapter 16, \*Spatial Extensions\*](#), provides information about these data types.

Several of the data type descriptions use these conventions:

- *M* indicates the maximum display width for integer types. For floating-point and fixed-point types, *M* is the total number of digits. For string types, *M* is the maximum length. The maximum allowable value of *M* depends on the data type.
- *D* applies to floating-point and fixed-point types and indicates the number of digits following the decimal point. The maximum possible value is 30, but should be no greater than *M*-2.
- Square brackets (‘[’ and ‘]’) indicate optional parts of type definitions.

## 11.1. Data Type Overview

### 11.1.1. Overview of Numeric Types

A summary of the numeric data types follows. For additional information, see [Section 11.2, “Numeric Types”](#). Storage requirements are given in [Section 11.5, “Data Type Storage Requirements”](#).

*M* indicates the maximum display width. The maximum legal display width is 255. Display width is unrelated to the range of values a type can contain, as described in [Section 11.2, “Numeric Types”](#).

If you specify `ZEROFILL` for a numeric column, MySQL automatically adds the `UNSIGNED` attribute to the column.

`SERIAL` is an alias for `BIGINT UNSIGNED NOT NULL AUTO_INCREMENT UNIQUE`.

`SERIAL DEFAULT VALUE` in the definition of an integer column is an alias for `NOT NULL AUTO_INCREMENT UNIQUE`.

**Warning:** When you use subtraction between integer values where one is of type `UNSIGNED`, the result is unsigned unless the `NO_UNSIGNED_SUBTRACTION` SQL mode is enabled. See [Section 12.8, “Cast Functions and Operators”](#).

- `BIT`

In versions of MySQL up to and including 4.1, `BIT` is a synonym for `TINYINT(1)`.

- `TINYINT[(M)] [UNSIGNED] [ZEROFILL]`

A very small integer. The signed range is `-128` to `127`. The unsigned range is `0` to `255`.

- `BOOL`, `BOOLEAN`



These types are synonyms for `TINYINT(1)`. The synonym `BOOLEAN` was added in MySQL 4.1.0. A value of zero is considered false. Non-zero values are considered true:

```
mysql> SELECT IF(0, 'true', 'false');
+-----+
| IF(0, 'true', 'false') |
+-----+
| false                   |
+-----+

mysql> SELECT IF(1, 'true', 'false');
+-----+
| IF(1, 'true', 'false') |
+-----+
| true                    |
+-----+

mysql> SELECT IF(2, 'true', 'false');
+-----+
| IF(2, 'true', 'false') |
+-----+
| true                    |
+-----+
```

However, the values `TRUE` and `FALSE` are merely aliases for `1` and `0`, respectively, as shown here:

```
mysql> SELECT IF(0 = FALSE, 'true', 'false');
+-----+
| IF(0 = FALSE, 'true', 'false') |
+-----+
| true                             |
+-----+

mysql> SELECT IF(1 = TRUE, 'true', 'false');
+-----+
| IF(1 = TRUE, 'true', 'false') |
+-----+
| true                             |
+-----+

mysql> SELECT IF(2 = TRUE, 'true', 'false');
+-----+
| IF(2 = TRUE, 'true', 'false') |
+-----+
| false                            |
+-----+

mysql> SELECT IF(2 = FALSE, 'true', 'false');
+-----+
| IF(2 = FALSE, 'true', 'false') |
+-----+
| false                            |
+-----+
```

The last two statements display the results shown because `2` is equal to neither `1` nor `0`.

We intend to implement full boolean type handling, in accordance with standard SQL, in a future MySQL release.

- `SMALLINT[(M)] [UNSIGNED] [ZEROFILL]`

A small integer. The signed range is `-32768` to `32767`. The unsigned range is `0` to `65535`.

- `MEDIUMINT[(M)] [UNSIGNED] [ZEROFILL]`

A medium-sized integer. The signed range is `-8388608` to `8388607`. The unsigned range is `0` to `16777215`.

- `INT[(M)] [UNSIGNED] [ZEROFILL]`

A normal-size integer. The signed range is `-2147483648` to `2147483647`. The unsigned range is `0` to `4294967295`.

- `INTEGER[(M)] [UNSIGNED] [ZEROFILL]`

This type is a synonym for `INT`.

- `BIGINT[(M)] [UNSIGNED] [ZEROFILL]`

A large integer. The signed range is `-9223372036854775808` to `9223372036854775807`. The unsigned range is `0` to `18446744073709551615`.

Some things you should be aware of with respect to `BIGINT` columns:

- All arithmetic is done using signed `BIGINT` or `DOUBLE` values, so you should not use unsigned big integers larger than `9223372036854775807` (63 bits) except with bit functions! If you do that, some of the last digits in the result may be wrong because of rounding errors when converting a `BIGINT` value to a `DOUBLE`.

MySQL 4.0 can handle `BIGINT` in the following cases:

- When using integers to store large unsigned values in a `BIGINT` column.
  - In `MIN(col_name)` or `MAX(col_name)`, where `col_name` refers to a `BIGINT` column.
  - When using operators (`+`, `-`, `*`, and so on) where both operands are integers.
  - You can always store an exact integer value in a `BIGINT` column by storing it using a string. In this case, MySQL performs a string-to-number conversion that involves no intermediate double-precision representation.
  - The `-`, `+`, and `*` operators use `BIGINT` arithmetic when both operands are integer values. This means that if you multiply two big integers (or results from functions that return integers), you may get unexpected results when the result is larger than `9223372036854775807`.
- `FLOAT[(M,D)] [UNSIGNED] [ZEROFILL]`

A small (single-precision) floating-point number. Allowable values are `-3.402823466E+38` to `-1.175494351E-38`, `0`, and `1.175494351E-38` to `3.402823466E+38`. These are the theoretical limits, based on the IEEE standard. The actual range might be slightly smaller depending on your hardware or operating system.

`M` is the total number of decimal digits and `D` is the number of digits following the decimal point. If `M` and `D` are omitted, values are stored to the limits allowed by the hardware. A single-precision floating-point number is accurate to approximately 7 decimal places.

`UNSIGNED`, if specified, disallows negative values.

Using `FLOAT` might give you some unexpected problems because all calculations in MySQL are done with double precision. See [Section A.5.7, “Solving Problems with No Matching Rows”](#).

- `DOUBLE[(M,D)] [UNSIGNED] [ZEROFILL]`

A normal-size (double-precision) floating-point number. Allowable values are `-1.7976931348623157E+308` to `-2.2250738585072014E-308`, `0`, and `2.2250738585072014E-308` to `1.7976931348623157E+308`. These are the theoretical limits, based on the IEEE standard. The actual range might be slightly smaller depending on your

hardware or operating system.

$M$  is the total number of decimal digits and  $D$  is the number of digits following the decimal point. If  $M$  and  $D$  are omitted, values are stored to the limits allowed by the hardware. A double-precision floating-point number is accurate to approximately 15 decimal places.

`UNSIGNED`, if specified, disallows negative values.

- `DOUBLE PRECISION[ ( $M$ ,  $D$ ) ] [ UNSIGNED ] [ ZEROFILL ], REAL[ ( $M$ ,  $D$ ) ] [ UNSIGNED ] [ ZEROFILL ]`

These types are synonyms for `DOUBLE`. Exception: If the `REAL_AS_FLOAT` SQL mode is enabled, `REAL` is a synonym for `FLOAT` rather than `DOUBLE`.

- `FLOAT( $p$ ) [ UNSIGNED ] [ ZEROFILL ]`

A floating-point number.  $p$  represents the precision in bits, but MySQL uses this value only to determine whether to use `FLOAT` or `DOUBLE` for the resulting data type. If  $p$  is from 0 to 24, the data type becomes `FLOAT` with no  $M$  or  $D$  values. If  $p$  is from 25 to 53, the data type becomes `DOUBLE` with no  $M$  or  $D$  values. The range of the resulting column is the same as for the single-precision `FLOAT` or double-precision `DOUBLE` data types described earlier in this section.

As of MySQL 3.23, this data type holds true floating-point values. In earlier MySQL versions, `FLOAT( $p$ )` always has two decimals.

`FLOAT( $p$ )` syntax is provided for ODBC compatibility.

- `DECIMAL[ ( $M$  [,  $D$ ] ) ] [ UNSIGNED ] [ ZEROFILL ]`

An unpacked fixed-point number. Behaves like a `CHAR` column; “unpacked” means the number is stored as a string, using one character for each digit of the value.  $M$  is the total number of digits and  $D$  is the number of digits after the decimal point. The decimal point and (for negative numbers) the ‘-’ sign are not counted in  $M$ , although space for them is reserved. If  $D$  is 0, values have no decimal point or fractional part. The maximum range of `DECIMAL` values is the same as for `DOUBLE`, but the actual range for a given `DECIMAL` column may be constrained by the choice of  $M$  and  $D$ . If  $D$  is omitted, the default is 0. If  $M$  is omitted, the default is 10.

`UNSIGNED`, if specified, disallows negative values.

**Note:** Before MySQL 3.23, the value of  $M$  must be large enough to include the space needed for the sign and the decimal point characters.

- `DEC[ ( $M$  [,  $D$ ] ) ] [ UNSIGNED ] [ ZEROFILL ], NUMERIC[ ( $M$  [,  $D$ ] ) ] [ UNSIGNED ] [ ZEROFILL ], FIXED[ ( $M$  [,  $D$ ] ) ] [ UNSIGNED ] [ ZEROFILL ]`

These types are synonyms for `DECIMAL`. The `FIXED` synonym was added in MySQL 4.1.0 for compatibility with other database systems.

## 11.1.2. Overview of Date and Time Types

A summary of the temporal data types follows. For additional information, see [Section 11.3, “Date and Time Types”](#). Storage requirements are given in [Section 11.5, “Data Type Storage Requirements”](#).

For the `DATETIME` and `DATE` range descriptions, “supported” means that although earlier values might work, there is no guarantee.

The `SUM()` and `AVG()` aggregate functions do not work with temporal values. (They convert the values to numbers, which loses the part after the first non-numeric character.) To work around this problem,

you can convert to numeric units, perform the aggregate operation, and convert back to a temporal value. Examples:

```
SELECT SEC_TO_TIME(SUM(TIME_TO_SEC(time_col))) FROM tbl_name;
SELECT FROM_DAYS(SUM(TO_DAYS(date_col))) FROM tbl_name;
```

- **DATE**

A date. The supported range is '1000-01-01' to '9999-12-31'. MySQL displays **DATE** values in 'YYYY-MM-DD' format, but allows you to assign values to **DATE** columns using either strings or numbers.

- **DATETIME**

A date and time combination. The supported range is '1000-01-01 00:00:00' to '9999-12-31 23:59:59'. MySQL displays **DATETIME** values in 'YYYY-MM-DD HH:MM:SS' format, but allows you to assign values to **DATETIME** columns using either strings or numbers.

- **TIMESTAMP[ (M) ]**

A timestamp. The range is '1970-01-01 00:00:00' to partway through the year 2037.

A **TIMESTAMP** column is useful for recording the date and time of an **INSERT** or **UPDATE** operation. By default, the first **TIMESTAMP** column in a table is automatically set to the date and time of the most recent operation if you do not assign it a value yourself. You can also set any **TIMESTAMP** column to the current date and time by assigning it a **NULL** value. Variations on automatic initialization and update properties are described in [Section 11.3.1.2, “TIMESTAMP Properties as of MySQL 4.1”](#).

In MySQL 4.1, **TIMESTAMP** is returned as a string with the format 'YYYY-MM-DD HH:MM:SS'. Display widths (used as described in the following paragraphs) are no longer supported; the display width is fixed at 19 characters. To obtain the value as a number, you should add **+0** to the timestamp column.

In MySQL 4.0 and earlier, **TIMESTAMP** values are displayed in **YYYYMMDDHHMMSS**, **YYMMDDH-HMMSS**, **YYYYMMDD**, or **YYMMDD** format, depending on whether *M* is 14 (or missing), 12, 8, or 6, but allows you to assign values to **TIMESTAMP** columns using either strings or numbers. The *M* argument affects only how a **TIMESTAMP** column is displayed, not storage. Its values always are stored using four bytes each. From MySQL 4.0.12, the **--new** option can be used to make the server behave as in MySQL 4.1.

Note that **TIMESTAMP (M)** columns where *M* is 8 or 14 are reported to be numbers, whereas other **TIMESTAMP (M)** columns are reported to be strings. This is just to ensure that you can reliably dump and restore the table with these types.

**Note:** The behavior of **TIMESTAMP** columns changed considerably in MySQL 4.1. For complete information on the differences with regard to this data type in MySQL 4.1 and later versions (as opposed to MySQL 4.0 and earlier versions), be sure to see [Section 11.3.1.1, “TIMESTAMP Properties Prior to MySQL 4.1”](#), and [Section 11.3.1.2, “TIMESTAMP Properties as of MySQL 4.1”](#).

- **TIME**

A time. The range is '-838:59:59' to '838:59:59'. MySQL displays **TIME** values in 'HH:MM:SS' format, but allows you to assign values to **TIME** columns using either strings or numbers.

- **YEAR[ ( 2 | 4 ) ]**

A year in two-digit or four-digit format. The default is four-digit format. In four-digit format, the allowable values are 1901 to 2155, and 0000. In two-digit format, the allowable values are 70 to 69, representing years from 1970 to 2069. MySQL displays `YEAR` values in `YYYY` format, but allows you to assign values to `YEAR` columns using either strings or numbers. The `YEAR` type is unavailable prior to MySQL 3.22.

### 11.1.3. Overview of String Types

A summary of the string data types follows. For additional information, see [Section 11.4, “String Types”](#). Storage requirements are given in [Section 11.5, “Data Type Storage Requirements”](#).

In some cases, MySQL may change a string column to a type different from that given in a `CREATE TABLE` or `ALTER TABLE` statement. See [Section 13.1.5.1, “Silent Column Specification Changes”](#).

In MySQL 4.1 and up, string data types include some features that you may not have encountered in working with versions of MySQL prior to 4.1:

- As of version 4.1, MySQL interprets length specifications in character column definitions in character units. (Before MySQL 4.1, column lengths were interpreted in bytes.) This applies to `CHAR`, `VARCHAR`, and the `TEXT` types.
- Column definitions for many string data types can include attributes that specify the character set or collation of the column. These attributes apply to the `CHAR`, `VARCHAR`, the `TEXT` types, `ENUM`, and `SET` data types:
  - The `CHARACTER SET` attribute specifies the character set, and the `COLLATE` attribute specifies a collation for the the character set. For example:

```
CREATE TABLE t
(
  c1 VARCHAR(20) CHARACTER SET utf8,
  c2 TEXT CHARACTER SET latin1 COLLATE latin1_general_cs
);
```

This table definition creates a column named `c1` that has a character set of `utf8` with the default collation for that character set, and a column named `c2` that has a character set of `latin1` and a case-sensitive collation.

`CHARSET` is a synonym for `CHARACTER SET`.

- From MySQL 4.1.0 on, the `ASCII` attribute is shorthand for `CHARACTER SET latin1`.
- From MySQL 4.1.1 on, the `UNICODE` attribute is shorthand for `CHARACTER SET ucs2`.
- As of MySQL 4.1.2, the `BINARY` attribute is shorthand for specifying the binary collation of the column character set. In this case, sorting and comparison are based on numeric character values. (Before MySQL 4.1.2, `BINARY` caused was disallowed for the `TEXT` types. For `CHAR` and `VARCHAR`, `BINARY` caused a column to store binary strings and sorting and comparison were based on numeric byte values. This is the same as using character values for single-byte character sets, but not for multi-byte character sets.)
- Character column sorting and comparison are based on the character set assigned to the column. (Before MySQL 4.1, sorting and comparison were based on the collation of the server character set.) For the `CHAR`, `VARCHAR`, `TEXT`, `ENUM`, and `SET` data types, you can declare a column with a binary collation or the `BINARY` attribute to cause sorting and comparison to use the underlying character code values rather than a lexical ordering.

[Chapter 10, Character Set Support](#), provides additional information about use of character sets in MySQL 4.1 and up.

- `[NATIONAL] CHAR(M) [CHARACTER SET charset_name] [COLLATE collation_name]`

A fixed-length string that is always right-padded with spaces to the specified length when stored. *M* represents the column length. The range of *M* is 0 to 255 characters (1 to 255 prior to MySQL 3.23).

**Note:** Trailing spaces are removed when `CHAR` values are retrieved.

In MySQL 4.1, a `CHAR` column with a length specification greater than 255 is converted to the smallest `TEXT` type that can hold values of the given length. For example, `CHAR(500)` is converted to `TEXT`, and `CHAR(200000)` is converted to `MEDIUMTEXT`. This is a compatibility feature. However, this conversion causes the column to become a variable-length column, and also affects trailing-space removal.

`CHAR` is shorthand for `CHARACTER NATIONAL CHAR` (or its equivalent short form, `NCHAR`) is the standard SQL way to define that a `CHAR` column should use some predefined character set. MySQL 4.1 and up uses `utf8` as this predefined character set. [Section 10.3.6, “National Character Set”](#).

From MySQL 4.1.2 on, the `CHAR BYTE` data type is an alias for the `BINARY` data type. This is a compatibility feature.

MySQL allows you to create a column of type `CHAR(0)`. This is useful primarily when you have to be compliant with old applications that depend on the existence of a column but that do not actually use its value. `CHAR(0)` is also quite nice when you need a column that can take only two values: A column that is defined as `CHAR(0) NULL` occupies only one bit and can take only the values `NULL` and `' '` (the empty string).

- `CHAR [CHARACTER SET charset_name] [COLLATE collation_name]`

This type is a synonym for `CHAR(1)`.

- `[NATIONAL] VARCHAR(M) [CHARACTER SET charset_name] [COLLATE collation_name]`

A variable-length string. *M* represents the maximum column length. The range of *M* is 1 to 255 before MySQL 4.0.2, and 0 to 255 as of MySQL 4.0.2.

**Note:** Trailing spaces are removed when `VARCHAR` values are stored. This differs from the standard SQL specification.

In MySQL 4.1, a `VARCHAR` column with a length specification greater than 255 is converted to the smallest `TEXT` type that can hold values of the given length. For example, `VARCHAR(500)` is converted to `TEXT`, and `VARCHAR(200000)` is converted to `MEDIUMTEXT`. This is a compatibility feature. However, this conversion affects trailing-space removal.

`VARCHAR` is shorthand for `CHARACTER VARYING`.

- `BINARY(M)`

The `BINARY` type is similar to the `CHAR` type, but stores binary byte strings rather than non-binary character strings.

This type was added in MySQL 4.1.2.

- `VARBINARY(M)`

The `VARBINARY` type is similar to the `VARCHAR` type, but stores binary byte strings rather than non-binary character strings.

This type was added in MySQL 4.1.2.

- `TINYBLOB`

A `BLOB` column with a maximum length of 255 ( $2^8 - 1$ ) bytes.

- `TINYTEXT [CHARACTER SET charset_name] [COLLATE collation_name]`

A `TEXT` column with a maximum length of 255 ( $2^8 - 1$ ) characters.

- `BLOB(M)`

A `BLOB` column with a maximum length of 65,535 ( $2^{16} - 1$ ) bytes.

Beginning with MySQL 4.1, an optional length *M* can be given for this type. MySQL creates the column as the smallest `BLOB` type large enough to hold values *M* bytes long.

- `TEXT(M) [CHARACTER SET charset_name] [COLLATE collation_name]`

A `TEXT` column with a maximum length of 65,535 ( $2^{16} - 1$ ) characters.

Beginning with MySQL 4.1, an optional length *M* can be given for this type. MySQL creates the column as the smallest `TEXT` type large enough to hold values *M* characters long.

- `MEDIUMBLOB`

A `BLOB` column with a maximum length of 16,777,215 ( $2^{24} - 1$ ) bytes.

- `MEDIUMTEXT [CHARACTER SET charset_name] [COLLATE collation_name]`

A `TEXT` column with a maximum length of 16,777,215 ( $2^{24} - 1$ ) characters.

- `LOBLOB`

A `BLOB` column with a maximum length of 4,294,967,295 or 4GB ( $2^{32} - 1$ ) bytes. Up to MySQL 3.23, the client/server protocol and `MyISAM` tables had a limit of 16MB per communication packet or table row. From MySQL 4.0, the maximum allowed length of `LOBLOB` columns depends on the configured maximum packet size in the client/server protocol and available memory.

- `LONGTEXT [CHARACTER SET charset_name] [COLLATE collation_name]`

A `TEXT` column with a maximum length of 4,294,967,295 or 4GB ( $2^{32} - 1$ ) characters. Up to MySQL 3.23, the client/server protocol and `MyISAM` tables had a limit of 16MB per communication packet or table row. From MySQL 4.0, the maximum allowed length of `LONGTEXT` columns depends on the configured maximum packet size in the client/server protocol and available memory.

- `ENUM('value1', 'value2', ...) [CHARACTER SET charset_name] [COLLATE collation_name]`

An enumeration. A string object that can have only one value, chosen from the list of values `'value1', 'value2', ..., NULL` or the special `' '` error value. An `ENUM` column can have a maximum of 65,535 distinct values. `ENUM` values are represented internally as integers.

- `SET('value1', 'value2', ...) [CHARACTER SET charset_name] [COLLATE collation_name]`

A set. A string object that can have zero or more values, each of which must be chosen from the list of values `'value1'`, `'value2'`, ... A `SET` column can have a maximum of 64 members. `SET` values are represented internally as integers.

### 11.1.4. Data Type Default Values

The `DEFAULT value` clause in a data type specification indicates a default value for a column. With one exception, the default value must be a constant; it cannot be a function or an expression. This means, for example, that you cannot set the default for a date column to be the value of a function such as `NOW()` or `CURRENT_DATE`. The exception is that you can specify `CURRENT_TIMESTAMP` as the default for a `TIMESTAMP` column as of MySQL 4.1.2. See [Section 11.3.1.2, “TIMESTAMP Properties as of MySQL 4.1”](#).

If a column definition includes no explicit `DEFAULT` value, MySQL determines the default value as follows:

If the column can take `NULL` as a value, the column is defined with an explicit `DEFAULT NULL` clause.

If the column cannot take `NULL` as the value, MySQL defines the column with an explicit `DEFAULT` clause, using the implicit default value for the column data type. Implicit defaults are defined as follows:

- For numeric types other than integer types declared with the `AUTO_INCREMENT` attribute, the default is `0`. For an `AUTO_INCREMENT` column, the default value is the next value in the sequence.
- For date and time types other than `TIMESTAMP`, the default is the appropriate “zero” value for the type. For the first `TIMESTAMP` column in a table, the default value is the current date and time. See [Section 11.3, “Date and Time Types”](#).
- For string types other than `ENUM`, the default value is the empty string. For `ENUM`, the default is the first enumeration value.

`BLOB` and `TEXT` columns cannot be assigned a default value.

For a given table, you can use the `SHOW CREATE TABLE` statement to see which columns have an explicit `DEFAULT` clause.

## 11.2. Numeric Types

MySQL supports all of the standard SQL numeric data types. These types include the exact numeric data types (`INTEGER`, `SMALLINT`, `DECIMAL`, and `NUMERIC`), as well as the approximate numeric data types (`FLOAT`, `REAL`, and `DOUBLE PRECISION`). The keyword `INT` is a synonym for `INTEGER`, and the keyword `DEC` is a synonym for `DECIMAL`. For numeric type storage requirements, see [Section 11.5, “Data Type Storage Requirements”](#).

As an extension to the SQL standard, MySQL also supports the integer types `TINYINT`, `MEDIUMINT`, and `BIGINT`. The following table shows the required storage and range for each of the integer types.

Type	Bytes	Minimum Value (Signed/Unsigned)	Maximum Value (Signed/Unsigned)
<code>TINYINT</code>	1	-128	127
		0	255
<code>SMALLINT</code>	2	-32768	32767



		0	65535
MEDIUMINT	3	-8388608	8388607
		0	16777215
INT	4	-2147483648	2147483647
		0	4294967295
BIGINT	8	-9223372036854775808	9223372036854775807
		0	18446744073709551615

Another extension is supported by MySQL for optionally specifying the display width of an integer value in parentheses following the base keyword for the type (for example, `INT(4)`). This optional display width specification is used to left-pad the display of values having a width less than the width specified for the column.

The display width does *not* constrain the range of values that can be stored in the column, nor the number of digits that are displayed for values having a width exceeding that specified for the column.

When used in conjunction with the optional extension attribute `ZEROFILL`, the default padding of spaces is replaced with zeros. For example, for a column declared as `INT(5) ZEROFILL`, a value of 4 is retrieved as `00004`. Note that if you store larger values than the display width in an integer column, you may experience problems when MySQL generates temporary tables for some complicated joins, because in these cases MySQL assumes that the data fits into the original column width.

**Note:** The `ZEROFILL` attribute is stripped when a column is involved in expressions or `UNION` queries.

All integer types can have an optional (non-standard) attribute `UNSIGNED`. Unsigned values can be used when you want to allow only non-negative numbers in a column and you need a larger upper numeric range for the column. For example, if an `INT` column is `UNSIGNED`, the size of the column's range is the same but its endpoints shift from `-2147483648` and `2147483647` up to `0` and `4294967295`.

As of MySQL 4.0.2, floating-point and fixed-point types also can be `UNSIGNED`. As with integer types, this attribute prevents negative values from being stored in the column. However, unlike the integer types, the upper range of column values remains the same.

If you specify `ZEROFILL` for a numeric column, MySQL automatically adds the `UNSIGNED` attribute to the column.

For floating-point data types, MySQL uses four bytes for single-precision values and eight bytes for double-precision values.

The `FLOAT` and `DOUBLE` data types are used to represent approximate numeric data values. For `FLOAT` the SQL standard allows an optional specification of the precision (but not the range of the exponent) in bits following the keyword `FLOAT` in parentheses. MySQL also supports this optional precision specification, but the precision value is used only to determine storage size. A precision from 0 to 23 results in a four-byte single-precision `FLOAT` column. A precision from 24 to 53 results in an eight-byte double-precision `DOUBLE` column.

MySQL allows a non-standard syntax: `FLOAT(M,D)` or `REAL(M,D)` or `DOUBLE PRECISION(M,D)`. Here, “`(M,D)`” means that values are displayed with up to `M` digits in total, of which `D` digits may be after the decimal point. For example, a column defined as `FLOAT(7,4)` will look like `-999.9999` when displayed. MySQL performs rounding when storing values, so if you insert `999.00009` into a `FLOAT(7,4)` column, the approximate result is `999.0001`.

MySQL treats `DOUBLE` as a synonym for `DOUBLE PRECISION` (a non-standard extension). MySQL also treats `REAL` as a synonym for `DOUBLE PRECISION` (a non-standard variation), unless the `REAL_AS_FLOAT` SQL mode is enabled.

For maximum portability, code requiring storage of approximate numeric data values should use `FLOAT` or `DOUBLE PRECISION` with no specification of precision or number of digits.

The `DECIMAL` and `NUMERIC` data types are used to store exact numeric data values. In MySQL, `NUMERIC` is implemented as `DECIMAL`. These types are used to store values for which it is important to preserve exact precision, for example with monetary data.

Through version 4.1, MySQL stores `DECIMAL` and `NUMERIC` values as strings, rather than in binary format. One character is used for each digit of the value, the decimal point (if the scale is greater than 0), and the ‘-’ sign (for negative numbers).

When declaring a `DECIMAL` or `NUMERIC` column, the precision and scale can be (and usually is) specified; for example:

```
salary DECIMAL(5,2)
```

In this example, 5 is the precision and 2 is the scale. The precision represents the number of significant digits that are stored for values, and the scale represents the number of digits that can be stored following the decimal point. If the scale is 0, `DECIMAL` and `NUMERIC` values contain no decimal point or fractional part.

Standard SQL requires that the `salary` column be able to store any value with five digits and two decimals. In this case, therefore, the range of values that can be stored in the `salary` column is from `-999.99` to `999.99`. In versions up to and including 4.1, MySQL varies from this limit in two ways due to the use of string format for value storage:

- On the positive end of the range, the column actually can store numbers up to `9999.99`. For positive numbers, MySQL uses the byte reserved for the sign to extend the upper end of the range.
- `DECIMAL` columns in MySQL before 3.23 are stored differently and cannot represent all the values required by standard SQL. This is because for a type of `DECIMAL(M,D)`, the value of `M` includes the bytes for the sign and the decimal point. The range of the `salary` column before MySQL 3.23 would be `-9.99` to `99.99`.

In standard SQL, the syntax `DECIMAL(M)` is equivalent to `DECIMAL(M,0)`. Similarly, the syntax `DECIMAL` is equivalent to `DECIMAL(M,0)`, where the implementation is allowed to decide the value of `M`. As of MySQL 3.23.6, both of these variant forms of the `DECIMAL` and `NUMERIC` data types are supported. The default value of `M` is 10. Before 3.23.6, `M` and `D` both must be specified explicitly.

The maximum range of `DECIMAL` and `NUMERIC` values is the same as for `DOUBLE`, but the actual range for a given `DECIMAL` or `NUMERIC` column can be constrained by the precision or scale for a given column. When such a column is assigned a value with more digits following the decimal point than are allowed by the specified scale, the value is converted to that scale. (The precise behavior is operating system-specific, but generally the effect is truncation to the allowable number of digits.)

When asked to store a value in a numeric column that is outside the data type's allowable range, MySQL clips the value to the appropriate endpoint of the range and stores the resulting value instead.

For example, when an out-of-range value is assigned to an integer column, MySQL stores the value representing the corresponding endpoint of the column data type range. If you store 256 into a `TINYINT` or `TINYINT UNSIGNED` column, MySQL stores 127 or 255, respectively. When a floating-point or fixed-point column is assigned a value that exceeds the range implied by the specified (or default) precision and scale, MySQL stores the value representing the corresponding endpoint of that range.

Conversions that occur due to clipping are reported as “warnings” for `ALTER TABLE`, `LOAD DATA INFILE`, `UPDATE`, and multiple-row `INSERT` statements.

## 11.3. Date and Time Types

The date and time types for representing temporal values are `DATETIME`, `DATE`, `TIMESTAMP`, `TIME`, and `YEAR`. Each temporal type has a range of legal values, as well as a “zero” value that is used when you specify an illegal value that MySQL cannot represent. The `TIMESTAMP` type has special automatic updating behavior, described later on. For temporary type storage requirements, see [Section 11.5, “Data Type Storage Requirements”](#).

MySQL version through 4.1 accept certain “illegal” values for dates, such as `'1999-11-31'`. This is useful when you want to store a possibly incorrect value specified by a user (for example, in a web form) in the database for future processing. MySQL verifies only that the month is in the range from 0 to 12 and that the day is in the range from 0 to 31. These ranges are defined to include zero because MySQL allows you to store dates where the day or month and day are zero in a `DATE` or `DATETIME` column. This is extremely useful for applications that need to store a birthdate for which you do not know the exact date. In this case, you simply store the date as `'1999-00-00'` or `'1999-01-00'`. If you store dates such as these, you should not expect to get correct results for functions such as `DATE_SUB()` or `DATE_ADD` that require complete dates.

MySQL also allows you to store `'0000-00-00'` as a “dummy date.” This is in some cases more convenient, and uses less data and index space, than storing `NULL` values.

Here are some general considerations to keep in mind when working with date and time types:

- MySQL retrieves values for a given date or time type in a standard output format, but it attempts to interpret a variety of formats for input values that you supply (for example, when you specify a value to be assigned to or compared to a date or time type). Only the formats described in the following sections are supported. It is expected that you supply legal values. Unpredictable results may occur if you use values in other formats.
- Dates containing two-digit year values are ambiguous because the century is unknown. MySQL interprets two-digit year values using the following rules:
  - Year values in the range `00-69` are converted to `2000-2069`.
  - Year values in the range `70-99` are converted to `1970-1999`.
- Although MySQL tries to interpret values in several formats, dates always must be given in year-month-day order (for example, `'98-09-04'`), rather than in the month-day-year or day-month-year orders commonly used elsewhere (for example, `'09-04-98'`, `'04-09-98'`).
- MySQL automatically converts a date or time type value to a number if the value is used in a numeric context and vice versa.
- By default, when MySQL encounters a value for a date or time type that is out of range or otherwise illegal for the type (as described at the beginning of this section), it converts the value to the “zero” value for that type. The exception is that out-of-range `TIME` values are clipped to the appropriate endpoint of the `TIME` range.

Data Type	“Zero” Value
<code>DATETIME</code>	<code>'0000-00-00 00:00:00'</code>
<code>DATE</code>	<code>'0000-00-00'</code>
<code>TIMESTAMP</code> (4.1 and up)	<code>'0000-00-00 00:00:00'</code>
<code>TIMESTAMP</code> (before 4.1)	<code>0000000000000000</code>
<code>TIME</code>	<code>'00:00:00'</code>
<code>YEAR</code>	<code>0000</code>

- The “zero” values are special, but you can store or refer to them explicitly using the values shown in the table. You can also do this using the values '0' or 0, which are easier to write.
- “Zero” date or time values used through MyODBC are converted automatically to `NULL` in MyODBC 2.50.12 and above, because ODBC cannot handle such values.

### 11.3.1. The `DATETIME`, `DATE`, and `TIMESTAMP` Types

The `DATETIME`, `DATE`, and `TIMESTAMP` types are related. This section describes their characteristics, how they are similar, and how they differ.

The `DATETIME` type is used when you need values that contain both date and time information. MySQL retrieves and displays `DATETIME` values in 'YYYY-MM-DD HH:MM:SS' format. The supported range is '1000-01-01 00:00:00' to '9999-12-31 23:59:59'.

The `DATE` type is used when you need only a date value, without a time part. MySQL retrieves and displays `DATE` values in 'YYYY-MM-DD' format. The supported range is '1000-01-01' to '9999-12-31'.

For the `DATETIME` and `DATE` range descriptions, “supported” means that although earlier values might work, there is no guarantee.

The `TIMESTAMP` data type has varying properties, depending on the MySQL version. These properties are described later in this section.

You can specify `DATETIME`, `DATE`, and `TIMESTAMP` values using any of a common set of formats:

- As a string in either 'YYYY-MM-DD HH:MM:SS' or 'YY-MM-DD HH:MM:SS' format. A “relaxed” syntax is allowed: Any punctuation character may be used as the delimiter between date parts or time parts. For example, '98-12-31 11:30:45', '98.12.31 11+30+45', '98/12/31 11\*30\*45', and '98@12@31 11^30^45' are equivalent.
- As a string in either 'YYYY-MM-DD' or 'YY-MM-DD' format. A “relaxed” syntax is allowed here, too. For example, '98-12-31', '98.12.31', '98/12/31', and '98@12@31' are equivalent.
- As a string with no delimiters in either 'YYYYMMDDHHMMSS' or 'YYMMDDHHMMSS' format, provided that the string makes sense as a date. For example, '19970523091528' and '970523091528' are interpreted as '1997-05-23 09:15:28', but '971122129015' is illegal (it has a nonsensical minute part) and becomes '0000-00-00 00:00:00'.
- As a string with no delimiters in either 'YYYYMMDD' or 'YYMMDD' format, provided that the string makes sense as a date. For example, '19970523' and '970523' are interpreted as '1997-05-23', but '971332' is illegal (it has nonsensical month and day parts) and becomes '0000-00-00'.
- As a number in either `YYYYMMDDHHMMSS` or `YYMMDDHHMMSS` format, provided that the number makes sense as a date. For example, 19830905132800 and 830905132800 are interpreted as '1983-09-05 13:28:00'.
- As a number in either `YYYYMMDD` or `YYMMDD` format, provided that the number makes sense as a date. For example, 19830905 and 830905 are interpreted as '1983-09-05'.
- As the result of a function that returns a value that is acceptable in a `DATETIME`, `DATE`, or `TIMESTAMP` context, such as `NOW()` or `CURRENT_DATE`.

Illegal `DATETIME`, `DATE`, or `TIMESTAMP` values are converted to the “zero” value of the appropriate type (`'0000-00-00 00:00:00'`, `'0000-00-00'`, or `0000000000000000`).

For values specified as strings that include date part delimiters, it is not necessary to specify two digits for month or day values that are less than 10. `'1979-6-9'` is the same as `'1979-06-09'`. Similarly, for values specified as strings that include time part delimiters, it is not necessary to specify two digits for hour, minute, or second values that are less than 10. `'1979-10-30 1:2:3'` is the same as `'1979-10-30 01:02:03'`.

Values specified as numbers should be 6, 8, 12, or 14 digits long. If a number is 8 or 14 digits long, it is assumed to be in `YYMMDD` or `YYMMDDHHMMSS` format and that the year is given by the first 4 digits. If the number is 6 or 12 digits long, it is assumed to be in `YYMMDD` or `YYMMDDHHMMSS` format and that the year is given by the first 2 digits. Numbers that are not one of these lengths are interpreted as though padded with leading zeros to the closest length.

Values specified as non-delimited strings are interpreted using their length as given. If the string is 8 or 14 characters long, the year is assumed to be given by the first 4 characters. Otherwise, the year is assumed to be given by the first 2 characters. The string is interpreted from left to right to find year, month, day, hour, minute, and second values, for as many parts as are present in the string. This means you should not use strings that have fewer than 6 characters. For example, if you specify `'9903'`, thinking that represents March, 1999, MySQL inserts a “zero” date into your table. This occurs because the year and month values are 99 and 03, but the day part is completely missing, so the value is not a legal date. However, as of MySQL 3.23, you can explicitly specify a value of zero to represent missing month or day parts. For example, you can use `'990300'` to insert the value `'1999-03-00'`.

You can to some extent assign values of one date type to an object of a different date type. However, there may be some alteration of the value or loss of information:

- If you assign a `DATE` value to a `DATETIME` or `TIMESTAMP` object, the time part of the resulting value is set to `'00:00:00'` because the `DATE` value contains no time information.
- If you assign a `DATETIME` or `TIMESTAMP` value to a `DATE` object, the time part of the resulting value is deleted because the `DATE` type stores no time information.
- Remember that although `DATETIME`, `DATE`, and `TIMESTAMP` values all can be specified using the same set of formats, the types do not all have the same range of values. For example, `TIMESTAMP` values cannot be earlier than 1970 or later than 2037. This means that a date such as `'1968-01-01'`, while legal as a `DATETIME` or `DATE` value, is not valid as a `TIMESTAMP` value and is converted to 0.

Be aware of certain pitfalls when specifying date values:

- The relaxed format allowed for values specified as strings can be deceiving. For example, a value such as `'10:11:12'` might look like a time value because of the `:` delimiter, but if used in a date context is interpreted as the year `'2010-11-12'`. The value `'10:45:15'` is converted to `'0000-00-00'` because `'45'` is not a legal month.
- The MySQL server performs only basic checking on the validity of a date: The ranges for year, month, and day are 1000 to 9999, 00 to 12, and 00 to 31, respectively. Any date containing parts not within these ranges is subject to conversion to `'0000-00-00'`. Please note that this still allows you to store invalid dates such as `'2002-04-31'`. To ensure that a date is valid, perform a check in your application.
- Dates containing two-digit year values are ambiguous because the century is unknown. MySQL interprets two-digit year values using the following rules:

- Year values in the range 00–69 are converted to 2000–2069.
- Year values in the range 70–99 are converted to 1970–1999.

### 11.3.1.1. **TIMESTAMP** Properties Prior to MySQL 4.1

The **TIMESTAMP** data type provides a type that you can use to automatically mark **INSERT** or **UPDATE** operations with the current date and time. If you have multiple **TIMESTAMP** columns in a table, only the first one is updated automatically. (From MySQL 4.1.2 on, you can specify which **TIMESTAMP** column updates; see [Section 11.3.1.2, “TIMESTAMP Properties as of MySQL 4.1”](#).)

Automatic updating of the first **TIMESTAMP** column in a table occurs under any of the following conditions:

- You explicitly set the column to **NULL**.
- The column is not specified explicitly in an **INSERT** or **LOAD DATA INFILE** statement.
- The column is not specified explicitly in an **UPDATE** statement and some other column changes value. An **UPDATE** that sets a column to the value it does not cause the **TIMESTAMP** column to be updated; if you set a column to its current value, MySQL ignores the update for efficiency.

A **TIMESTAMP** column other than the first also can be assigned the current date and time by setting it to **NULL** or to any function that produces the current date and time (**NOW()**, **CURRENT\_TIMESTAMP**).

Note that the information in the following discussion applies to **TIMESTAMP** columns only for tables not created with **MAXDB** mode enabled, because such columns are created as **DATETIME** columns.

You can set any **TIMESTAMP** column to a value different from the current date and time by setting it explicitly to the desired value. This is true even for the first **TIMESTAMP** column. You can use this property if, for example, you want a **TIMESTAMP** to be set to the current date and time when you create a row, but not to be changed whenever the row is updated later:

- Let MySQL set the column when the row is created. This initializes it to the current date and time.
- When you perform subsequent updates to other columns in the row, set the **TIMESTAMP** column explicitly to its current value:

```
UPDATE tbl_name
SET timestamp_col = timestamp_col,
    other_col1 = new_value1,
    other_col2 = new_value2, ...
```

Another way to maintain a column that records row-creation time is to use a **DATETIME** column that you initialize to **NOW()** when the row is created and do not modify for subsequent updates.

**TIMESTAMP** values may range from the beginning of 1970 to partway through the year 2037, with a resolution of one second. Values are displayed as numbers. When you store a value in a **TIMESTAMP** column, it is assumed to be represented in the current time zone, and is converted to UTC for storage. When you retrieve the value, it is converted from UTC back to the local time zone for display. Before MySQL 4.1.3, the server has a single time zone. As of 4.1.3, clients can set their own time zones on a per-connection basis, as described in [Section 5.10.8, “MySQL Server Time Zone Support”](#).

Prior to version 4.1, the format in which MySQL retrieves and displays `TIMESTAMP` values depends on the display size, as illustrated in the following table. The “full” `TIMESTAMP` format is 14 digits, but `TIMESTAMP` columns may be created with shorter display sizes:

Data Type	Display Format
<code>TIMESTAMP ( 14 )</code>	<code>YYYYMMDDHHMMSS</code>
<code>TIMESTAMP ( 12 )</code>	<code>YYMMDDHHMMSS</code>
<code>TIMESTAMP ( 10 )</code>	<code>YYMMDDHHMM</code>
<code>TIMESTAMP ( 8 )</code>	<code>YYYYMMDD</code>
<code>TIMESTAMP ( 6 )</code>	<code>YYMMDD</code>
<code>TIMESTAMP ( 4 )</code>	<code>YYMM</code>
<code>TIMESTAMP ( 2 )</code>	<code>YY</code>

All `TIMESTAMP` columns have the same storage size, regardless of display size. The most common display sizes are 6, 8, 12, and 14. You can specify an arbitrary display size at table creation time, but values of 0 or greater than 14 are coerced to 14. Odd-valued sizes in the range from 1 to 13 are coerced to the next higher even number.

`TIMESTAMP` columns store legal values using the full precision with which the value was specified, regardless of the display size. This has several implications:

- Always specify year, month, and day, even if your column types are `TIMESTAMP ( 4 )` or `TIMESTAMP ( 2 )`. Otherwise, the value is not a legal date and 0 is stored.
- If you use `ALTER TABLE` to widen a narrow `TIMESTAMP` column, information is displayed that previously was “hidden.”
- Similarly, narrowing a `TIMESTAMP` column does not cause information to be lost, except in the sense that less information is shown when the values are displayed.
- If you are planning to use `mysqldump` for the database, do not use `TIMESTAMP ( 4 )` or `TIMESTAMP ( 2 )`. The display format for these data types are not legal dates and 0 will be stored instead. This inconsistency is fixed starting with MySQL 4.1, where display width is ignored. To prepare for transition to versions after 4.0, you should change to use display widths of 6 or more, which will produce a legal display format. You can change the display width of `TIMESTAMP` data types, without losing any information, by using `ALTER TABLE` as indicated above.

If you need to print the timestamps for external applications, you can use `MID ( )` to extract the relevant part of the timestamp: for example, to imitate the `TIMESTAMP ( 4 )` display format.

- Although `TIMESTAMP` values are stored to full precision, the only function that operates directly on the underlying stored value is `UNIX_TIMESTAMP ( )`. Other functions operate on the formatted retrieved value. This means you cannot use a function such as `HOUR ( )` or `SECOND ( )` unless the relevant part of the `TIMESTAMP` value is included in the formatted value. For example, the `HH` part of a `TIMESTAMP` column is not displayed unless the display size is at least 10, so trying to use `HOUR ( )` on shorter `TIMESTAMP` values produces a meaningless result.

In MySQL 4.1, `TIMESTAMP` display format changes to be the same as `DATETIME`, that is, as a string in `'YYYY-MM-DD HH:MM:SS'` format rather than as a number in `YYYYMMDDHHMMSS` format. To test applications written for MySQL 4.0 for compatibility with this change, you can set the `new` system variable to 1. This variable is available beginning with MySQL 4.0.12. It can be set at server startup by specifying the `--new` option to `mysqld`. At runtime, a user who has the `SUPER` privilege can set the glob-

al value with a `SET` statement:

```
mysql> SET GLOBAL new = 1;
```

Any client can set its session value of `new` as follows:

```
mysql> SET new = 1;
```

The general effect of setting `new` to 1 is that values for a `TIMESTAMP` column display as strings rather than as numbers. Also, `DESCRIBE` displays the column definition as `timestamp(19)`, rather than as `timestamp(14)`.

However, the effect differs somewhat for `TIMESTAMP` columns that are created while `new` is set to 1. In this case, column values display as strings and `DESCRIBE` shows the definition as `timestamp(19)`, regardless of the current value of `new`.

In other words, with `new=1`, all `TIMESTAMP` values display as strings and `DESCRIBE` shows a display width of 19. For columns created while `new=1`, they continue to display as strings and to have a display width of 19 even if `new` is set to 0.

For a `TIMESTAMP` column that displays as a string, you can display it as a number by retrieving it as `col_name+0`.

### 11.3.1.2. `TIMESTAMP` Properties as of MySQL 4.1

In MySQL 4.1 and up, the properties of the `TIMESTAMP` data type change in the ways described in this section.

From MySQL 4.1.0 on, `TIMESTAMP` display format differs from that of earlier MySQL releases:

- `TIMESTAMP` columns are displayed in the same format as `DATETIME` columns. In other words, the display width is fixed at 19 characters, and the format is `YYYY-MM-DD HH:MM:SS`.
- Display widths (used as described in the preceding section) are no longer supported. In other words, for declarations such as `TIMESTAMP(2)`, `TIMESTAMP(4)`, and so on, the display width is ignored.

Beginning with MySQL 4.1.1, the MySQL server can also be run with the `MAXDB` SQL mode enabled. When the server runs with this mode enabled, `TIMESTAMP` is identical with `DATETIME`. That is, if this mode is enabled at the time that a table is created, `TIMESTAMP` columns are created as `DATE-TIME` columns. As a result, such columns use `DATETIME` display format, have the same range of values, and there is no automatic initialization or updating to the current date and time.

To enable `MAXDB` mode, set the server SQL mode to `MAXDB` at startup using the `-sql-mode=MAXDB` server option or by setting the global `sql_mode` variable at runtime:

```
mysql> SET GLOBAL sql_mode=MAXDB;
```

A client can cause the server to run in `MAXDB` mode for its own connection as follows:

```
mysql> SET SESSION sql_mode=MAXDB;
```

Beginning with MySQL 4.1.2, you have more flexible control over when automatic `TIMESTAMP` initialization and updating occur and which column should have those behaviors:



- For one `TIMESTAMP` column in a table, you can assign the current timestamp as the default value and the auto-update value. It is possible to have the current timestamp be the default value for initializing the column, for the auto-update value, or both. It is not possible to have the current timestamp be the default value for one column and the auto-update value for another column.
- You can specify which `TIMESTAMP` column to automatically initialize or update to the current date and time. This need not be the first `TIMESTAMP` column.

The following discussion describes the revised syntax and behavior. Note that this information applies only to `TIMESTAMP` columns for tables not created with `MAXDB` mode enabled. As noted earlier in this section, `MAXDB` mode causes columns to be created as `DATETIME` columns.

The following items summarize the pre-4.1.2 properties for `TIMESTAMP` initialization and updating:

The first `TIMESTAMP` column in table row automatically is set to the current timestamp when the record is created if the column is set to `NULL` or is not specified at all.

The first `TIMESTAMP` column in table row automatically is updated to the current timestamp when the value of any other column in the row is changed, unless the `TIMESTAMP` column explicitly is assigned a value other than `NULL`.

If a `DEFAULT` value is specified for the first `TIMESTAMP` column when the table is created, it is silently ignored.

Other `TIMESTAMP` columns in the table can be set to the current `TIMESTAMP` by assigning `NULL` to them, but they do not update automatically.

As of 4.1.2, you have more flexibility in deciding which `TIMESTAMP` column automatically is initialized and updated to the current timestamp. The rules are as follows:

If a `DEFAULT` value is specified for the first `TIMESTAMP` column in a table, it is not ignored. The default can be `CURRENT_TIMESTAMP` or a constant date and time value.

`DEFAULT NULL` is the same as `DEFAULT CURRENT_TIMESTAMP` for the *first* `TIMESTAMP` column. For any other `TIMESTAMP` column, `DEFAULT NULL` is treated as `DEFAULT 0`.

Any single `TIMESTAMP` column in a table can be used as the one that is initialized to the current timestamp or updated automatically.

In a `CREATE TABLE` statement, the first `TIMESTAMP` column can be declared in any of the following ways:

- With both `DEFAULT CURRENT_TIMESTAMP` and `ON UPDATE CURRENT_TIMESTAMP` clauses, the column has the current timestamp for its default value, and is automatically updated.
- With neither `DEFAULT` nor `ON UPDATE` clauses, it is the same as `DEFAULT CURRENT_TIMESTAMP ON UPDATE CURRENT_TIMESTAMP`.
- With a `DEFAULT CURRENT_TIMESTAMP` clause and no `ON UPDATE` clause, the column has the current timestamp for its default value but is not automatically updated.
- With no `DEFAULT` clause and with an `ON UPDATE CURRENT_TIMESTAMP` clause, the column has a default of 0 and is automatically updated.
- With a constant `DEFAULT` value, the column has the given default. If the column has an `ON UPDATE CURRENT_TIMESTAMP` clause, it is automatically updated, otherwise not.

In other words, you can use the current timestamp for both the initial value and the auto-update value, or either one, or neither. (For example, you can specify `ON UPDATE` to get auto-update without also having the column auto-initialized.)

`CURRENT_TIMESTAMP` or any of its synonyms (`CURRENT_TIMESTAMP()`, `NOW()`, `LOCALTIME`, `LOCALTIME()`, `LOCALTIMESTAMP`, or `LOCALTIMESTAMP()`) can be used in the `DEFAULT` and `ON UPDATE` clauses. They all mean “the current timestamp.” (`UTC_TIMESTAMP` is not allowed. Its range of values does not align with those of the `TIMESTAMP` column anyway unless the current time zone is `UTC`.)

The order of the `DEFAULT` and `ON UPDATE` attributes does not matter. If both `DEFAULT` and `ON UPDATE` are specified for a `TIMESTAMP` column, either can precede the other. For example, these statements are equivalent:

```
CREATE TABLE t (ts TIMESTAMP);
CREATE TABLE t (ts TIMESTAMP DEFAULT CURRENT_TIMESTAMP
ON UPDATE CURRENT_TIMESTAMP);
CREATE TABLE t (ts TIMESTAMP ON UPDATE CURRENT_TIMESTAMP
DEFAULT CURRENT_TIMESTAMP);
```

To specify automatic default or updating for a `TIMESTAMP` column other than the first one, you must suppress the automatic initialization and update behaviors for the first `TIMESTAMP` column by explicitly assigning it a constant `DEFAULT` value (for example, `DEFAULT 0` or `DEFAULT '2003-01-01 00:00:00'`). Then for the other `TIMESTAMP` column, the rules are the same as for the first `TIMESTAMP` column, except that if you omit both of the `DEFAULT` and `ON UPDATE` clauses, no automatic initialization or updating occurs.

Example. These statements are equivalent:

```
CREATE TABLE t (
  ts1 TIMESTAMP DEFAULT 0,
  ts2 TIMESTAMP DEFAULT CURRENT_TIMESTAMP
ON UPDATE CURRENT_TIMESTAMP);
CREATE TABLE t (
  ts1 TIMESTAMP DEFAULT 0,
  ts2 TIMESTAMP ON UPDATE CURRENT_TIMESTAMP
DEFAULT CURRENT_TIMESTAMP);
```

Beginning with MySQL 4.1.3, you can set the current time zone on a per-connection basis, as described in [Section 5.10.8, “MySQL Server Time Zone Support”](#). `TIMESTAMP` values still are stored in `UTC`, but are converted from the current time zone for storage, and converted back to the current time zone for retrieval. As long as the time zone setting remains constant, you get back the same value you store. If you store a `TIMESTAMP` value, and then change the time zone and retrieve the value, the retrieved value is different than the value you stored. This occurs because the same time zone was not used for conversion in both directions. The current time zone is available as the value of the `time_zone` system variable.

Beginning with MySQL 4.1.6, you can include the `NULL` attribute in the definition of a `TIMESTAMP` column to allow the column to contain `NULL` values. For example:

```
CREATE TABLE t (
  ts1 TIMESTAMP NULL DEFAULT NULL,
  ts2 TIMESTAMP NULL DEFAULT 0,
  ts3 TIMESTAMP NULL DEFAULT CURRENT_TIMESTAMP
);
```

Before MySQL 4.1.6 (and even as of 4.1.6 if the `NULL` attribute is not specified), setting the column to `NULL` sets it to the current timestamp. Note that a `TIMESTAMP` column which allows `NULL` values *not* take on the current timestamp except under one of the following conditions:

- Its default value is defined as `CURRENT_TIMESTAMP`
- `NOW()` or `CURRENT_TIMESTAMP` is inserted into the column

In other words, a `TIMESTAMP` column defined as `NULL` will auto-initialize only if it is created using a definition such as the following:

```
CREATE TABLE t (ts TIMESTAMP NULL DEFAULT CURRENT_TIMESTAMP);
```

Otherwise — that is, if the `TIMESTAMP` column is defined to allow `NULL` values but not using `DEFAULT TIMESTAMP`, as shown here...

```
CREATE TABLE t1 (ts TIMESTAMP NULL DEFAULT NULL);
CREATE TABLE t2 (ts TIMESTAMP NULL DEFAULT '0000-00-00 00:00:00');
```

...then you must explicitly insert a value corresponding to the current date and time, for example:

```
INSERT INTO t1 VALUES (NOW());
INSERT INTO t2 VALUES (CURRENT_TIMESTAMP);
```

## 11.3.2. The `TIME` Type

MySQL retrieves and displays `TIME` values in `'HH:MM:SS'` format (or `'HHH:MM:SS'` format for large hours values). `TIME` values may range from `'-838:59:59'` to `'838:59:59'`. The hours part may be so large because the `TIME` type can be used not only to represent a time of day (which must be less than 24 hours), but also elapsed time or a time interval between two events (which may be much greater than 24 hours, or even negative).

You can specify `TIME` values in a variety of formats:

- As a string in `'D HH:MM:SS.fraction'` format. You can also use one of the following “relaxed” syntaxes: `'HH:MM:SS.fraction'`, `'HH:MM:SS'`, `'HH:MM'`, `'D HH:MM:SS'`, `'D HH:MM'`, `'D HH'`, or `'SS'`. Here `D` represents days and can have a value from 0 to 34. Note that MySQL does not store the fraction part.
- As a string with no delimiters in `'HHMMSS'` format, provided that it makes sense as a time. For example, `'101112'` is understood as `'10:11:12'`, but `'109712'` is illegal (it has a nonsensical minute part) and becomes `'00:00:00'`.
- As a number in `HHMMSS` format, provided that it makes sense as a time. For example, `101112` is understood as `'10:11:12'`. The following alternative formats are also understood: `SS`, `MMSS`, `HHMMSS`, `HHMMSS.fraction`. Note that MySQL does not store the fractional part.
- As the result of a function that returns a value that is acceptable in a `TIME` context, such as `CURRENT_TIME`.

For `TIME` values specified as strings that include a time part delimiter, it is not necessary to specify two digits for hours, minutes, or seconds values that are less than 10. `'8:3:2'` is the same as `'08:03:02'`.

Be careful about assigning “short” `TIME` values to a `TIME` column. Without colons, MySQL interprets values using the assumption that the rightmost digits represent seconds. (MySQL interprets `TIME` values as elapsed time rather than as time of day.) For example, you might think of `'1112'` and `1112` as meaning `'11:12:00'` (12 minutes after 11 o'clock), but MySQL interprets them as `'00:11:12'` (11 minutes, 12 seconds). Similarly, `'12'` and `12` are interpreted as `'00:00:12'`. `TIME` values with

colons, by contrast, are always treated as time of the day. That is '11:12' means '11:12:00', not '00:11:12'.

By default, values that lie outside the `TIME` range but are otherwise legal are clipped to the closest endpoint of the range. For example, '-850:00:00' and '850:00:00' are converted to '-838:59:59' and '838:59:59'. Illegal `TIME` values are converted to '00:00:00'. Note that because '00:00:00' is itself a legal `TIME` value, there is no way to tell, from a value of '00:00:00' stored in a table, whether the original value was specified as '00:00:00' or whether it was illegal.

For more restrictive treatment of invalid `TIME` values, enable strict SQL mode to cause errors to occur. See [Section 5.2.5, “SQL Modes”](#).

### 11.3.3. The `YEAR` Type

The `YEAR` type is a one-byte type used for representing years.

MySQL retrieves and displays `YEAR` values in `YYYY` format. The range is 1901 to 2155.

You can specify `YEAR` values in a variety of formats:

- As a four-digit string in the range '1901' to '2155'.
- As a four-digit number in the range 1901 to 2155.
- As a two-digit string in the range '00' to '99'. Values in the ranges '00' to '69' and '70' to '99' are converted to `YEAR` values in the ranges 2000 to 2069 and 1970 to 1999.
- As a two-digit number in the range 1 to 99. Values in the ranges 1 to 69 and 70 to 99 are converted to `YEAR` values in the ranges 2001 to 2069 and 1970 to 1999. Note that the range for two-digit numbers is slightly different from the range for two-digit strings, because you cannot specify zero directly as a number and have it be interpreted as 2000. You must specify it as a string '0' or '00' or it is interpreted as 0000.
- As the result of a function that returns a value that is acceptable in a `YEAR` context, such as `NOW()`.

Illegal `YEAR` values are converted to 0000.

### 11.3.4. Y2K Issues and Date Types

As discussed in [Section 1.4.5, “Year 2000 Compliance”](#), MySQL itself is year 2000 (Y2K) safe. However, particular input values presented to MySQL may not be Y2K safe. Any value containing a two-digit year is ambiguous, because the century is unknown. Such values must be interpreted into four-digit form because MySQL stores years internally using four digits.

For `DATETIME`, `DATE`, `TIMESTAMP`, and `YEAR` types, MySQL interprets dates with ambiguous year values using the following rules:

- Year values in the range 00–69 are converted to 2000–2069.
- Year values in the range 70–99 are converted to 1970–1999.

Remember that these rules are only heuristics that provide reasonable guesses as to what your data values mean. If the rules used by MySQL do not produce the correct values, you should provide unambiguous input containing four-digit year values.

`ORDER BY` properly sorts `TIMESTAMP` or `YEAR` values that have two-digit years.

Some functions like `MIN()` and `MAX()` convert a `TIMESTAMP` or `YEAR` to a number. This means that a value with a two-digit year does not work properly with these functions. The fix in this case is to convert the `TIMESTAMP` or `YEAR` to four-digit year format or use something like `MIN(DATE_ADD(timestamp, INTERVAL 0 DAY))`.

## 11.4. String Types

The string types are `CHAR`, `VARCHAR`, `BINARY`, `VARBINARY`, `BLOB`, `TEXT`, `ENUM`, and `SET`. This section describes how these types work and how to use them in your queries. For string type storage requirements, see [Section 11.5, “Data Type Storage Requirements”](#).

### 11.4.1. The `CHAR` and `VARCHAR` Types

The `CHAR` and `VARCHAR` types are similar, but differ in the way they are stored and retrieved. No letter-case conversion takes place during storage or retrieval.

The `CHAR` and `VARCHAR` types are declared with a length that indicates the maximum number of characters you want to store. For example, `CHAR(30)` can hold up to 30 characters. (Before MySQL 4.1, the length is interpreted as number of bytes.)

The length of a `CHAR` column is fixed to the length that you declare when you create the table. The length can be any value from 0 to 255. (Before MySQL 3.23, the length of `CHAR` may be from 1 to 255.) When `CHAR` values are stored, they are right-padded with spaces to the specified length. When `CHAR` values are retrieved, trailing spaces are removed.

Values in `VARCHAR` columns are variable-length strings. The length can be specified as a value from 1 to 255 before MySQL 4.0.2 and 0 to 255 as of MySQL 4.0.2.

In contrast to `CHAR`, `VARCHAR` values are stored using only as many characters as are needed, plus one byte to record the length (two bytes for columns that are declared with a length longer than 255).

`VARCHAR` values are not padded when they are stored. Trailing spaces in MySQL version up to and including 4.1 are removed from values when stored in a `VARCHAR` column; this also means that the spaces are absent from retrieved values.

If you assign a value to a `CHAR` or `VARCHAR` column that exceeds the column's maximum length, the value is truncated to fit. If the truncated characters are not spaces, a warning is generated.

If you need a data type for which trailing spaces are not removed, consider using a `BLOB` or `TEXT` type. If you want to store binary values such as results from an encryption or compression function that might contain arbitrary byte values, use a `BLOB` column rather than a `CHAR` or `VARCHAR` column, to avoid potential problems with trailing space removal that would change data values.

The following table illustrates the differences between `CHAR` and `VARCHAR` by showing the result of storing various string values into `CHAR(4)` and `VARCHAR(4)` columns:

Value	<code>CHAR(4)</code>	Storage Required	<code>VARCHAR(4)</code>	Storage Required
' '	'   '	4 bytes	' '	1 byte
'ab'	'ab  '	4 bytes	'ab'	3 bytes
'abcd'	'abcd'	4 bytes	'abcd'	5 bytes
'abcdefgh'	'abcd'	4 bytes	'abcd'	5 bytes

If a given value is stored into the `CHAR(4)` and `VARCHAR(4)` columns, the values retrieved from the columns are not always the same because trailing spaces are removed from `CHAR` columns upon retrieval.

As of MySQL 4.1, values in `CHAR` and `VARCHAR` columns are sorted and compared according to the character set collation assigned to the column. Before MySQL 4.1, sorting and comparison are based on the collation of the server character set; you can declare the column with the `BINARY` attribute to cause sorting and comparison to be based on the numeric values of the bytes in column values. `BINARY` does not affect how column values are stored or retrieved.

Note that all MySQL collations are of type `PADSPACE`. This means that all `CHAR` and `VARCHAR` values in MySQL are compared without regard to any trailing spaces. For example:

```
mysql> CREATE TABLE names (myname CHAR(10), yourname VARCHAR(10));
Query OK, 0 rows affected (0.09 sec)

mysql> INSERT INTO names VALUES ('Monty ', 'Monty ');
Query OK, 1 row affected (0.00 sec)

mysql> SELECT myname = 'Monty ', yourname = 'Monty ' FROM names;
+-----+-----+
| myname = 'Monty ' | yourname = 'Monty ' |
+-----+-----+
| 1 | 1 |
+-----+-----+
1 row in set (0.00 sec)
```

Note that this is true for all MySQL versions, and it is not affected by the trimming of trailing spaces from `VARCHAR` values before storing them. Nor does the server SQL mode make any difference in this regard.

For those cases where trailing pad characters are stripped or comparisons ignore them, if a column has an index that requires unique values, inserting into the column values that differ only in number of trailing pad characters will result in a duplicate-key error. For example, if a table contains 'a', an attempt to store 'a ' causes a duplicate-key error.

The `BINARY` attribute is sticky. This means that if a column marked `BINARY` is used in an expression, the whole expression is treated as a `BINARY` value.

MySQL may silently change the type of a `CHAR` or `VARCHAR` column at table creation time. See [Section 13.1.5.1, “Silent Column Specification Changes”](#).

## 11.4.2. The `BINARY` and `VARBINARY` Types

The `BINARY` and `VARBINARY` types are similar to `CHAR` and `VARCHAR`, except that they contain binary strings rather than non-binary strings. That is, they contain byte strings rather than character strings. This means that they have no character set, and sorting and comparison are based on the numeric values of the bytes in the values.

The allowable maximum length is the same for `BINARY` and `VARBINARY` as it is for `CHAR` and `VARCHAR`, except that the length for `BINARY` and `VARBINARY` is a length in bytes rather than in characters.

Before MySQL 4.1.2, `BINARY(M)` and `VARBINARY(M)` are treated as `CHAR(M) BINARY` and `VARCHAR(M) BINARY`. As of MySQL 4.1.2, the `BINARY` and `VARBINARY` data types are distinct from the `CHAR BINARY` and `VARCHAR BINARY` data types. For the latter types, the `BINARY` attribute does not cause the column to be treated as a binary string column. Instead, it causes the binary collation for the column character set to be used, and the column itself contains non-binary character strings rather than binary byte strings. For example, in 4.1.2 and up, `CHAR(5) BINARY` is treated as `CHAR(5) CHARACTER SET latin1 COLLATE latin1_bin`, assuming that the default character set is `latin1`. This differs from `BINARY(5)`, which stores 5-bytes binary strings that have no

character set or collation.

The handling of trailing spaces is the same for `BINARY` and `VARBINARY` as it is for `CHAR` and `VARCHAR`. When `BINARY` values are stored, they are right-padded with spaces to the specified length. When `BINARY` values are retrieved, trailing spaces are removed. For `VARBINARY`, trailing spaces are removed when values are stored.

For those cases where trailing pad bytes are stripped or comparisons ignore them, if a column has an index that requires unique values, inserting into the column values that differ only in number of trailing pad bytes will result in a duplicate-key error. For example, if a table contains 'a', an attempt to store 'a ' causes a duplicate-key error. Trailing spaces are significant in comparisons.

You should consider the preceding padding and stripping characteristics carefully if you plan to use one of these data types for storing binary data and you require that the value retrieved be exactly the same as the value stored. The following example illustrates how space-padding of `BINARY` values affects column value comparisons:

```
mysql> CREATE TABLE t (c BINARY(3));
Query OK, 0 rows affected (0.00 sec)

mysql> INSERT INTO t SET c = 'a ';
Query OK, 1 row affected (0.00 sec)

mysql> SELECT HEX(c), c = 'a', c = 'a ' from t;
+-----+-----+-----+
| HEX(c) | c = 'a' | c = 'a ' |
+-----+-----+-----+
| 61     | 1       | 0        |
+-----+-----+-----+
1 row in set (0.00 sec)
```

If the value retrieved must be the same as the value specified for storage with no padding, it might be preferable to use one of the `BLOB` data types instead.

### 11.4.3. The `BLOB` and `TEXT` Types

A `BLOB` is a binary large object that can hold a variable amount of data. The four `BLOB` types are `TINYBLOB`, `BLOB`, `MEDIUMBLOB`, and `LONGBLOB`. These differ only in the maximum length of the values they can hold. The four `TEXT` types are `TINYTEXT`, `TEXT`, `MEDIUMTEXT`, and `LONGTEXT`. These correspond to the four `BLOB` types and have the same maximum lengths and storage requirements. See [Section 11.5, “Data Type Storage Requirements”](#). No lettercase conversion for `TEXT` or `BLOB` columns takes place during storage or retrieval.

`BLOB` columns are treated as binary strings (byte strings). `TEXT` columns are treated as non-binary strings (character strings). `BLOB` columns have no character set, and sorting and comparison are based on the numeric values of the bytes in column values. `TEXT` columns have a character set, and values are sorted and compared based on the collation of the character set assigned to the column as of MySQL 4.1. Before 4.1, `TEXT` sorting and comparison are based on the collation of the server character set.

If a `TEXT` column is indexed, index entry comparisons are space-padded at the end. This means that, if the index requires unique values, duplicate-key errors will occur for values that differ only in the number of trailing spaces. For example, if a table contains 'a', an attempt to store 'a ' causes a duplicate-key error. This is not true for `BLOB` columns.

If you assign a value to a `BLOB` or `TEXT` column that exceeds the data type's maximum length, the value is truncated to fit. If the truncated characters are not spaces, a warning is generated. See [Section 5.2.5, “SQL Modes”](#).

In most respects, you can regard a `BLOB` column as a `VARBINARY` column that can be as big as you like. Similarly, you can regard a `TEXT` column as a `VARCHAR` column. `BLOB` and `TEXT` differ from `VARBINARY` and `VARCHAR` in the following ways:

- There is no trailing-space removal for `BLOB` and `TEXT` columns when values are stored or retrieved. This differs from `VARBINARY` and `VARCHAR`, for which trailing spaces are removed when values are stored.

Note that `TEXT` is on comparison space extended to fit the compared object, exactly like `CHAR` and `VARCHAR`.

- You can have indexes on `BLOB` and `TEXT` columns only as of MySQL 3.23.2 for `MyISAM` tables or MySQL 4.0.14 for `InnoDB` tables. Previous versions of MySQL did not support indexing these data types.
- For indexes on `BLOB` and `TEXT` columns, you must specify an index prefix length. For `CHAR` and `VARCHAR`, a prefix length is optional. See [Section 7.4.3, “Column Indexes”](#).
- `BLOB` and `TEXT` columns cannot have `DEFAULT` values.

From MySQL 4.1.0 on, `LONG` and `LONG VARCHAR` map to the `MEDIUMTEXT` data type. This is a compatibility feature. If you use the `BINARY` attribute with a `TEXT` data type, the column is assigned the binary collation of the column character set.

MySQL Connector/ODBC defines `BLOB` values as `LONGVARBINARY` and `TEXT` values as `LONGVARCHAR`.

Because `BLOB` and `TEXT` values can be extremely long, you might encounter some constraints in using them:

- Only the first `max_sort_length` bytes of the column are used when sorting. The default value of `max_sort_length` is 1024. This value can be changed using the `--max_sort_length=N` option when starting the `mysqld` server. See [Section 5.2.2, “System Variables”](#).

As of MySQL 4.0.3, you can make more bytes significant in sorting or grouping by increasing the value of `max_sort_length` at runtime. Any client can change the value of its session `max_sort_length` variable:

```
mysql> SET max_sort_length = 2000;
mysql> SELECT id, comment FROM t
-> ORDER BY comment;
```

Another way to use `GROUP BY` or `ORDER BY` on a `BLOB` or `TEXT` column containing long values when you want more than `max_sort_length` bytes to be significant is to convert the column value into a fixed-length object. The standard way to do this is with the `SUBSTRING()` function. For example, the following statement causes 2000 bytes of the `comment` column to be taken into account for sorting:

```
mysql> SELECT id, SUBSTRING(comment,1,2000) FROM t
-> ORDER BY SUBSTRING(comment,1,2000);
```

Before MySQL 3.23.2, you can group on an expression involving `BLOB` or `TEXT` values by using a column alias or by specifying the column position:

```
mysql> SELECT id, SUBSTRING(comment,1,2000) AS b
-> FROM tbl_name GROUP BY b;
mysql> SELECT id, SUBSTRING(comment,1,2000)
-> FROM tbl_name GROUP BY 2;
```

- The maximum size of a `BLOB` or `TEXT` object is determined by its type, but the largest value you actually can transmit between the client and server is determined by the amount of available memory



and the size of the communications buffers. You can change the message buffer size by changing the value of the `max_allowed_packet` variable, but you must do so for both the server and your client program. For example, both `mysql` and `mysqldump` allow you to change the client-side `max_allowed_packet` value. See [Section 7.5.2, “Tuning Server Parameters”](#), [Section 8.6, “mysql — The MySQL Command-Line Tool”](#), and [Section 8.12, “mysqldump — A Database Backup Program”](#). You may also want to compare the packet sizes and the size of the data objects you are storing with the storage requirements, see [Section 11.5, “Data Type Storage Requirements”](#)

Each `BLOB` or `TEXT` value is represented internally by a separately allocated object. This is in contrast to all other data types, for which storage is allocated once per column when the table is opened.

In some cases, it may be desirable to store binary data such as media files in `BLOB` or `TEXT` columns. You may find MySQL's string handling functions useful for working with such data. See [Section 12.3, “String Functions”](#). For security and other reasons, it is usually preferable to do so using application code rather than allowing application users the `FILE` privilege. You can discuss specifics for various languages and platforms in the MySQL Forums (<http://forums.mysql.com/>).

### 11.4.4. The `ENUM` Type

An `ENUM` is a string object with a value chosen from a list of allowed values that are enumerated explicitly in the column specification at table creation time.

An enumeration value must be a quoted string literal; it may not be an expression, even one that evaluates to a string value. This means that you also may not employ a user variable as an enumeration value.

The value may also be the empty string ( `' '` ) or `NULL` under certain circumstances:

- If you insert an invalid value into an `ENUM` (that is, a string not present in the list of allowed values), the empty string is inserted instead as a special error value. This string can be distinguished from a “normal” empty string by the fact that this string has the numerical value 0. More about this later.

If strict SQL mode is enabled, attempts to insert invalid `ENUM` values result in an error.

- If an `ENUM` column is declared to allow `NULL`, the `NULL` value is a legal value for the column, and the default value is `NULL`. If an `ENUM` column is declared `NOT NULL`, its default value is the first element of the list of allowed values.

Each enumeration value has an index:

- Values from the list of allowable elements in the column specification are numbered beginning with 1.
- The index value of the empty string error value is 0. This means that you can use the following `SELECT` statement to find rows into which invalid `ENUM` values were assigned:

```
mysql> SELECT * FROM tbl_name WHERE enum_col=0;
```

- The index of the `NULL` value is `NULL`.
- The term “index” here refers only to position within the list of enumeration values. It has nothing to do with table indexes.

For example, a column specified as `ENUM('one', 'two', 'three')` can have any of the values

shown here. The index of each value is also shown:

Value	Index
NULL	NULL
' '	0
'one'	1
'two'	2
'three'	3

An enumeration can have a maximum of 65,535 elements.

Starting from MySQL 3.23.51, trailing spaces are automatically deleted from `ENUM` member values in the table definition when a table is created.

When retrieved, values stored into an `ENUM` column are displayed using the lettercase that was used in the column definition. Before MySQL 4.1.1, lettercase is irrelevant when you assign values to an `ENUM` column. As of 4.1.1, `ENUM` columns can be assigned a character set and collation. For binary or case-sensitive collations, lettercase does matter when you assign values to the column.

If you retrieve an `ENUM` value in a numeric context, the column value's index is returned. For example, you can retrieve numeric values from an `ENUM` column like this:

```
mysql> SELECT enum_col+0 FROM tbl_name;
```

If you store a number into an `ENUM` column, the number is treated as the index into the possible values, and the value stored is the enumeration member with that index. (However, this does *not* work with `LOAD DATA`, which treats all input as strings.) If the numeric value is quoted, it is still interpreted as an index if there is no matching string in the list of enumeration values. For these reasons, it is not advisable to define an `ENUM` column with enumeration values that look like numbers, because this can easily become confusing. For example, the following column has enumeration members with string values of '0', '1', and '2', but numeric index values of 1, 2, and 3:

```
numbers ENUM('0','1','2')
```

If you store 2, it is interpreted as an index value, and becomes '1' (the value with index 2). If you store '2', it matches an enumeration value, so it is stored as '2'. If you store '3', it does not match any enumeration value, so it is treated as an index and becomes '2' (the value with index 3).

```
mysql> INSERT INTO t (numbers) VALUES(2),('2'),('3');
mysql> SELECT * FROM t;
```

```
+-----+
| numbers |
+-----+
| 1       |
| 2       |
| 2       |
+-----+
```

`ENUM` values are sorted according to the order in which the enumeration members were listed in the column specification. (In other words, `ENUM` values are sorted according to their index numbers.) For example, 'a' sorts before 'b' for `ENUM('a', 'b')`, but 'b' sorts before 'a' for `ENUM('b', 'a')`. The empty string sorts before non-empty strings, and `NULL` values sort before all other enumeration values. If you expect sorting to be done alphabetically, you should specify the `ENUM` list in alphabetical order. You can also use `GROUP BY CAST(col AS CHAR)` or `GROUP BY CONCAT(col)` to make sure that the column is sorted lexically rather than by index number.

If you want to determine all possible values for an `ENUM` column, use `SHOW COLUMNS FROM`

`tbl_name LIKE enum_col` and parse the `ENUM` definition in the `Type` column of the output.

## 11.4.5. The `SET` Type

A `SET` is a string object that can have zero or more values, each of which must be chosen from a list of allowed values specified when the table is created. `SET` column values that consist of multiple set members are specified with members separated by commas (`,`). A consequence of this is that `SET` member values should not themselves contain commas.

For example, a column specified as `SET('one', 'two') NOT NULL` can have any of these values:

```
' '
'one'
'two'
'one,two'
```

A `SET` can have a maximum of 64 different members.

Starting from MySQL 3.23.51, trailing spaces are automatically deleted from `SET` member values in the table definition when a table is created.

When retrieved, values stored into a `SET` column are displayed using the lettercase that was used in the column definition. Before MySQL 4.1.1, lettercase is irrelevant when you assign values to an `SET` column. As of 4.1.1, `SET` columns can be assigned a character set and collation. For binary or case-sensitive collations, lettercase does matter when you assign values to the column.

MySQL stores `SET` values numerically, with the low-order bit of the stored value corresponding to the first set member. If you retrieve a `SET` value in a numeric context, the value retrieved has bits set corresponding to the set members that make up the column value. For example, you can retrieve numeric values from a `SET` column like this:

```
mysql> SELECT set_col+0 FROM tbl_name;
```

If a number is stored into a `SET` column, the bits that are set in the binary representation of the number determine the set members in the column value. For a column specified as `SET('a', 'b', 'c', 'd')`, the members have the following decimal and binary values:

SET Member	Decimal Value	Binary Value
'a'	1	0001
'b'	2	0010
'c'	4	0100
'd'	8	1000

If you assign a value of 9 to this column, that is 1001 in binary, the first and fourth `SET` members 'a' and 'd' are selected and the resulting value is 'a,d'.

For a value containing more than one `SET` element, it does not matter what order the elements are listed in when you insert the value. It also does not matter how many times a given element is listed in the value. When the value is retrieved later, each element in the value appears once, with elements listed according to the order in which they were specified at table creation time. For example, suppose that a column is specified as `SET('a', 'b', 'c', 'd')`:

```
mysql> CREATE TABLE myset (col SET('a', 'b', 'c', 'd'));
```

If you insert the values 'a,d', 'd,a', 'a,d,d', 'a,d,a', and 'd,a,d':

```
mysql> INSERT INTO myset (col) VALUES
-> ('a,d'), ('d,a'), ('a,d,a'), ('a,d,d'), ('d,a,d');
Query OK, 5 rows affected (0.01 sec)
Records: 5 Duplicates: 0 Warnings: 0
```

Then all of these values appear as 'a,d' when retrieved:

```
mysql> SELECT col FROM myset;
+-----+
| col   |
+-----+
| a,d   |
| a,d   |
| a,d   |
| a,d   |
| a,d   |
+-----+
5 rows in set (0.04 sec)
```

If you set a `SET` column to an unsupported value, the value is ignored and a warning is issued:

```
mysql> INSERT INTO myset (col) VALUES ('a,d,d,s');
Query OK, 1 row affected, 1 warning (0.03 sec)

mysql> SHOW WARNINGS;
+-----+-----+-----+
| Level | Code | Message                                     |
+-----+-----+-----+
| Warning | 1265 | Data truncated for column 'col' at row 1 |
+-----+-----+-----+
1 row in set (0.04 sec)

mysql> SELECT col FROM myset;
+-----+
| col   |
+-----+
| a,d   |
| a,d   |
| a,d   |
| a,d   |
| a,d   |
| a,d   |
+-----+
6 rows in set (0.01 sec)
```

If strict SQL mode is enabled, attempts to insert invalid `SET` values result in an error.

`SET` values are sorted numerically. `NULL` values sort before non-`NULL` `SET` values.

Normally, you search for `SET` values using the `FIND_IN_SET()` function or the `LIKE` operator:

```
mysql> SELECT * FROM tbl_name WHERE FIND_IN_SET('value',set_col)>0;
mysql> SELECT * FROM tbl_name WHERE set_col LIKE '%value%';
```

The first statement finds rows where `set_col` contains the `value` set member. The second is similar, but not the same: It finds rows where `set_col` contains `value` anywhere, even as a substring of another set member.

The following statements also are legal:

```
mysql> SELECT * FROM tbl_name WHERE set_col & 1;
mysql> SELECT * FROM tbl_name WHERE set_col = 'val1,val2';
```

The first of these statements looks for values containing the first set member. The second looks for an exact match. Be careful with comparisons of the second type. Comparing set values to `'val1,val2'` returns different results than comparing values to `'val2,val1'`. You should specify the values in the same order in which they are listed in the column definition.

If you want to determine all possible values for a `SET` column, use `SHOW COLUMNS FROM tbl_name LIKE set_col` and parse the `SET` definition in the `Type` column of the output.

## 11.5. Data Type Storage Requirements

The storage requirements for each of the data types supported by MySQL are listed here by category.

The maximum size of a row in a `MyISAM` table is 65,534 bytes. Each `BLOB` and `TEXT` column accounts for only five to nine bytes toward this size.

**Important:** For tables using the `NDBCluster` storage engine, there is the factor of *4-byte alignment* to be taken into account when calculating storage requirements. This means that all `NDB` data storage is done in multiples of 4 bytes. Thus, a column value that — in a table using a storage engine other than `NDB` — would take 15 bytes for storage, requires 16 bytes in an `NDB` table. This requirement applies in addition to any other considerations that are discussed in this section. For example, in `NDBCluster` tables, the `TINYINT`, `SMALLINT`, `MEDIUMINT`, and `INTEGER (INT)` column types each require 4 bytes storage per record.

In addition, when calculating storage requirements for Cluster tables, you must remember that every table using the `NDBCluster` storage engine requires a primary key; if no primary key is defined by the user, then a “hidden” primary key will be created by `NDB`. This hidden primary key consumes 31-35 bytes per table record.

When calculating Cluster memory requirements, you may find useful the `ndb_size.pl` utility which is available on [MySQLForge](http://forge.mysql.com/) [http://forge.mysql.com/]. This Perl script connects to a current MySQL (non-Cluster) database and creates a report on how much space that database would require if it used the `NDBCluster` storage engine.

### Storage Requirements for Numeric Types

Data Type	Storage Required
<code>TINYINT</code>	1 byte
<code>SMALLINT</code>	2 bytes
<code>MEDIUMINT</code>	3 bytes
<code>INT</code> , <code>INTEGER</code>	4 bytes
<code>BIGINT</code>	8 bytes
<code>FLOAT(p)</code>	4 bytes if $0 \leq p \leq 24$ , 8 bytes if $25 \leq p \leq 53$
<code>FLOAT</code>	4 bytes
<code>DOUBLE [PRECISION]</code> , <code>REAL</code>	8 bytes
<code>DECIMAL(M,D)</code> , <code>NUMERIC(M,D)</code>	Varies; see following discussion

In MySQL versions up to and including 4.1, `DECIMAL` columns are represented as strings and their storage requirements are:

- $M+2$  bytes, if  $D > 0$
- $M+1$  bytes, if  $D = 0$
- $D+2$ , if  $M < D$

### Storage Requirements for Date and Time Types

Data Type	Storage Required
DATE	3 bytes
DATETIME	8 bytes
TIMESTAMP	4 bytes
TIME	3 bytes
YEAR	1 byte

### Storage Requirements for String Types

Data Type	Storage Required
CHAR( <i>M</i> )	<i>M</i> bytes, $0 \leq M \leq 255$
VARCHAR( <i>M</i> )	<i>L</i> +1 bytes, where $L \leq M$ and $0 \leq M \leq 255$
BINARY( <i>M</i> )	<i>M</i> bytes, $0 \leq M \leq 255$
VARBINARY( <i>M</i> )	<i>L</i> +1 bytes, where $L \leq M$ and $0 \leq M \leq 255$
TINYBLOB, TINYTEXT	<i>L</i> +1 bytes, where $L < 2^8$
BLOB, TEXT	<i>L</i> +2 bytes, where $L < 2^{16}$
MEDIUMBLOB, MEDIUMTEXT	<i>L</i> +3 bytes, where $L < 2^{24}$
LONGBLOB, LONGTEXT	<i>L</i> +4 bytes, where $L < 2^{32}$
ENUM('value1', 'value2', ...)	1 or 2 bytes, depending on the number of enumeration values (65,535 values maximum)
SET('value1', 'value2', ...)	1, 2, 3, 4, or 8 bytes, depending on the number of set members (64 members maximum)

For the CHAR, VARCHAR, and TEXT types, *L* and *M* in the preceding table should be interpreted as number of bytes before MySQL 4.1 and as number of characters thereafter. Lengths for these types in column specifications indicate number of characters from MySQL 4.1 on. The number of extra bytes for recording lengths for variable-length data types is unchanged. For example, *L*+1 bytes to store a TINYTEXT value before MySQL 4.1 becomes *L* characters + 1 byte to store the length as of MySQL 4.1.

VARCHAR and the BLOB and TEXT types are variable-length types. For each, the storage requirements depend on the actual length of column values (represented by *L* in the preceding table), rather than on the type's maximum possible size. For example, a VARCHAR(10) column can hold a string with a maximum length of 10. The actual storage required is the length of the string (*L*), plus one byte to record the length of the string. For the string 'abcd', *L* is 4 and the storage requirement is five bytes.

As of MySQL 4.1, to calculate the number of bytes used to store a particular CHAR, VARCHAR, or TEXT column value, you must take into account the character set used for that column. In particular, when using the utf8 Unicode character set, you must keep in mind that not all utf8 characters use the same number of bytes. For a breakdown of the storage used for different categories of utf8 characters, see [Section 10.7, “Unicode Support”](#).

**Note:** The NDB Cluster engine supports only fixed-width columns. This means that a VARCHAR column from a table in a MySQL Cluster will behave almost as if it were of type CHAR (except that each record still has one extra byte overhead). For example, in an NDB table, each record in a column declared as VARCHAR(100) will take up 101 bytes for storage, regardless of the length of the string actually stored in any given record.

The BLOB and TEXT types require 1, 2, 3, or 4 bytes to record the length of the column value, depending on the maximum possible length of the type. See [Section 11.4.3, “The BLOB and TEXT Types”](#).

`TEXT` and `BLOB` columns are implemented differently in the `NDB Cluster` storage engine, wherein each record in a `TEXT` column is made up of two separate parts. One of these is of fixed size (256 bytes), and is actually stored in the original table. The other consists of any data in excess of 256 bytes, which is stored in a hidden table. The records in this second table are always 2,000 bytes long. This means that the size of a `TEXT` column is 256 if `size`  $\leq$  256 (where `size` represents the size of the record); otherwise, the size is  $256 + size + (2000 - (size - 256) \% 2000)$ .

The size of an `ENUM` object is determined by the number of different enumeration values. One byte is used for enumerations with up to 255 possible values. Two bytes are used for enumerations having between 256 and 65,535 possible values. See [Section 11.4.4, “The `ENUM` Type”](#).

The size of a `SET` object is determined by the number of different set members. If the set size is `N`, the object occupies  $(N + 7) / 8$  bytes, rounded up to 1, 2, 3, 4, or 8 bytes. A `SET` can have a maximum of 64 members. See [Section 11.4.5, “The `SET` Type”](#).

## 11.6. Choosing the Right Type for a Column

For the most efficient use of storage, try to use the most precise type in all cases. For example, if an integer column is used for values in the range from 1 to 99999, `MEDIUMINT UNSIGNED` is the best type. Of the types that represent all the required values, it uses the least amount of storage.

For earlier MySQL versions, accurate representation of monetary values was a common problem. In these MySQL versions, you should also use the `DECIMAL` type. In this case the value is stored as a string, so no loss of accuracy should occur on storage. Calculations on these `DECIMAL` values are however done using double-precision operations. If accuracy is not too important or if speed is important, the `DOUBLE` type may also be good enough.

For high precision, you can always convert to a fixed-point type stored in a `BIGINT`. This allows you to do all calculations with 64-bit integers and then convert results back to floating-point values only when necessary.

## 11.7. Using Data Types from Other Database Engines

To make it easier to use code written for SQL implementations from other vendors, MySQL maps data types as shown in the following table. These mappings make it easier to import table definitions from other database systems into MySQL:

Other Vendor Type	MySQL Type
<code>BINARY (M)</code>	<code>CHAR (M)</code> <code>BINARY</code> (before MySQL 4.1.2)
<code>BOOL</code>	<code>TINYINT</code>
<code>BOOLEAN</code>	<code>TINYINT</code>
<code>CHAR VARYING (M)</code>	<code>VARCHAR (M)</code>
<code>DEC</code>	<code>DECIMAL</code>
<code>FIXED</code>	<code>DECIMAL</code> (MySQL 4.1.0 on)
<code>FLOAT4</code>	<code>FLOAT</code>
<code>FLOAT8</code>	<code>DOUBLE</code>
<code>INT1</code>	<code>TINYINT</code>
<code>INT2</code>	<code>SMALLINT</code>
<code>INT3</code>	<code>MEDIUMINT</code>
<code>INT4</code>	<code>INT</code>
<code>INT8</code>	<code>BIGINT</code>

LONG VARBINARY	MEDIUMBLOB
LONG VARCHAR	MEDIUMTEXT
LONG	MEDIUMTEXT (MySQL 4.1.0 on)
MIDDLEINT	MEDIUMINT
NUMERIC	DECIMAL
VARBINARY( <i>M</i> )	VARCHAR( <i>M</i> ) BINARY (before MySQL 4.1.2)

As of MySQL 4.1.2, `BINARY` and `VARBINARY` are distinct data types and are not converted to `CHAR BINARY` and `VARCHAR BINARY`.

Data type mapping occurs at table creation time, after which the original type specifications are discarded. If you create a table with types used by other vendors and then issue a `DESCRIBE tbl_name` statement, MySQL reports the table structure using the equivalent MySQL types. For example:

```
mysql> CREATE TABLE t (a BOOL, b FLOAT8, c LONG VARCHAR, d NUMERIC);
Query OK, 0 rows affected (0.00 sec)

mysql> DESCRIBE t;
+-----+-----+-----+-----+-----+-----+
| Field | Type          | Null | Key | Default | Extra |
+-----+-----+-----+-----+-----+-----+
| a     | tinyint(1)   | YES  |     | NULL    |       |
| b     | double       | YES  |     | NULL    |       |
| c     | mediumtext   | YES  |     | NULL    |       |
| d     | decimal(10,0)| YES  |     | NULL    |       |
+-----+-----+-----+-----+-----+-----+
4 rows in set (0.01 sec)
```



---

# Chapter 12. Functions and Operators

Expressions can be used at several points in SQL statements, such as in the `ORDER BY` or `HAVING` clauses of `SELECT` statements, in the `WHERE` clause of a `SELECT`, `DELETE`, or `UPDATE` statement, or in `SET` statements. Expressions can be written using literal values, column values, `NULL`, built-in functions, user-defined functions, and operators. This chapter describes the functions and operators that are allowed for writing expressions in MySQL. Instructions for writing user-defined functions are given in [Section 19.2, “Adding New Functions to MySQL”](#).

An expression that contains `NULL` always produces a `NULL` value unless otherwise indicated in the documentation for a particular function or operator.

**Note:** By default, there must be no whitespace between a function name and the parenthesis following it. This helps the MySQL parser distinguish between function calls and references to tables or columns that happen to have the same name as a function. However, spaces around function arguments are permitted.

You can tell the MySQL server to accept spaces after function names by starting it with the `--sql-mode=IGNORE_SPACE` option. (See [Section 5.2.5, “SQL Modes”](#).) Individual client programs can request this behavior by using the `CLIENT_IGNORE_SPACE` option for `mysql_real_connect()`. In either case, all function names become reserved words.

For the sake of brevity, most examples in this chapter display the output from the `mysql` program in abbreviated form. Rather than showing examples in this format:

```
mysql> SELECT MOD(29,9);
+-----+
| mod(29,9) |
+-----+
|          2 |
+-----+
1 rows in set (0.00 sec)
```

This format is used instead:

```
mysql> SELECT MOD(29,9);
-> 2
```

## 12.1. Operators

### 12.1.1. Operator Precedence

Operator precedences are shown in the following list, from lowest precedence to the highest. Operators that are shown together on a line have the same precedence.

```
:=
||, OR, XOR
&&, AND
BETWEEN, CASE, WHEN, THEN, ELSE
=, <=>, >=, >, <=, <, <>, !=, IS, LIKE, REGEXP, IN
|
&
<<, >>
-, +
*, /, DIV, %, MOD
^
- (unary minus), ~ (unary bit inversion)
!, NOT
BINARY, COLLATE
```

The precedence of operators determines the order of evaluation of terms in an expression. To override

this order and group terms explicitly, use parentheses. For example:

```
mysql> SELECT 1+2*3;
-> 7
mysql> SELECT (1+2)*3;
-> 9
```

## 12.1.2. Type Conversion in Expression Evaluation

When an operator is used with operands of different types, type conversion occurs to make the operands compatible. Some conversions occur implicitly. For example, MySQL automatically converts numbers to strings as necessary, and vice versa.

```
mysql> SELECT 1+'1';
-> 2
mysql> SELECT CONCAT(2,' test');
-> '2 test'
```

It is also possible to perform explicit conversions. If you want to convert a number to a string explicitly, use the `CAST()` or `CONCAT()` function (`CAST()` is preferable, but is unavailable before MySQL 4.0.2):

```
mysql> SELECT 38.8, CAST(38.8 AS CHAR);
-> 38.8, '38.8'
mysql> SELECT 38.8, CONCAT(38.8);
-> 38.8, '38.8'
```

The following rules describe how conversion occurs for comparison operations:

- If one or both arguments are `NULL`, the result of the comparison is `NULL`, except for the `NULL`-safe `<=>` equality comparison operator. For `NULL <=> NULL`, the result is true.
- If both arguments in a comparison operation are strings, they are compared as strings.
- If both arguments are integers, they are compared as integers.
- Hexadecimal values are treated as binary strings if not compared to a number.
- If one of the arguments is a `TIMESTAMP` or `DATETIME` column and the other argument is a constant, the constant is converted to a timestamp before the comparison is performed. This is done to be more ODBC-friendly. Note that this is not done for the arguments to `IN()`! To be safe, always use complete datetime, date, or time strings when doing comparisons.
- In all other cases, the arguments are compared as floating-point (real) numbers.

The following examples illustrate conversion of strings to numbers for comparison operations:

```
mysql> SELECT 1 > '6x';
-> 0
mysql> SELECT 7 > '6x';
-> 1
mysql> SELECT 0 > 'x6';
-> 0
mysql> SELECT 0 = 'x6';
-> 1
```

Note that when you are comparing a string column with a number, MySQL cannot use an index on the column to quickly look up the value. If `str_col` is an indexed string column, the index cannot be used when performing the lookup in the following statement:

```
SELECT * FROM tbl_name WHERE str_col=1;
```

The reason for this is that there are many different strings that may convert to the value 1, such as '1', ' 1', or '1a'.

Comparisons that use floating-point numbers (or values that are converted to floating-point numbers) are approximate because such numbers are inexact. This might lead to results that appear inconsistent:

```
mysql> SELECT '18015376320243458' = 18015376320243458;
-> 1
mysql> SELECT '18015376320243459' = 18015376320243459;
-> 0
```

Such results can occur because the values are converted to floating-point numbers, which have only 53 bits of precision and are subject to rounding:

```
mysql> SELECT '18015376320243459'+0.0;
-> 1.8015376320243e+16
```

Furthermore, the conversion from string to floating-point and from integer to floating-point do not necessarily occur the same way. The integer may be converted to floating-point by the CPU, whereas the string is converted digit by digit in an operation that involves floating-point multiplications.

The results shown will vary on different systems, and can be affected by factors such as computer architecture or the compiler version or optimization level. One way to avoid such problems is to use `CAST()` so that a value will not be converted implicitly to a float-point number:

```
mysql> SELECT CAST('18015376320243459' AS UNSIGNED) = 18015376320243459;
-> 1
```

For more information about floating-point comparisons, see [Section A.5.8, “Problems with Floating-Point Comparisons”](#).

## 12.1.3. Comparison Functions and Operators

Comparison operations result in a value of 1 (`TRUE`), 0 (`FALSE`), or `NULL`. These operations work for both numbers and strings. Strings are automatically converted to numbers and numbers to strings as necessary.

Some of the functions in this section (such as `LEAST()` and `GREATEST()`) return values other than 1 (`TRUE`), 0 (`FALSE`), or `NULL`. However, the value they return is based on comparison operations performed according to the rules described in [Section 12.1.2, “Type Conversion in Expression Evaluation”](#).

To convert a value to a specific type for comparison purposes, you can use the `CAST()` function. String values can be converted to a different character set using `CONVERT()`. See [Section 12.8, “Cast Functions and Operators”](#).

By default, string comparisons are not case sensitive and use the current character set. The default is `latin1` (cp1252 West European), which also works well for English.

- =

Equal:

```
mysql> SELECT 1 = 0;
-> 0
mysql> SELECT '0' = 0;
-> 1
mysql> SELECT '0.0' = 0;
```

```
mysql> -> 1
mysql> SELECT '0.01' = 0;
        -> 0
mysql> SELECT '.01' = 0.01;
        -> 1
```

- `<=>`

`NULL`-safe equal. This operator performs an equality comparison like the `=` operator, but returns `1` rather than `NULL` if both operands are `NULL`, and `0` rather than `NULL` if one operand is `NULL`.

```
mysql> SELECT 1 <=> 1, NULL <=> NULL, 1 <=> NULL;
        -> 1, 1, 0
mysql> SELECT 1 = 1, NULL = NULL, 1 = NULL;
        -> 1, NULL, NULL
```

`<=>` was added in MySQL 3.23.0.

- `<>`, `!=`

Not equal:

```
mysql> SELECT '.01' <> '0.01';
        -> 1
mysql> SELECT .01 <> '0.01';
        -> 0
mysql> SELECT 'zapp' <> 'zappp';
        -> 1
```

- `<=`

Less than or equal:

```
mysql> SELECT 0.1 <= 2;
        -> 1
```

- `<`

Less than:

```
mysql> SELECT 2 < 2;
        -> 0
```

- `>=`

Greater than or equal:

```
mysql> SELECT 2 >= 2;
        -> 1
```

- `>`

Greater than:

```
mysql> SELECT 2 > 2;
        -> 0
```

- `IS NULL`, `IS NOT NULL`

Tests whether a value is or is not `NULL`.

```
mysql> SELECT 1 IS NULL, 0 IS NULL, NULL IS NULL;
-> 0, 0, 1
mysql> SELECT 1 IS NOT NULL, 0 IS NOT NULL, NULL IS NOT NULL;
-> 1, 1, 0
```

To work well with ODBC programs, MySQL supports the following extra features when using `IS NULL`:

- You can find the row that contains the most recent `AUTO_INCREMENT` value by issuing a statement of the following form immediately after generating the value:

```
SELECT * FROM tbl_name WHERE auto_col IS NULL
```

This behavior can be disabled by setting `SQL_AUTO_IS_NULL=0`. See [Section 13.5.3, “SET Syntax”](#).

- For `DATE` and `DATETIME` columns that are declared as `NOT NULL`, you can find the special date `'0000-00-00'` by using a statement like this:

```
SELECT * FROM tbl_name WHERE date_column IS NULL
```

This is needed to get some ODBC applications to work because ODBC does not support a `'0000-00-00'` date value.

- `expr BETWEEN min AND max`

If `expr` is greater than or equal to `min` and `expr` is less than or equal to `max`, `BETWEEN` returns 1, otherwise it returns 0. This is equivalent to the expression `(min <= expr AND expr <= max)` if all the arguments are of the same type. Otherwise type conversion takes place according to the rules described in [Section 12.1.2, “Type Conversion in Expression Evaluation”](#), but applied to all the three arguments. **Note:** Before MySQL 4.0.5, arguments were converted to the type of `expr` instead.

```
mysql> SELECT 1 BETWEEN 2 AND 3;
-> 0
mysql> SELECT 'b' BETWEEN 'a' AND 'c';
-> 1
mysql> SELECT 2 BETWEEN 2 AND '3';
-> 1
mysql> SELECT 2 BETWEEN 2 AND 'x-3';
-> 0
```

- `expr NOT BETWEEN min AND max`

This is the same as `NOT (expr BETWEEN min AND max)`.

- `COALESCE(value, ...)`

Returns the first non-`NULL` value in the list, or `NULL` if there are no non-`NULL` values.

```
mysql> SELECT COALESCE(NULL,1);
-> 1
mysql> SELECT COALESCE(NULL,NULL,NULL);
-> NULL
```

`COALESCE()` was added in MySQL 3.23.3.

- `GREATEST(value1,value2,...)`

With two or more arguments, returns the largest (maximum-valued) argument. The arguments are

compared using the same rules as for `LEAST()`.

```
mysql> SELECT GREATEST(2,0);
-> 2
mysql> SELECT GREATEST(34.0,3.0,5.0,767.0);
-> 767.0
mysql> SELECT GREATEST('B','A','C');
-> 'C'
```

`GREATEST()` returns `NULL` only if all arguments are `NULL`.

Before MySQL 3.22.5, you can use `MAX()` instead of `GREATEST()`.

- `expr IN (value, ...)`

Returns `1` if `expr` is equal to any of the values in the `IN` list, else returns `0`. If all values are constants, they are evaluated according to the type of `expr` and sorted. The search for the item then is done using a binary search. This means `IN` is very quick if the `IN` value list consists entirely of constants. Otherwise, type conversion takes place according to the rules described in [Section 12.1.2, “Type Conversion in Expression Evaluation”](#), but applied to all the arguments.

```
mysql> SELECT 2 IN (0,3,5,7);
-> 0
mysql> SELECT 'wefwf' IN ('wee','wefwf','weg');
-> 1
```

You should never mix quoted and unquoted values in an `IN` list because the comparison rules for quoted values (such as strings) and unquoted values (such as numbers) differ. Mixing types may therefore lead to inconsistent results. For example, do not write an `IN` expression like this:

```
SELECT val1 FROM tbl1 WHERE val1 IN (1,2,'a');
```

Instead, write it like this:

```
SELECT val1 FROM tbl1 WHERE val1 IN ('1','2','a');
```

The number of values in the `IN` list is only limited by the `max_allowed_packet` value.

To comply with the SQL standard, from MySQL 4.1.0 on `IN` returns `NULL` not only if the expression on the left hand side is `NULL`, but also if no match is found in the list and one of the expressions in the list is `NULL`.

From MySQL 4.1.0 on, `IN()` syntax can also be used to write certain types of subqueries. See [Section 13.2.8.3, “Subqueries with ANY, IN, and SOME”](#).

- `expr NOT IN (value, ...)`

This is the same as `NOT (expr IN (value, ...))`.

- `ISNULL(expr)`

If `expr` is `NULL`, `ISNULL()` returns `1`, otherwise it returns `0`.

```
mysql> SELECT ISNULL(1+1);
-> 0
mysql> SELECT ISNULL(1/0);
-> 1
```

`ISNULL()` can be used instead of `=` to test whether a value is `NULL`. (Comparing a value to `NULL`

using = always yields false.)

The `ISNULL()` function shares some special behaviors with the `IS NULL` comparison operator. See the description of `IS NULL`.

- `INTERVAL(N,N1,N2,N3,...)`

Returns 0 if  $N < N1$ , 1 if  $N < N2$  and so on or -1 if  $N$  is `NULL`. All arguments are treated as integers. It is required that  $N1 < N2 < N3 < \dots < Nn$  for this function to work correctly. This is because a binary search is used (very fast).

```
mysql> SELECT INTERVAL(23, 1, 15, 17, 30, 44, 200);
-> 3
mysql> SELECT INTERVAL(10, 1, 10, 100, 1000);
-> 2
mysql> SELECT INTERVAL(22, 23, 30, 44, 200);
-> 0
```

- `LEAST(value1,value2,...)`

With two or more arguments, returns the smallest (minimum-valued) argument. The arguments are compared using the following rules:

- If the return value is used in an `INTEGER` context or all arguments are integer-valued, they are compared as integers.
- If the return value is used in a `REAL` context or all arguments are real-valued, they are compared as reals.
- If any argument is a case-sensitive string, the arguments are compared as case-sensitive strings.
- In all other cases, the arguments are compared as case-insensitive strings.

```
mysql> SELECT LEAST(2,0);
-> 0
mysql> SELECT LEAST(34.0,3.0,5.0,767.0);
-> 3.0
mysql> SELECT LEAST('B','A','C');
-> 'A'
```

`LEAST()` returns `NULL` only if all arguments are `NULL`.

Before MySQL 3.22.5, you can use `MIN()` instead of `LEAST()`.

Note that the preceding conversion rules can produce strange results in some borderline cases:

```
mysql> SELECT CAST(LEAST(3600, 9223372036854775808.0) as SIGNED);
-> -9223372036854775808
```

This happens because MySQL reads `9223372036854775808.0` in an integer context. The integer representation is not good enough to hold the value, so it wraps to a signed integer.

## 12.1.4. Logical Operators

In SQL, all logical operators evaluate to `TRUE`, `FALSE`, or `NULL` (`UNKNOWN`). In MySQL, these are implemented as 1 (`TRUE`), 0 (`FALSE`), and `NULL`. Most of this is common to different SQL database servers, although some servers may return any non-zero value for `TRUE`.

- NOT, !

Logical NOT. Evaluates to 1 if the operand is 0, to 0 if the operand is non-zero, and NOT NULL returns NULL.

```
mysql> SELECT NOT 10;
-> 0
mysql> SELECT NOT 0;
-> 1
mysql> SELECT NOT NULL;
-> NULL
mysql> SELECT ! (1+1);
-> 0
mysql> SELECT ! 1+1;
-> 1
```

The last example produces 1 because the expression evaluates the same way as (!1)+1.

- AND, &&

Logical AND. Evaluates to 1 if all operands are non-zero and not NULL, to 0 if one or more operands are 0, otherwise NULL is returned.

```
mysql> SELECT 1 && 1;
-> 1
mysql> SELECT 1 && 0;
-> 0
mysql> SELECT 1 && NULL;
-> NULL
mysql> SELECT 0 && NULL;
-> 0
mysql> SELECT NULL && 0;
-> 0
```

Please note that MySQL versions prior to 4.0.5 stop evaluation when a NULL is encountered, rather than continuing the process to check for possible 0 values. This means that in these versions, SELECT (NULL AND 0) returns NULL instead of 0. As of MySQL 4.0.5, the code has been re-engineered so that the result is always as prescribed by the SQL standards while still using the optimization wherever possible.

- OR, ||

Logical OR. When both operands are non-NULL, the result is 1 if any operand is non-zero, and 0 otherwise. With a NULL operand, the result is 1 if the other operand is non-zero, and NULL otherwise. If both operands are NULL, the result is NULL.

```
mysql> SELECT 1 || 1;
-> 1
mysql> SELECT 1 || 0;
-> 1
mysql> SELECT 0 || 0;
-> 0
mysql> SELECT 0 || NULL;
-> NULL
mysql> SELECT 1 || NULL;
-> 1
```

- XOR

Logical XOR. Returns NULL if either operand is NULL. For non-NULL operands, evaluates to 1 if an odd number of operands is non-zero, otherwise 0 is returned.

```
mysql> SELECT 1 XOR 1;
-> 0
mysql> SELECT 1 XOR 0;
-> 1
```



```
mysql> SELECT 1 XOR NULL;
-> NULL
mysql> SELECT 1 XOR 1 XOR 1;
-> 1
```

$a \text{ XOR } b$  is mathematically equal to  $(a \text{ AND } (\text{NOT } b)) \text{ OR } ((\text{NOT } a) \text{ AND } b)$ .

XOR was added in MySQL 4.0.2.

## 12.2. Control Flow Functions

- `CASE value WHEN [compare_value] THEN result [WHEN [compare_value] THEN result ...] [ELSE result] END`

```
CASE WHEN [condition] THEN result [WHEN [condition] THEN result
...] [ELSE result] END
```

The first version returns the *result* where  $value=compare\_value$ . The second version returns the result for the first condition that is true. If there was no matching result value, the result after `ELSE` is returned, or `NULL` if there is no `ELSE` part.

```
mysql> SELECT CASE 1 WHEN 1 THEN 'one'
->      WHEN 2 THEN 'two' ELSE 'more' END;
-> 'one'
mysql> SELECT CASE WHEN 1>0 THEN 'true' ELSE 'false' END;
-> 'true'
mysql> SELECT CASE BINARY 'B'
->      WHEN 'a' THEN 1 WHEN 'b' THEN 2 END;
-> NULL
```

Before MySQL 4.1, the type of the return value (`INTEGER`, `DOUBLE`, or `STRING`) is the same as the type of the first returned value (the expression after the first `THEN`). From MySQL 4.1.0, the default return type is the compatible aggregated type of all return values.

Note that `CASE` evaluation depends also on the context in which it is used. If used in string context, the result is returned as a string. If used in numeric context, the result is returned decimal, real, or integer value.

`CASE` was added in MySQL 3.23.3.

- `IF(expr1,expr2,expr3)`

If *expr1* is `TRUE` ( $expr1 \neq 0$  and  $expr1 \neq \text{NULL}$ ) then `IF()` returns *expr2*; otherwise it returns *expr3*. `IF()` returns a numeric or string value, depending on the context in which it is used.

```
mysql> SELECT IF(1>2,2,3);
-> 3
mysql> SELECT IF(1<2,'yes','no');
-> 'yes'
mysql> SELECT IF(STRCMP('test','test1'),'no','yes');
-> 'no'
```

If only one of *expr2* or *expr3* is explicitly `NULL`, the result type of the `IF()` function is the type of non-`NULL` expression. (This behavior was implemented in MySQL 4.0.3.)

*expr1* is evaluated as an integer value, which means that if you are testing floating-point or string values, you should do so using a comparison operation.

```
mysql> SELECT IF(0.1,1,0);
-> 0
mysql> SELECT IF(0.1<>0,1,0);
-> 1
```

In the first case shown, `IF(0.1)` returns 0 because 0.1 is converted to an integer value, resulting in a test of `IF(0)`. This may not be what you expect. In the second case, the comparison tests the original floating-point value to see whether it is non-zero. The result of the comparison is used as an integer.

The default return type of `IF()` (which may matter when it is stored into a temporary table) is calculated in MySQL 3.23 as follows:

Expression	Return Value
<code>expr2</code> or <code>expr3</code> returns a string	string
<code>expr2</code> or <code>expr3</code> returns a floating-point value	floating-point
<code>expr2</code> or <code>expr3</code> returns an integer	integer

If `expr2` and `expr3` are both strings, the result is case sensitive if either string is case sensitive (starting from MySQL 3.23.51).

- `IFNULL(expr1, expr2)`

If `expr1` is not `NULL`, `IFNULL()` returns `expr1`; otherwise it returns `expr2`. `IFNULL()` returns a numeric or string value, depending on the context in which it is used.

```
mysql> SELECT IFNULL(1,0);
-> 1
mysql> SELECT IFNULL(NULL,10);
-> 10
mysql> SELECT IFNULL(1/0,10);
-> 10
mysql> SELECT IFNULL(1/0,'yes');
-> 'yes'
```

In MySQL 4.0.6 and above, the default result value of `IFNULL(expr1, expr2)` is the more “general” of the two expressions, in the order `STRING`, `REAL`, or `INTEGER`. The difference from earlier MySQL versions is mostly notable when you create a table based on expressions or MySQL has to internally store a value from `IFNULL()` in a temporary table.

```
mysql> CREATE TABLE tmp SELECT IFNULL(1,'test') AS test;
mysql> DESCRIBE tmp;
+-----+-----+-----+-----+-----+-----+
| Field | Type  | Null | Key | Default | Extra |
+-----+-----+-----+-----+-----+-----+
| test  | char(4) |      |     |         |       |
+-----+-----+-----+-----+-----+-----+
```

As of MySQL 4.0.6, the type for the `test` column is `CHAR(4)`, whereas in earlier versions the type would be `BIGINT`.

- `NULLIF(expr1, expr2)`

Returns `NULL` if `expr1 = expr2` is true, otherwise returns `expr1`. This is the same as `CASE WHEN expr1 = expr2 THEN NULL ELSE expr1 END`.

```
mysql> SELECT NULLIF(1,1);
-> NULL
mysql> SELECT NULLIF(1,2);
-> 1
```

Note that MySQL evaluates *expr1* twice if the arguments are not equal.

`NULLIF()` was added in MySQL 3.23.15.

## 12.3. String Functions

String-valued functions return `NULL` if the length of the result would be greater than the value of the `max_allowed_packet` system variable. See [Section 7.5.2, “Tuning Server Parameters”](#).

For functions that operate on string positions, the first position is numbered 1.

- `ASCII(str)`

Returns the numeric value of the leftmost character of the string *str*. Returns 0 if *str* is the empty string. Returns `NULL` if *str* is `NULL`. `ASCII()` works for characters with numeric values from 0 to 255.

```
mysql> SELECT ASCII('2');
-> 50
mysql> SELECT ASCII(2);
-> 50
mysql> SELECT ASCII('dx');
-> 100
```

See also the `ORD()` function.

- `BIN(N)`

Returns a string representation of the binary value of *N*, where *N* is a longlong (`BIGINT`) number. This is equivalent to `CONV(N, 10, 2)`. Returns `NULL` if *N* is `NULL`.

```
mysql> SELECT BIN(12);
-> '1100'
```

- `BIT_LENGTH(str)`

Returns the length of the string *str* in bits.

```
mysql> SELECT BIT_LENGTH('text');
-> 32
```

`BIT_LENGTH()` was added in MySQL 4.0.2.

- `CHAR(N, ... [USING charset_name])`

`CHAR()` interprets each argument *N* as an integer and returns a string consisting of the characters given by the code values of those integers. `NULL` values are skipped.

```
mysql> SELECT CHAR(77,121,83,81,'76');
-> 'MySQL'
mysql> SELECT CHAR(77,77.3,'77.3');
-> 'MMM'
```

`CHAR()` returns a string in the connection character set. As of MySQL 4.1.16, the optional `USING` clause may be used to produce a string in a given character set:

```
mysql> SELECT CHARSET(CHAR(0x65)), CHARSET(CHAR(0x65 USING utf8));
+-----+-----+
| CHARSET(CHAR(0x65)) | CHARSET(CHAR(0x65 USING utf8)) |
+-----+-----+
| latin1              | utf8                             |
+-----+-----+
```

- `CHAR_LENGTH(str)`

Returns the length of the string *str*, measured in characters. A multi-byte character counts as a single character. This means that for a string containing five two-byte characters, `LENGTH()` returns 10, whereas `CHAR_LENGTH()` returns 5.

- `CHARACTER_LENGTH(str)`

`CHARACTER_LENGTH()` is a synonym for `CHAR_LENGTH()`.

- `CONCAT(str1,str2,...)`

Returns the string that results from concatenating the arguments. May have one or more arguments. If all arguments are non-binary strings, the result is a non-binary string. If the arguments include any binary strings, the result is a binary string. A numeric argument is converted to its equivalent binary string form; if you want to avoid that, you can use an explicit type cast, as in this example:

```
SELECT CONCAT(CAST(int_col AS CHAR), char_col);
```

`CONCAT()` returns `NULL` if any argument is `NULL`.

```
mysql> SELECT CONCAT('My', 'S', 'QL');
-> 'MySQL'
mysql> SELECT CONCAT('My', NULL, 'QL');
-> NULL
mysql> SELECT CONCAT(14.3);
-> '14.3'
```

- `CONCAT_WS(separator,str1,str2,...)`

`CONCAT_WS()` stands for Concatenate With Separator and is a special form of `CONCAT()`. The first argument is the separator for the rest of the arguments. The separator is added between the strings to be concatenated. The separator can be a string, as can the rest of the arguments. If the separator is `NULL`, the result is `NULL`.

```
mysql> SELECT CONCAT_WS(',', 'First name', 'Second name', 'Last Name');
-> 'First name,Second name,Last Name'
mysql> SELECT CONCAT_WS(',', 'First name', NULL, 'Last Name');
-> 'First name,Last Name'
```

`CONCAT_WS()` skips any `NULL` values after the separator argument. Before MySQL 4.0.14, `CONCAT_WS()` skips empty strings as well as `NULL` values.

- `CONV(N,from_base,to_base)`

Converts numbers between different number bases. Returns a string representation of the number *N*, converted from base *from\_base* to base *to\_base*. Returns `NULL` if any argument is `NULL`. The argument *N* is interpreted as an integer, but may be specified as an integer or a string. The minimum base is 2 and the maximum base is 36. If *to\_base* is a negative number, *N* is regarded as a signed number. Otherwise, *N* is treated as unsigned. `CONV()` works with 64-bit precision.

```
mysql> SELECT CONV('a',16,2);
-> '1010'
mysql> SELECT CONV('6E',18,8);
```

```
mysql> SELECT CONV(-17,10,-18);
-> '172'
mysql> SELECT CONV(-17,10,-18);
-> '-H'
mysql> SELECT CONV(10+'10'+'10'+0xa,10,10);
-> '40'
```

- `ELT(N, str1, str2, str3, ...)`

Returns `str1` if `N = 1`, `str2` if `N = 2`, and so on. Returns `NULL` if `N` is less than 1 or greater than the number of arguments. `ELT()` is the complement of `FIELD()`.

```
mysql> SELECT ELT(1, 'ej', 'Heja', 'hej', 'foo');
-> 'ej'
mysql> SELECT ELT(4, 'ej', 'Heja', 'hej', 'foo');
-> 'foo'
```

- `EXPORT_SET(bits, on, off[, separator[, number_of_bits]])`

Returns a string such that for every bit set in the value `bits`, you get an `on` string and for every reset bit, you get an `off` string. Bits in `bits` are examined from right to left (from low-order to high-order bits). Strings are added to the result from left to right, separated by the `separator` string (the default being the comma character `,`). The number of bits examined is given by `number_of_bits` (defaults to 64).

```
mysql> SELECT EXPORT_SET(5,'Y','N',' ',4);
-> 'Y,N,Y,N'
mysql> SELECT EXPORT_SET(6,'1','0',' ',10);
-> '0,1,1,0,0,0,0,0,0,0'
```

- `FIELD(str, str1, str2, str3, ...)`

Returns the index (position) of `str` in the `str1, str2, str3, ...` list. Returns 0 if `str` is not found.

If all arguments to `FIELD()` are strings, all arguments are compared as strings. If all arguments are numbers, they are compared as numbers. Otherwise, the arguments are compared as double.

If `str` is `NULL`, the return value is 0 because `NULL` fails equality comparison with any value. `FIELD()` is the complement of `ELT()`.

```
mysql> SELECT FIELD('ej', 'Hej', 'ej', 'Heja', 'hej', 'foo');
-> 2
mysql> SELECT FIELD('fo', 'Hej', 'ej', 'Heja', 'hej', 'foo');
-> 0
```

- `FIND_IN_SET(str, strlist)`

Returns a value in the range of 1 to `N` if the string `str` is in the string list `strlist` consisting of `N` substrings. A string list is a string composed of substrings separated by `,` characters. If the first argument is a constant string and the second is a column of type `SET`, the `FIND_IN_SET()` function is optimized to use bit arithmetic. Returns 0 if `str` is not in `strlist` or if `strlist` is the empty string. Returns `NULL` if either argument is `NULL`. This function does not work properly if the first argument contains a comma (`,`) character.

```
mysql> SELECT FIND_IN_SET('b', 'a,b,c,d');
-> 2
```

- `FORMAT(X, D)`

Formats the number `X` to a format like `'#,###,###.##'`, rounded to `D` decimal places, and re-

turns the result as a string. If *D* is 0, the result has no decimal point or fractional part.

```
mysql> SELECT FORMAT(12332.123456, 4);
-> '12,332.1235'
mysql> SELECT FORMAT(12332.1,4);
-> '12,332.1000'
mysql> SELECT FORMAT(12332.2,0);
-> '12,332'
```

- `HEX(N_or_S)`

If *N\_or\_S* is a number, returns a string representation of the hexadecimal value of *N*, where *N* is a longlong (`BIGINT`) number. This is equivalent to `CONV(N, 10, 16)`.

From MySQL 4.0.1 and up, if *N\_or\_S* is a string, returns a hexadecimal string representation of *N\_or\_S* where each character in *N\_or\_S* is converted to two hexadecimal digits.

```
mysql> SELECT HEX(255);
-> 'FF'
mysql> SELECT 0x616263;
-> 'abc'
mysql> SELECT HEX('abc');
-> 616263
```

- `INSERT(str, pos, len, newstr)`

Returns the string *str*, with the substring beginning at position *pos* and *len* characters long replaced by the string *newstr*. Returns the original string if *pos* is not within the length of the string. Replaces the rest of the string from position *pos* is *len* is not within the length of the rest of the string. Returns `NULL` if any argument is `NULL`.

```
mysql> SELECT INSERT('Quadratic', 3, 4, 'What');
-> 'QuWhattic'
mysql> SELECT INSERT('Quadratic', -1, 4, 'What');
-> 'Quadratic'
mysql> SELECT INSERT('Quadratic', 3, 100, 'What');
-> 'QuWhat'
```

This function is multi-byte safe.

- `INSTR(str, substr)`

Returns the position of the first occurrence of substring *substr* in string *str*. This is the same as the two-argument form of `LOCATE()`, except that the order of the arguments is reversed.

```
mysql> SELECT INSTR('foobarbar', 'bar');
-> 4
mysql> SELECT INSTR('xbar', 'foobar');
-> 0
```

This function is multi-byte safe. In MySQL 3.23, this function is case sensitive. For 4.0 on, it is case sensitive only if either argument is a binary string.

- `LCASE(str)`

`LCASE()` is a synonym for `LOWER()`.

- `LEFT(str, len)`

Returns the leftmost *len* characters from the string *str*.

```
mysql> SELECT LEFT('foobarbar', 5);
-> 'fooba'
```

- `LENGTH(str)`

Returns the length of the string *str*, measured in bytes. A multi-byte character counts as multiple bytes. This means that for a string containing five two-byte characters, `LENGTH()` returns 10, whereas `CHAR_LENGTH()` returns 5.

```
mysql> SELECT LENGTH('text');
-> 4
```

- `LOAD_FILE(file_name)`

Reads the file and returns the file contents as a string. To use this function, the file must be located on the server host, you must specify the full pathname to the file, and you must have the `FILE` privilege. The file must be readable by all and its size less than `max_allowed_packet` bytes.

If the file does not exist or cannot be read because one of the preceding conditions is not satisfied, the function returns `NULL`.

```
mysql> UPDATE t
      SET blob_col=LOAD_FILE('/tmp/picture')
      WHERE id=1;
```

Before MySQL 3.23, you must read the file inside your application and create an `INSERT` statement to update the database with the file contents. If you are using the MySQL++ library, one way to do this can be found in the MySQL++ manual, available at <http://tangentsoft.net/mysql++/doc/>.

- `LOCATE(substr, str)`, `LOCATE(substr, str, pos)`

The first syntax returns the position of the first occurrence of substring *substr* in string *str*. The second syntax returns the position of the first occurrence of substring *substr* in string *str*, starting at position *pos*. Returns 0 if *substr* is not in *str*.

```
mysql> SELECT LOCATE('bar', 'foobarbar');
-> 4
mysql> SELECT LOCATE('xbar', 'foobar');
-> 0
mysql> SELECT LOCATE('bar', 'foobarbar', 5);
-> 7
```

This function is multi-byte safe. In MySQL 3.23, this function is case sensitive. For 4.0 on, it is case sensitive only if either argument is a binary string.

- `LOWER(str)`

Returns the string *str* with all characters changed to lowercase according to the current character set mapping. The default is `latin1` (cp1252 West European).

```
mysql> SELECT LOWER('QUADRATICALLY');
-> 'quadratically'
```

This function is multi-byte safe.

- `LPAD(str, len, padstr)`

Returns the string *str*, left-padded with the string *padstr* to a length of *len* characters. If *str* is longer than *len*, the return value is shortened to *len* characters.

```
mysql> SELECT LPAD('hi',4,'??');
-> '??hi'
```

```
mysql> SELECT LPAD('hi',1,'??');
-> 'h'
```

- `LTRIM(str)`

Returns the string *str* with leading space characters removed.

```
mysql> SELECT LTRIM(' barbar');
-> 'barbar'
```

This function is multi-byte safe.

- `MAKE_SET(bits,str1,str2,...)`

Returns a set value (a string containing substrings separated by ',' characters) consisting of the strings that have the corresponding bit in *bits* set. *str1* corresponds to bit 0, *str2* to bit 1, and so on. `NULL` values in *str1*, *str2*, ... are not appended to the result.

```
mysql> SELECT MAKE_SET(1,'a','b','c');
-> 'a'
mysql> SELECT MAKE_SET(1 | 4,'hello','nice','world');
-> 'hello,world'
mysql> SELECT MAKE_SET(1 | 4,'hello','nice',NULL,'world');
-> 'hello'
mysql> SELECT MAKE_SET(0,'a','b','c');
-> ''
```

- `MID(str,pos,len)`

`MID(str,pos,len)` is a synonym for `SUBSTRING(str,pos,len)`.

- `OCT(N)`

Returns a string representation of the octal value of *N*, where *N* is a longlong (`BIGINT`) number. This is equivalent to `CONV(N, 10, 8)`. Returns `NULL` if *N* is `NULL`.

```
mysql> SELECT OCT(12);
-> '14'
```

- `OCTET_LENGTH(str)`

`OCTET_LENGTH()` is a synonym for `LENGTH()`.

- `ORD(str)`

If the leftmost character of the string *str* is a multi-byte character, returns the code for that character, calculated from the numeric values of its constituent bytes using this formula:

```
(1st byte code)
+ (2nd byte code × 256)
+ (3rd byte code × 2562) ...
```

If the leftmost character is not a multi-byte character, `ORD()` returns the same value as the `ASCII()` function.

```
mysql> SELECT ORD('2');
-> 50
```

- `POSITION(substr IN str)`



`POSITION(substr IN str)` is a synonym for `LOCATE(substr, str)`.

- `QUOTE(str)`

Quotes a string to produce a result that can be used as a properly escaped data value in an SQL statement. The string is returned enclosed by single quotes and with each instance of single quote ('), backslash (\), ASCII NUL, and Control-Z preceded by a backslash. If the argument is `NULL`, the return value is the word "NULL" without enclosing single quotes. The `QUOTE()` function was added in MySQL 4.0.3.

```
mysql> SELECT QUOTE('Don\'t!');
-> 'Don\'t!'
mysql> SELECT QUOTE(NULL);
-> NULL
```

- `REPEAT(str, count)`

Returns a string consisting of the string `str` repeated `count` times. If `count` is less than 1, returns an empty string. Returns `NULL` if `str` or `count` are `NULL`.

```
mysql> SELECT REPEAT('MySQL', 3);
-> 'MySQLMySQLMySQL'
```

- `REPLACE(str, from_str, to_str)`

Returns the string `str` with all occurrences of the string `from_str` replaced by the string `to_str`. `REPLACE()` performs a case-sensitive match when searching for `from_str`.

```
mysql> SELECT REPLACE('www.mysql.com', 'w', 'Ww');
-> 'WwWwWw.mysql.com'
```

This function is multi-byte safe.

- `REVERSE(str)`

Returns the string `str` with the order of the characters reversed.

```
mysql> SELECT REVERSE('abc');
-> 'cba'
```

This function is multi-byte safe.

- `RIGHT(str, len)`

Returns the rightmost `len` characters from the string `str`.

```
mysql> SELECT RIGHT('foobarbar', 4);
-> 'rbar'
```

This function is multi-byte safe.

- `RPAD(str, len, padstr)`

Returns the string `str`, right-padded with the string `padstr` to a length of `len` characters. If `str` is longer than `len`, the return value is shortened to `len` characters.

```
mysql> SELECT RPAD('hi', 5, '?');
-> 'hi???'
mysql> SELECT RPAD('hi', 1, '?');
-> 'hi'
```

```
-> 'h'
```

This function is multi-byte safe.

- `RTRIM(str)`

Returns the string *str* with trailing space characters removed.

```
mysql> SELECT RTRIM('barbar ');
-> 'barbar'
```

This function is multi-byte safe.

- `SOUNDEX(str)`

Returns a soundex string from *str*. Two strings that sound almost the same should have identical soundex strings. A standard soundex string is four characters long, but the `SOUNDEX()` function returns an arbitrarily long string. You can use `SUBSTRING()` on the result to get a standard soundex string. All non-alphabetic characters in *str* are ignored. All international alphabetic characters outside the A-Z range are treated as vowels.

```
mysql> SELECT SOUNDEX('Hello');
-> 'H400'
mysql> SELECT SOUNDEX('Quadratically');
-> 'Q36324'
```

**Note:** This function implements the original Soundex algorithm, not the more popular enhanced version (also described by D. Knuth). The difference is that original version discards vowels first and duplicates second, whereas the enhanced version discards duplicates first and vowels second.

- `expr1 SOUNDS LIKE expr2`

This is the same as `SOUNDEX(expr1) = SOUNDEX(expr2)`. It is available beginning with MySQL 4.1.0.

- `SPACE(N)`

Returns a string consisting of *N* space characters.

```
mysql> SELECT SPACE(6);
-> '      '
```

- `SUBSTRING(str, pos)`, `SUBSTRING(str FROM pos)`, `SUBSTRING(str, pos, len)`, `SUBSTRING(str FROM pos FOR len)`

The forms without a *len* argument return a substring from string *str* starting at position *pos*. The forms with a *len* argument return a substring *len* characters long from string *str*, starting at position *pos*. The forms that use `FROM` are standard SQL syntax. Beginning with MySQL 4.1.0, it is possible to use a negative value for *pos*. In this case, the beginning of the substring is *pos* characters from the end of the string, rather than the beginning. A negative value may be used for *pos* in any of the forms of this function.

```
mysql> SELECT SUBSTRING('Quadratically',5);
-> 'ratically'
mysql> SELECT SUBSTRING('foobarbar' FROM 4);
-> 'barbar'
mysql> SELECT SUBSTRING('Quadratically',5,6);
-> 'ratika'
mysql> SELECT SUBSTRING('Sakila', -3);
-> 'ila'
mysql> SELECT SUBSTRING('Sakila', -5, 3);
```

```
mysql> -> 'aki'
mysql> SELECT SUBSTRING('Sakila' FROM -4 FOR 2);
mysql> -> 'ki'
```

This function is multi-byte safe.

If *len* is less than 1, the result is the empty string.

`SUBSTR()` is a synonym for `SUBSTRING()`, added in MySQL 4.1.1.

- `SUBSTRING_INDEX(str, delim, count)`

Returns the substring from string *str* before *count* occurrences of the delimiter *delim*. If *count* is positive, everything to the left of the final delimiter (counting from the left) is returned. If *count* is negative, everything to the right of the final delimiter (counting from the right) is returned. `SUBSTRING_INDEX()` performs a case-sensitive match when searching for *delim*.

```
mysql> SELECT SUBSTRING_INDEX('www.mysql.com', '.', 2);
mysql> -> 'www.mysql'
mysql> SELECT SUBSTRING_INDEX('www.mysql.com', '.', -2);
mysql> -> 'mysql.com'
```

This function is multi-byte safe.

- `TRIM([{BOTH | LEADING | TRAILING} [remstr] FROM] str), TRIM([remstr FROM] str)`

Returns the string *str* with all *remstr* prefixes or suffixes removed. If none of the specifiers `BOTH`, `LEADING`, or `TRAILING` is given, `BOTH` is assumed. *remstr* is optional and, if not specified, spaces are removed.

```
mysql> SELECT TRIM(' bar ');
mysql> -> 'bar'
mysql> SELECT TRIM(LEADING 'x' FROM 'xxxbarxxx');
mysql> -> 'barxxx'
mysql> SELECT TRIM(BOTH 'x' FROM 'xxxbarxxx');
mysql> -> 'bar'
mysql> SELECT TRIM(TRAILING 'xyz' FROM 'barxyz');
mysql> -> 'barx'
```

This function is multi-byte safe.

- `UCASE(str)`

`UCASE()` is a synonym for `UPPER()`.

- `UNHEX(str)`

Performs the inverse operation of `HEX(str)`. That is, it interprets each pair of hexadecimal digits in the argument as a number and converts it to the character represented by the number. The resulting characters are returned as a binary string.

```
mysql> SELECT UNHEX('4D7953514C');
mysql> -> 'MySQL'
mysql> SELECT HEX(UNHEX('string'));
mysql> -> 'string'
mysql> SELECT UNHEX(HEX('1267'));
mysql> -> '1267'
```

`UNHEX()` was added in MySQL 4.1.2.

- `UPPER(str)`

Returns the string *str* with all characters changed to uppercase according to the current character set mapping. The default is `latin1` (cp1252 West European).

```
mysql> SELECT UPPER('Hej');
-> 'HEJ'
```

This function is multi-byte safe.

## 12.3.1. String Comparison Functions

If a string function is given a binary string as an argument, the resulting string is also a binary string. A number converted to a string is treated as a binary string. This affects only comparisons.

Normally, if any expression in a string comparison is case sensitive, the comparison is performed in case-sensitive fashion.

- `expr LIKE pat [ESCAPE 'escape_char']`

Pattern matching using SQL simple regular expression comparison. Returns `1` (`TRUE`) or `0` (`FALSE`). If either *expr* or *pat* is `NULL`, the result is `NULL`.

The pattern need not be a literal string. For example, it can be specified as a string expression or table column.

Per the SQL standard, `LIKE` performs matching on a per-character basis, thus it can produce results different from the `=` comparison operator:

```
mysql> SELECT 'ä' LIKE 'ae' COLLATE latin1_german2_ci;
+-----+
| 'ä' LIKE 'ae' COLLATE latin1_german2_ci |
+-----+
|                                           0 |
+-----+
mysql> SELECT 'ä' = 'ae' COLLATE latin1_german2_ci;
+-----+
| 'ä' = 'ae' COLLATE latin1_german2_ci |
+-----+
|                                           1 |
+-----+
```

With `LIKE` you can use the following two wildcard characters in the pattern:

Character	Description
<code>%</code>	Matches any number of characters, even zero characters
<code>_</code>	Matches exactly one character

```
mysql> SELECT 'David!' LIKE 'David_';
-> 1
mysql> SELECT 'David!' LIKE '%D%v%';
-> 1
```

To test for literal instances of a wildcard character, precede it by the escape character. If you do not specify the `ESCAPE` character, `'\'` is assumed.

String	Description
--------	-------------

<code>\%</code>	Matches one ‘%’ character
<code>\_</code>	Matches one ‘_’ character

```
mysql> SELECT 'David!' LIKE 'David\_';
-> 0
mysql> SELECT 'David_' LIKE 'David\_';
-> 1
```

To specify a different escape character, use the `ESCAPE` clause:

```
mysql> SELECT 'David_' LIKE 'David|_' ESCAPE '|';
-> 1
```

The following two statements illustrate that string comparisons are not case sensitive unless one of the operands is a binary string:

```
mysql> SELECT 'abc' LIKE 'ABC';
-> 1
mysql> SELECT 'abc' LIKE BINARY 'ABC';
-> 0
```

In MySQL, `LIKE` is allowed on numeric expressions. (This is an extension to the standard SQL `LIKE`.)

```
mysql> SELECT 10 LIKE '1%';
-> 1
```

**Note:** Because MySQL uses C escape syntax in strings (for example, ‘\n’ to represent a newline character), you must double any ‘\’ that you use in `LIKE` strings. For example, to search for ‘\n’, specify it as ‘\\n’. To search for ‘\’, specify it as ‘\\\\’; this is because the backslashes are stripped once by the parser and again when the pattern match is made, leaving a single backslash to be matched against. (Exception: At the end of the pattern string, backslash can be specified as ‘\\’.

At the end of the string, backslash stands for itself because there is nothing following to escape.)

- `expr NOT LIKE pat [ESCAPE 'escape_char']`

This is the same as `NOT (expr LIKE pat [ESCAPE 'escape_char'])`.

- `expr NOT REGEXP pat, expr NOT RLIKE pat`

This is the same as `NOT (expr REGEXP pat)`.

- `expr REGEXP pat expr RLIKE pat`

Performs a pattern match of a string expression `expr` against a pattern `pat`. The pattern can be an extended regular expression. The syntax for regular expressions is discussed in [Appendix G, Regular Expressions](#). Returns 1 if `expr` matches `pat`; otherwise it returns 0. If either `expr` or `pat` is `NULL`, the result is `NULL`. `RLIKE` is a synonym for `REGEXP`, provided for `mSQL` compatibility.

The pattern need not be a literal string. For example, it can be specified as a string expression or table column.

**Note:** Because MySQL uses the C escape syntax in strings (for example, ‘\n’ to represent the newline character), you must double any ‘\’ that you use in your `REGEXP` strings.

As of MySQL 3.23.4, `REGEXP` is not case sensitive, except when used with binary strings.

```
mysql> SELECT 'Monty!' REGEXP 'm%y%';
-> 0
mysql> SELECT 'Monty!' REGEXP '.*';
-> 1
mysql> SELECT 'new*\n*line' REGEXP 'new\\*\\.\\*line';
-> 1
mysql> SELECT 'a' REGEXP 'A', 'a' REGEXP BINARY 'A';
-> 1 0
mysql> SELECT 'a' REGEXP '^[a-d]';
-> 1
```

`REGEXP` and `RLIKE` use the current character set when deciding the type of a character. The default is `latin1` (cp1252 West European). **Warning:** These operators are not multi-byte safe.

- `STRCMP(expr1,expr2)`

`STRCMP()` returns 0 if the strings are the same, -1 if the first argument is smaller than the second according to the current sort order, and 1 otherwise.

```
mysql> SELECT STRCMP('text', 'text2');
-> -1
mysql> SELECT STRCMP('text2', 'text');
-> 1
mysql> SELECT STRCMP('text', 'text');
-> 0
```

As of MySQL 4.0, `STRCMP()` uses the current character set when performing comparisons. This makes the default comparison behavior case insensitive unless one or both of the operands are binary strings. Before MySQL 4.0, `STRCMP()` is case sensitive.

## 12.4. Numeric Functions

### 12.4.1. Arithmetic Operators

The usual arithmetic operators are available. The precision of the result is determined according to the following rules:

- Note that in the case of `-`, `+`, and `*`, the result is calculated with `BIGINT` (64-bit) precision if both arguments are integers.
- If one of the arguments is an unsigned integer, and the other argument is also an integer, the result is an unsigned integer.
- If any of the operands of a `+`, `-`, `/`, `*`, `%` is a real or string value, then the precision of the result is the precision of the argument with the maximum precision.
- In multiplication and division, the precision of the result when using two integer values is the precision of the first argument + the value of the `div_precision_increment` global variable. For example, the expression `5.05 / 0.0014` would have a precision of six decimal places (`4.047976`).

These rules are applied for each operation, such that nested calculations imply the precision of each component. Hence, `(14620 / 9432456) / (24250 / 9432456)`, would resolve first to `(0.0014) / (0.0026)`, with the final result having 8 decimal places (`0.57692308`).

Because of these rules and the method they are applied, care should be taken to ensure that components and sub-components of a calculation use the appropriate level of precision. See [Section 12.8, “Cast](#)

Functions and Operators”.

- +

Addition:

```
mysql> SELECT 3+5;
-> 8
```

- -

Subtraction:

```
mysql> SELECT 3-5;
-> -2
```

- -

Unary minus. This operator changes the sign of the argument.

```
mysql> SELECT - 2;
-> -2
```

**Note:** If this operator is used with a `BIGINT`, the return value is also a `BIGINT`. This means that you should avoid using `-` on integers that may have the value of  $-2^{63}$ .

- \*

Multiplication:

```
mysql> SELECT 3*5;
-> 15
mysql> SELECT 18014398509481984*18014398509481984.0;
-> 324518553658426726783156020576256.0
mysql> SELECT 18014398509481984*18014398509481984;
-> 0
```

The result of the last expression is incorrect because the result of the integer multiplication exceeds the 64-bit range of `BIGINT` calculations. (See [Section 11.2, “Numeric Types”](#).)

- /

Division:

```
mysql> SELECT 3/5;
-> 0.60
```

Division by zero produces a `NULL` result:

```
mysql> SELECT 102/(1-1);
-> NULL
```

A division is calculated with `BIGINT` arithmetic only if performed in a context where its result is converted to an integer.

- DIV

Integer division. Similar to `FLOOR( )`, but is safe with `BIGINT` values.

```
mysql> SELECT 5 DIV 2;
```

```
-> 2
```

`DIV` was implemented in MySQL 4.1.0.

## 12.4.2. Mathematical Functions

All mathematical functions return `NULL` in the event of an error.

- `ABS(X)`

Returns the absolute value of `X`.

```
mysql> SELECT ABS(2);
-> 2
mysql> SELECT ABS(-32);
-> 32
```

This function is safe to use with `BIGINT` values.

- `ACOS(X)`

Returns the arc cosine of `X`, that is, the value whose cosine is `X`. Returns `NULL` if `X` is not in the range `-1` to `1`.

```
mysql> SELECT ACOS(1);
-> 0.000000
mysql> SELECT ACOS(1.0001);
-> NULL
mysql> SELECT ACOS(0);
-> 1.570796
```

- `ASIN(X)`

Returns the arc sine of `X`, that is, the value whose sine is `X`. Returns `NULL` if `X` is not in the range `-1` to `1`.

```
mysql> SELECT ASIN(0.2);
-> 0.201358
mysql> SELECT ASIN('foo');
-> 0.000000
```

- `ATAN(X)`

Returns the arc tangent of `X`, that is, the value whose tangent is `X`.

```
mysql> SELECT ATAN(2);
-> 1.107149
mysql> SELECT ATAN(-2);
-> -1.107149
```

- `ATAN(Y, X)`, `ATAN2(Y, X)`

Returns the arc tangent of the two variables `X` and `Y`. It is similar to calculating the arc tangent of `Y / X`, except that the signs of both arguments are used to determine the quadrant of the result.

```
mysql> SELECT ATAN(-2, 2);
-> -0.785398
mysql> SELECT ATAN2(PI(), 0);
-> 1.570796
```



- `CEILING(X)`, `CEIL(X)`

Returns the smallest integer value not less than  $X$ .

```
mysql> SELECT CEILING(1.23);
-> 2
mysql> SELECT CEIL(-1.23);
-> -1
```

Note that the return value is converted to a `BIGINT`.

The `CEIL()` alias was added in MySQL 4.0.6.

- `COS(X)`

Returns the cosine of  $X$ , where  $X$  is given in radians.

```
mysql> SELECT COS(PI());
-> -1.000000
```

- `COT(X)`

Returns the cotangent of  $X$ .

```
mysql> SELECT COT(12);
-> -1.57267341
mysql> SELECT COT(0);
-> NULL
```

- `CRC32(expr)`

Computes a cyclic redundancy check value and returns a 32-bit unsigned value. The result is `NULL` if the argument is `NULL`. The argument is expected to be a string and (if possible) is treated as one if it is not.

```
mysql> SELECT CRC32('MySQL');
-> 3259397556
```

`CRC32()` is available as of MySQL 4.1.0.

- `DEGREES(X)`

Returns the argument  $X$ , converted from radians to degrees.

```
mysql> SELECT DEGREES(PI());
-> 180.000000
```

- `EXP(X)`

Returns the value of  $e$  (the base of natural logarithms) raised to the power of  $X$ .

```
mysql> SELECT EXP(2);
-> 7.389056
mysql> SELECT EXP(-2);
-> 0.135335
```

- `FLOOR(X)`

Returns the largest integer value not greater than  $X$ .

```
mysql> SELECT FLOOR(1.23);
```

```
mysql> SELECT FLOOR(-1.23);
-> -2
```

Note that the return value is converted to a `BIGINT`.

- `FORMAT(X,D)`

Formats the number `X` to a format like '`#,###,###.##`', rounded to `D` decimal places, and returns the result as a string. For details, see [Section 12.3, “String Functions”](#).

- `LN(X)`

Returns the natural logarithm of `X`; that is, the base- $e$  logarithm of `X`.

```
mysql> SELECT LN(2);
-> 0.693147
mysql> SELECT LN(-2);
-> NULL
```

This function was added in MySQL 4.0.3. It is synonymous with `LOG(X)`.

- `LOG(X)`, `LOG(B,X)`

If called with one parameter, this function returns the natural logarithm of `X`.

```
mysql> SELECT LOG(2);
-> 0.693147
mysql> SELECT LOG(-2);
-> NULL
```

If called with two parameters, this function returns the logarithm of `X` for an arbitrary base `B`.

```
mysql> SELECT LOG(2,65536);
-> 16.000000
mysql> SELECT LOG(1,100);
-> NULL
```

The arbitrary base option was added in MySQL 4.0.3. `LOG(B,X)` is equivalent to `LOG(X) / LOG(B)`.

- `LOG2(X)`

Returns the base-2 logarithm of `X`.

```
mysql> SELECT LOG2(65536);
-> 16.000000
mysql> SELECT LOG2(-100);
-> NULL
```

`LOG2()` is useful for finding out how many bits a number would require for storage. This function was added in MySQL 4.0.3. In earlier versions, you can use `LOG(X) / LOG(2)` instead.

- `LOG10(X)`

Returns the base-10 logarithm of `X`.

```
mysql> SELECT LOG10(2);
-> 0.301030
mysql> SELECT LOG10(100);
-> 2.000000
mysql> SELECT LOG10(-100);
-> NULL
```

- `MOD(N, M)`, `N % M`, `N MOD M`

Modulo operation. Returns the remainder of  $N$  divided by  $M$ .

```
mysql> SELECT MOD(234, 10);
-> 4
mysql> SELECT 253 % 7;
-> 1
mysql> SELECT MOD(29,9);
-> 2
mysql> SELECT 29 MOD 9;
-> 2
```

This function is safe to use with `BIGINT` values. The `N MOD M` syntax works only as of MySQL 4.1.0.

As of MySQL 4.1.7, `MOD()` works on values that have a fractional part and returns the exact remainder after division:

```
mysql> SELECT MOD(34.5, 3);
-> 1.5
```

Before MySQL 4.1.7, `MOD()` rounds arguments with a fractional value to integers and returns an integer result:

```
mysql> SELECT MOD(34.5, 3);
-> 2
```

- `PI()`

Returns the value of  $\pi$  (pi). The default number of decimal places displayed is five, but MySQL uses the full double-precision value internally.

```
mysql> SELECT PI();
-> 3.141593
mysql> SELECT PI()+0.00000000000000000000;
-> 3.141592653589793116
```

- `POW(X, Y)`, `POWER(X, Y)`

Returns the value of  $X$  raised to the power of  $Y$ .

```
mysql> SELECT POW(2, 2);
-> 4.000000
mysql> SELECT POW(2, -2);
-> 0.250000
```

- `RADIANS(X)`

Returns the argument  $X$ , converted from degrees to radians. (Note that  $\pi$  radians equals 180 degrees.)

```
mysql> SELECT RADIANS(90);
-> 1.570796
```

- `RAND()`, `RAND(N)`

Returns a random floating-point value  $v$  between 0 and 1 inclusive (that is, in the range  $0 \leq v \leq 1.0$ ). If an integer argument  $N$  is specified, it is used as the seed value, which produces a repeatable

sequence of column values.

```
mysql> SELECT RAND();
-> 0.9233482386203
mysql> SELECT RAND(20);
-> 0.15888261251047
mysql> SELECT RAND(20);
-> 0.15888261251047
mysql> SELECT RAND();
-> 0.63553050033332
mysql> SELECT RAND();
-> 0.70100469486881
mysql> SELECT RAND(20);
-> 0.15888261251047
```

To obtain a random integer  $R$  in the range  $i \leq R \leq j$ , use the expression `FLOOR(i + RAND() * (j - i))`. For example, to obtain a random integer in the range of 7 to 12 inclusive, you could use the following statement:

```
SELECT FLOOR(7 + (RAND() * 5));
```

You cannot use a column with `RAND()` values in an `ORDER BY` clause, because `ORDER BY` would evaluate the column multiple times. However, as of MySQL 3.23, you can retrieve rows in random order like this:

```
mysql> SELECT * FROM tbl_name ORDER BY RAND();
```

`ORDER BY RAND()` combined with `LIMIT` is useful for selecting a random sample from a set of rows:

```
mysql> SELECT * FROM table1, table2 WHERE a=b AND c<d -> ORDER BY RAND() LIMIT 1000;
```

Note that `RAND()` in a `WHERE` clause is re-evaluated every time the `WHERE` is executed.

`RAND()` is not meant to be a perfect random generator, but instead is a fast way to generate *ad hoc* random numbers which is portable between platforms for the same MySQL version.

- `ROUND(X)`, `ROUND(X, D)`

Returns the argument  $X$ , rounded to the nearest integer. With two arguments, returns  $X$  rounded to  $D$  decimal places.  $D$  can be negative to cause  $D$  digits left of the decimal point of the value  $X$  to become zero.

```
mysql> SELECT ROUND(-1.23);
-> -1
mysql> SELECT ROUND(-1.58);
-> -2
mysql> SELECT ROUND(1.58);
-> 2
mysql> SELECT ROUND(1.298, 1);
-> 1.3
mysql> SELECT ROUND(1.298, 0);
-> 1
mysql> SELECT ROUND(23.298, -1);
-> 20
```

The return type is the same type as that of the first argument (assuming that it is integer, double, or decimal). This means that for an integer argument, the result is an integer (no decimal places).

The behavior of `ROUND()` when the argument is halfway between two integers depends on the C library implementation. Different implementations round to the nearest even number, always up, always down, or always toward zero. If you need one kind of rounding, you should use a well-defined function such as `TRUNCATE()` or `FLOOR()` instead.

- `SIGN(X)`

Returns the sign of the argument as `-1`, `0`, or `1`, depending on whether `X` is negative, zero, or positive.

```
mysql> SELECT SIGN(-32);
-> -1
mysql> SELECT SIGN(0);
-> 0
mysql> SELECT SIGN(234);
-> 1
```

- `SIN(X)`

Returns the sine of `X`, where `X` is given in radians.

```
mysql> SELECT SIN(PI());
-> 1.2246063538224e-16
mysql> SELECT ROUND(SIN(PI()));
-> 0
```

- `SQRT(X)`

Returns the square root of a non-negative number `X`.

```
mysql> SELECT SQRT(4);
-> 2
mysql> SELECT SQRT(20);
-> 4.4721359549996
mysql> SELECT SQRT(-16);
-> NULL
```

- `TAN(X)`

Returns the tangent of `X`, where `X` is given in radians.

```
mysql> SELECT TAN(PI());
-> -1.2246063538224e-16
mysql> SELECT TAN(PI()+1);
-> 1.5574077246549
```

- `TRUNCATE(X,D)`

Returns the number `X`, truncated to `D` decimal places. If `D` is `0`, the result has no decimal point or fractional part. `D` can be negative to cause `D` digits left of the decimal point of the value `X` to become zero.

```
mysql> SELECT TRUNCATE(1.223,1);
-> 1.2
mysql> SELECT TRUNCATE(1.999,1);
-> 1.9
mysql> SELECT TRUNCATE(1.999,0);
-> 1
mysql> SELECT TRUNCATE(-1.999,1);
-> -1.9
mysql> SELECT TRUNCATE(122,-2);
-> 100
mysql> SELECT TRUNCATE(10.28*100,0);
-> 1027
```

Starting from MySQL 3.23.51, all numbers are rounded toward zero.

## 12.5. Date and Time Functions

This section describes the functions that can be used to manipulate temporal values. See [Section 11.3, “Date and Time Types”](#), for a description of the range of values each date and time type has and the valid formats in which values may be specified.

Here is an example that uses date functions. The following query selects all rows with a *date\_col* value from within the last 30 days:

```
mysql> SELECT something FROM tbl_name
-> WHERE DATE_SUB(CURDATE(),INTERVAL 30 DAY) <= date_col;
```

Note that the query also selects rows with dates that lie in the future.

Functions that expect date values usually accept datetime values and ignore the time part. Functions that expect time values usually accept datetime values and ignore the date part.

Functions that return the current date or time each are evaluated only once per query at the start of query execution. This means that multiple references to a function such as `NOW()` within a single query always produce the same result (for our purposes a single query also includes a call to a stored routine or trigger and all sub-routines called by that routine/trigger). This principle also applies to `CURDATE()`, `CURTIME()`, `UTC_DATE()`, `UTC_TIME()`, `UTC_TIMESTAMP()`, and to any of their synonyms.

Beginning with MySQL 4.1.3, the `CURRENT_TIMESTAMP()`, `CURRENT_TIME()`, `CURRENT_DATE()`, and `FROM_UNIXTIME()` functions return values in the connection's current time zone, which is available as the value of the `time_zone` system variable. In addition, `UNIX_TIMESTAMP()` assumes that its argument is a datetime value in the current time zone. See [Section 5.10.8, “MySQL Server Time Zone Support”](#).

Some date functions can be used with “zero” dates or incomplete dates such as `'2001-11-00'`, whereas others cannot. Functions that extract parts of dates typically work with incomplete dates. For example:

```
mysql> SELECT DAYOFMONTH('2001-11-00'), MONTH('2005-00-00');
-> 0, 0
```

Other functions expect complete dates and return `NULL` for incomplete dates. These include functions that perform date arithmetic or that map parts of dates to names. For example:

```
mysql> SELECT DATE_ADD('2006-05-00',INTERVAL 1 DAY);
-> NULL
mysql> SELECT DAYNAME('2006-05-00');
-> NULL
```

- `ADDDATE(date, INTERVAL expr unit), ADDDATE(expr, days)`

When invoked with the `INTERVAL` form of the second argument, `ADDDATE()` is a synonym for `DATE_ADD()`. The related function `SUBDATE()` is a synonym for `DATE_SUB()`. For information on the `INTERVAL unit` argument, see the discussion for `DATE_ADD()`.

```
mysql> SELECT DATE_ADD('1998-01-02', INTERVAL 31 DAY);
-> '1998-02-02'
mysql> SELECT ADDDATE('1998-01-02', INTERVAL 31 DAY);
-> '1998-02-02'
```

As of MySQL 4.1.1, the second syntax is allowed. When invoked with the `days` form of the second argument, MySQL treats it as an integer number of days to be added to `expr`.

```
mysql> SELECT ADDDATE('1998-01-02', 31);
```

```
-> '1998-02-02'
```

- `ADDTIME(expr1,expr2)`

`ADDTIME()` adds *expr2* to *expr1* and returns the result. *expr1* is a time or datetime expression, and *expr2* is a time expression.

```
mysql> SELECT ADDTIME('1997-12-31 23:59:59.999999',
-> '1 1:1:1.000002');
mysql> SELECT ADDTIME('01:00:00.999999', '02:00:00.999998');
-> '03:00:01.999997'
```

`ADDTIME()` was added in MySQL 4.1.1.

- `CONVERT_TZ(dt,from_tz,to_tz)`

`CONVERT_TZ()` converts a datetime value *dt* from the time zone given by *from\_tz* to the time zone given by *to\_tz* and returns the resulting value. Time zones are specified as described in [Section 5.10.8, “MySQL Server Time Zone Support”](#). This function returns `NULL` if the arguments are invalid.

If the value falls out of the supported range of the `TIMESTAMP` type when converted from *from\_tz* to UTC, no conversion occurs. The `TIMESTAMP` range is described in [Section 11.1.2, “Overview of Date and Time Types”](#).

```
mysql> SELECT CONVERT_TZ('2004-01-01 12:00:00','GMT','MET');
-> '2004-01-01 13:00:00'
mysql> SELECT CONVERT_TZ('2004-01-01 12:00:00','+00:00','+10:00');
-> '2004-01-01 22:00:00'
```

**Note:** To use named time zones such as `'MET'` or `'Europe/Moscow'`, the time zone tables must be properly set up. See [Section 5.10.8, “MySQL Server Time Zone Support”](#), for instructions.

`CONVERT_TZ()` was added in MySQL 4.1.3.

If you intend to use `CONVERT_TZ()` while other tables are locked with `LOCK TABLES`, you must also lock the `mysql.time_zone_name` table.

- `CURDATE()`

Returns the current date as a value in `'YYYY-MM-DD'` or `YYYYMMDD` format, depending on whether the function is used in a string or numeric context.

```
mysql> SELECT CURDATE();
-> '1997-12-15'
mysql> SELECT CURDATE() + 0;
-> 19971215
```

- `CURRENT_DATE`, `CURRENT_DATE()`

`CURRENT_DATE` and `CURRENT_DATE()` are synonyms for `CURDATE()`.

- `CURTIME()`

Returns the current time as a value in `'HH:MM:SS'` or `HHMMSS` format, depending on whether the function is used in a string or numeric context.

```
mysql> SELECT CURTIME();
-> '23:50:26'
mysql> SELECT CURTIME() + 0;
```

```
-> 235026
```

- `CURRENT_TIME`, `CURRENT_TIME()`

`CURRENT_TIME` and `CURRENT_TIME()` are synonyms for `CURTIME()`.

- `CURRENT_TIMESTAMP`, `CURRENT_TIMESTAMP()`

`CURRENT_TIMESTAMP` and `CURRENT_TIMESTAMP()` are synonyms for `NOW()`.

- `DATE(expr)`

Extracts the date part of the date or datetime expression `expr`.

```
mysql> SELECT DATE('2003-12-31 01:02:03');
-> '2003-12-31'
```

`DATE()` is available as of MySQL 4.1.1.

- `DATEDIFF(expr1,expr2)`

`DATEDIFF()` returns `expr1 - expr2` expressed as a value in days from one date to the other. `expr1` and `expr2` are date or date-and-time expressions. Only the date parts of the values are used in the calculation.

```
mysql> SELECT DATEDIFF('1997-12-31 23:59:59','1997-12-30');
-> 1
mysql> SELECT DATEDIFF('1997-11-30 23:59:59','1997-12-31');
-> -31
```

`DATEDIFF()` was added in MySQL 4.1.1.

- `DATE_ADD(date,INTERVAL expr unit)`, `DATE_SUB(date,INTERVAL expr unit)`

These functions perform date arithmetic. `date` is a `DATETIME` or `DATE` value specifying the starting date. `expr` is an expression specifying the interval value to be added or subtracted from the starting date. `expr` is a string; it may start with a '-' for negative intervals. `unit` is a keyword indicating the units in which the expression should be interpreted.

The `INTERVAL` keyword and the `unit` specifier are not case sensitive.

The following table shows the expected form of the `expr` argument for each `unit` value.

<code>unit</code> Value	Expected <code>expr</code> Format	Version
<code>MICROSECOND</code>	<code>MICROSECONDS</code>	4.1.1
<code>SECOND</code>	<code>SECONDS</code>	Pre-4.1
<code>MINUTE</code>	<code>MINUTES</code>	Pre-4.1
<code>HOURL</code>	<code>HOURS</code>	Pre-4.1
<code>DAY</code>	<code>DAYS</code>	Pre-4.1
<code>WEEK</code>	<code>WEEKS</code>	5.0.0
<code>MONTH</code>	<code>MONTHS</code>	Pre-4.1
<code>QUARTER</code>	<code>QUARTERS</code>	5.0.0
<code>YEAR</code>	<code>YEARS</code>	Pre-4.1
<code>SECOND_MICROSECOND</code>	' <code>SECONDS.MICROSECONDS</code> '	4.1.1



MINUTE_MICROSECOND	'MINUTES.MICROSECONDS'	4.1.1
MINUTE_SECOND	'MINUTES:SECONDS'	4.1.1
HOURL_MICROSECOND	'HOURS.MICROSECONDS'	4.1.1
HOURL_SECOND	'HOURS:MINUTES:SECONDS'	4.1.1
HOURL_MINUTE	'HOURS:MINUTES'	Pre-4.1
DAY_MICROSECOND	'DAYS.MICROSECONDS'	4.1.1
DAY_SECOND	'DAYS HOURS:MINUTES:SECONDS'	Pre-4.1
DAY_MINUTE	'DAYS HOURS:MINUTES'	Pre-4.1
DAY_HOUR	'DAYS HOURS'	Pre-4.1
YEAR_MONTH	'YEARS-MONTHS'	Pre-4.1

The *type* values DAY\_MICROSECOND, HOURL\_MICROSECOND, MINUTE\_MICROSECOND, SECOND\_MICROSECOND, and MICROSECOND are allowed as of MySQL 4.1.1.

MySQL allows any punctuation delimiter in the *expr* format. Those shown in the table are the suggested delimiters. If the *date* argument is a DATE value and your calculations involve only YEAR, MONTH, and DAY parts (that is, no time parts), the result is a DATE value. Otherwise, the result is a DATETIME value.

As of MySQL 3.23, date arithmetic also can be performed using INTERVAL together with the + or - operator:

```
date + INTERVAL expr unit
date - INTERVAL expr unit
```

INTERVAL *expr* *unit* is allowed on either side of the + operator if the expression on the other side is a date or datetime value. For the - operator, INTERVAL *expr* *unit* is allowed only on the right side, because it makes no sense to subtract a date or datetime value from an interval.

```
mysql> SELECT '1997-12-31 23:59:59' + INTERVAL 1 SECOND;
-> '1998-01-01 00:00:00'
mysql> SELECT INTERVAL 1 DAY + '1997-12-31';
-> '1998-01-01'
mysql> SELECT '1998-01-01' - INTERVAL 1 SECOND;
-> '1997-12-31 23:59:59'
mysql> SELECT DATE_ADD('1997-12-31 23:59:59',
-> INTERVAL 1 SECOND);
-> '1998-01-01 00:00:00'
mysql> SELECT DATE_ADD('1997-12-31 23:59:59',
-> INTERVAL 1 DAY);
-> '1998-01-01 23:59:59'
mysql> SELECT DATE_ADD('1997-12-31 23:59:59',
-> INTERVAL '1:1' MINUTE_SECOND);
-> '1998-01-01 00:01:00'
mysql> SELECT DATE_SUB('1998-01-01 00:00:00',
-> INTERVAL '1 1:1:1' DAY_SECOND);
-> '1997-12-30 22:58:59'
mysql> SELECT DATE_ADD('1998-01-01 00:00:00',
-> INTERVAL '-1 10' DAY_HOUR);
-> '1997-12-30 14:00:00'
mysql> SELECT DATE_SUB('1998-01-02', INTERVAL 31 DAY);
-> '1997-12-02'
mysql> SELECT DATE_ADD('1992-12-31 23:59:59.000002',
-> INTERVAL '1.999999' SECOND_MICROSECOND);
-> '1993-01-01 00:00:01.000001'
```

If you specify an interval value that is too short (does not include all the interval parts that would be expected from the *unit* keyword), MySQL assumes that you have left out the leftmost parts of the

interval value. For example, if you specify a *unit* of `DAY_SECOND`, the value of *expr* is expected to have days, hours, minutes, and seconds parts. If you specify a value like `'1:10'`, MySQL assumes that the days and hours parts are missing and the value represents minutes and seconds. In other words, `'1:10' DAY_SECOND` is interpreted in such a way that it is equivalent to `'1:10' MINUTE_SECOND`. This is analogous to the way that MySQL interprets `TIME` values as representing elapsed time rather than as a time of day.

If you add to or subtract from a date value something that contains a time part, the result is automatically converted to a datetime value:

```
mysql> SELECT DATE_ADD('1999-01-01', INTERVAL 1 DAY);
-> '1999-01-02'
mysql> SELECT DATE_ADD('1999-01-01', INTERVAL 1 HOUR);
-> '1999-01-01 01:00:00'
```

If you add `MONTH`, `YEAR_MONTH`, or `YEAR` and the resulting date has a day that is larger than the maximum day for the new month, the day is adjusted to the maximum days in the new month:

```
mysql> SELECT DATE_ADD('1998-01-30', INTERVAL 1 MONTH);
-> '1998-02-28'
```

Date arithmetic operations require complete dates and do not work with incomplete dates such as `'2006-07-00'` or badly malformed dates:

```
mysql> SELECT DATE_ADD('2006-07-00', INTERVAL 1 DAY);
-> NULL
mysql> SELECT '2005-03-32' + INTERVAL 1 MONTH;
-> NULL
```

- `DATE_FORMAT(date, format)`

Formats the *date* value according to the *format* string.

The following specifiers may be used in the *format* string. As of MySQL 3.23, the `'%'` character is required before format specifier characters. In earlier versions of MySQL, `'%'` was optional.

Specifier	Description
<code>%a</code>	Abbreviated weekday name ( <code>Sun..Sat</code> )
<code>%b</code>	Abbreviated month name ( <code>Jan..Dec</code> )
<code>%c</code>	Month, numeric ( <code>0..12</code> )
<code>%D</code>	Day of the month with English suffix ( <code>0th, 1st, 2nd, 3rd, ...</code> )
<code>%d</code>	Day of the month, numeric ( <code>00..31</code> )
<code>%e</code>	Day of the month, numeric ( <code>0..31</code> )
<code>%f</code>	Microseconds ( <code>000000..999999</code> )
<code>%H</code>	Hour ( <code>00..23</code> )
<code>%h</code>	Hour ( <code>01..12</code> )
<code>%I</code>	Hour ( <code>01..12</code> )
<code>%i</code>	Minutes, numeric ( <code>00..59</code> )
<code>%j</code>	Day of year ( <code>001..366</code> )
<code>%k</code>	Hour ( <code>0..23</code> )
<code>%l</code>	Hour ( <code>1..12</code> )
<code>%M</code>	Month name ( <code>January..December</code> )
<code>%m</code>	Month, numeric ( <code>00..12</code> )

<code>%P</code>	AM or PM
<code>%r</code>	Time, 12-hour ( <code>hh:mm:ss</code> followed by AM or PM)
<code>%S</code>	Seconds (00..59)
<code>%s</code>	Seconds (00..59)
<code>%T</code>	Time, 24-hour ( <code>hh:mm:ss</code> )
<code>%U</code>	Week (00..53), where Sunday is the first day of the week
<code>%u</code>	Week (00..53), where Monday is the first day of the week
<code>%V</code>	Week (01..53), where Sunday is the first day of the week; used with <code>%X</code>
<code>%v</code>	Week (01..53), where Monday is the first day of the week; used with <code>%x</code>
<code>%W</code>	Weekday name (Sunday..Saturday)
<code>%w</code>	Day of the week (0=Sunday..6=Saturday)
<code>%X</code>	Year for the week where Sunday is the first day of the week, numeric, four digits; used with <code>%V</code>
<code>%x</code>	Year for the week, where Monday is the first day of the week, numeric, four digits; used with <code>%v</code>
<code>%Y</code>	Year, numeric, four digits
<code>%y</code>	Year, numeric (two digits)
<code>%%</code>	A literal '%' character
<code>%x</code>	<code>x</code> , for any 'x' not listed above

The `%v`, `%V`, `%x`, and `%X` format specifiers are available as of MySQL 3.23.8. `%f` is available as of MySQL 4.1.1.

Ranges for the month and day specifiers begin with zero due to the fact that MySQL allows the storing of incomplete dates such as '2004-00-00' (as of MySQL 3.23).

```
mysql> SELECT DATE_FORMAT('1997-10-04 22:23:00', '%W %M %Y');
-> 'Saturday October 1997'
mysql> SELECT DATE_FORMAT('1997-10-04 22:23:00', '%H:%i:%s');
-> '22:23:00'
mysql> SELECT DATE_FORMAT('1997-10-04 22:23:00',
    '%D %y %a %d %m %b %j');
-> '4th 97 Sat 04 10 Oct 277'
mysql> SELECT DATE_FORMAT('1997-10-04 22:23:00',
    '%H %k %I %r %T %S %w');
-> '22 22 10 10:23:00 PM 22:23:00 00 6'
mysql> SELECT DATE_FORMAT('1999-01-01', '%X %V');
-> '1998 52'
mysql> SELECT DATE_FORMAT('2006-06-00', '%d');
-> '00'
```

- `DAY(date)`

`DAY()` is a synonym for `DAYOFMONTH()`. It is available as of MySQL 4.1.1.

- `DAYNAME(date)`

Returns the name of the weekday for `date`.

```
mysql> SELECT DAYNAME('1998-02-05');
-> 'Thursday'
```

- `DAYOFMONTH(date)`

Returns the day of the month for *date*, in the range 0 to 31.

```
mysql> SELECT DAYOFMONTH('1998-02-03');
-> 3
```

- `DAYOFWEEK(date)`

Returns the weekday index for *date* (1 = Sunday, 2 = Monday, ..., 7 = Saturday). These index values correspond to the ODBC standard.

```
mysql> SELECT DAYOFWEEK('1998-02-03');
-> 3
```

- `DAYOFYEAR(date)`

Returns the day of the year for *date*, in the range 1 to 366.

```
mysql> SELECT DAYOFYEAR('1998-02-03');
-> 34
```

- `EXTRACT(unit FROM date)`

The `EXTRACT()` function uses the same kinds of unit specifiers as `DATE_ADD()` or `DATE_SUB()`, but extracts parts from the date rather than performing date arithmetic.

```
mysql> SELECT EXTRACT(YEAR FROM '1999-07-02');
-> 1999
mysql> SELECT EXTRACT(YEAR_MONTH FROM '1999-07-02 01:02:03');
-> 199907
mysql> SELECT EXTRACT(DAY_MINUTE FROM '1999-07-02 01:02:03');
-> 20102
mysql> SELECT EXTRACT(MICROSECOND
-> FROM '2003-01-02 10:30:00.00123');
-> 123
```

`EXTRACT()` was added in MySQL 3.23.0.

- `FROM_DAYS(N)`

Given a day number *N*, returns a `DATE` value.

```
mysql> SELECT FROM_DAYS(729669);
-> '1997-10-07'
```

Use `FROM_DAYS()` with caution on old dates. It is not intended for use with values that precede the advent of the Gregorian calendar (1582). See [Section 12.6, “What Calendar Is Used By MySQL?”](#).

- `FROM_UNIXTIME(unix_timestamp)`, `FROM_UNIXTIME(unix_timestamp, format)`

Returns a representation of the *unix\_timestamp* argument as a value in 'YYYY-MM-DD HH:MM:SS' or YYYYMMDDHHMMSS format, depending on whether the function is used in a string or numeric context. *unix\_timestamp* is an internal timestamp value such as is produced by the `UNIX_TIMESTAMP()` function.

If *format* is given, the result is formatted according to the *format* string, which is used the same way as listed in the entry for the `DATE_FORMAT()` function.

```
mysql> SELECT FROM_UNIXTIME(875996580);
-> '1997-10-04 22:23:00'
mysql> SELECT FROM_UNIXTIME(875996580) + 0;
-> 19971004222300
```

```
mysql> SELECT FROM_UNIXTIME(UNIX_TIMESTAMP(),
->    '%Y %D %M %h:%i:%s %x');
-> '2003 6th August 06:22:58 2003'
```

Note: If you use `UNIX_TIMESTAMP()` and `FROM_UNIXTIME()` to convert between `TIMESTAMP` values and Unix timestamp values, the conversion is lossy because the mapping is not one-to-one in both directions. For details, see the description of the `UNIX_TIMESTAMP()` function.

- `GET_FORMAT( DATE | TIME | DATETIME , 'EUR' | 'USA' | 'JIS' | 'ISO' | 'INTERNAL' )`

Returns a format string. This function is useful in combination with the `DATE_FORMAT()` and the `STR_TO_DATE()` functions.

The possible values for the first and second arguments result in several possible format strings (for the specifiers used, see the table in the `DATE_FORMAT()` function description). ISO format refers to ISO 9075, not ISO 8601.

Function Call	Result
<code>GET_FORMAT( DATE , 'USA' )</code>	<code>'%m.%d.%Y'</code>
<code>GET_FORMAT( DATE , 'JIS' )</code>	<code>'%Y-%m-%d'</code>
<code>GET_FORMAT( DATE , 'ISO' )</code>	<code>'%Y-%m-%d'</code>
<code>GET_FORMAT( DATE , 'EUR' )</code>	<code>'%d.%m.%Y'</code>
<code>GET_FORMAT( DATE , 'INTERNAL' )</code>	<code>'%Y%m%d'</code>
<code>GET_FORMAT( DATETIME , 'USA' )</code>	<code>'%Y-%m-%d %H.%i.%s'</code>
<code>GET_FORMAT( DATETIME , 'JIS' )</code>	<code>'%Y-%m-%d %H:%i:%s'</code>
<code>GET_FORMAT( DATETIME , 'ISO' )</code>	<code>'%Y-%m-%d %H:%i:%s'</code>
<code>GET_FORMAT( DATETIME , 'EUR' )</code>	<code>'%Y-%m-%d %H.%i.%s'</code>
<code>GET_FORMAT( DATETIME , 'INTERNAL' )</code>	<code>'%Y%m%d%H%i%s'</code>
<code>GET_FORMAT( TIME , 'USA' )</code>	<code>'%h:%i:%s %p'</code>
<code>GET_FORMAT( TIME , 'JIS' )</code>	<code>'%H:%i:%s'</code>
<code>GET_FORMAT( TIME , 'ISO' )</code>	<code>'%H:%i:%s'</code>
<code>GET_FORMAT( TIME , 'EUR' )</code>	<code>'%H.%i.%s'</code>
<code>GET_FORMAT( TIME , 'INTERNAL' )</code>	<code>'%H%i%s'</code>

As of MySQL 4.1.4, `TIMESTAMP` can also be used as the first argument to `GET_FORMAT()`, in which case the function returns the same values as for `DATETIME`.

```
mysql> SELECT DATE_FORMAT('2003-10-03',GET_FORMAT(DATE,'EUR'));
-> '03.10.2003'
mysql> SELECT STR_TO_DATE('10.31.2003',GET_FORMAT(DATE,'USA'));
-> '2003-10-31'
```

`GET_FORMAT()` is available as of MySQL 4.1.1.

- `HOUR( time )`

Returns the hour for `time`. The range of the return value is 0 to 23 for time-of-day values. However, the range of `TIME` values actually is much larger, so `HOUR` can return values greater than 23.

```
mysql> SELECT HOUR('10:05:03');
```

```
mysql> SELECT HOUR('272:59:59');
-> 10
-> 272
```

- `LAST_DAY(date)`

Takes a date or datetime value and returns the corresponding value for the last day of the month. Returns `NULL` if the argument is invalid.

```
mysql> SELECT LAST_DAY('2003-02-05');
-> '2003-02-28'
mysql> SELECT LAST_DAY('2004-02-05');
-> '2004-02-29'
mysql> SELECT LAST_DAY('2004-01-01 01:01:01');
-> '2004-01-31'
mysql> SELECT LAST_DAY('2003-03-32');
-> NULL
```

`LAST_DAY()` is available as of MySQL 4.1.1.

- `LOCALTIME, LOCALTIME()`

`LOCALTIME` and `LOCALTIME()` are synonyms for `NOW()`.

- `LOCALTIMESTAMP, LOCALTIMESTAMP()`

`LOCALTIMESTAMP` and `LOCALTIMESTAMP()` are synonyms for `NOW()`.

They were added in MySQL 4.0.6.

- `MAKEDATE(year, dayofyear)`

Returns a date, given year and day-of-year values. `dayofyear` must be greater than 0 or the result is `NULL`.

```
mysql> SELECT MAKEDATE(2001,31), MAKEDATE(2001,32);
-> '2001-01-31', '2001-02-01'
mysql> SELECT MAKEDATE(2001,365), MAKEDATE(2004,365);
-> '2001-12-31', '2004-12-30'
mysql> SELECT MAKEDATE(2001,0);
-> NULL
```

`MAKEDATE()` is available as of MySQL 4.1.1.

- `MAKETIME(hour, minute, second)`

Returns a time value calculated from the `hour`, `minute`, and `second` arguments.

```
mysql> SELECT MAKETIME(12,15,30);
-> '12:15:30'
```

`MAKETIME()` is available as of MySQL 4.1.1.

- `MICROSECOND(expr)`

Returns the microseconds from the time or datetime expression `expr` as a number in the range from 0 to 999999.

```
mysql> SELECT MICROSECOND('12:00:00.123456');
-> 123456
mysql> SELECT MICROSECOND('1997-12-31 23:59:59.000010');
-> 10
```

`MICROSECOND()` is available as of MySQL 4.1.1.

- `MINUTE(time)`

Returns the minute for *time*, in the range 0 to 59.

```
mysql> SELECT MINUTE('98-02-03 10:05:03');
-> 5
```

- `MONTH(date)`

Returns the month for *date*, in the range 0 to 12.

```
mysql> SELECT MONTH('1998-02-03');
-> 2
```

- `MONTHNAME(date)`

Returns the full name of the month for *date*.

```
mysql> SELECT MONTHNAME('1998-02-05');
-> 'February'
```

- `NOW()`

Returns the current date and time as a value in 'YYYY-MM-DD HH:MM:SS' or YYYYMMDDHHMMSS format, depending on whether the function is used in a string or numeric context.

```
mysql> SELECT NOW();
-> '1997-12-15 23:50:26'
mysql> SELECT NOW() + 0;
-> 19971215235026
```

- `PERIOD_ADD(P,N)`

Adds *N* months to period *P* (in the format `YYMM` or `YYYYMM`). Returns a value in the format `YYMM`. Note that the period argument *P* is *not* a date value.

```
mysql> SELECT PERIOD_ADD(9801,2);
-> 199803
```

- `PERIOD_DIFF(P1,P2)`

Returns the number of months between periods *P1* and *P2*. *P1* and *P2* should be in the format `YYMM` or `YYYYMM`. Note that the period arguments *P1* and *P2* are *not* date values.

```
mysql> SELECT PERIOD_DIFF(9802,199703);
-> 11
```

- `QUARTER(date)`

Returns the quarter of the year for *date*, in the range 1 to 4.

```
mysql> SELECT QUARTER('98-04-01');
-> 2
```

- `SECOND(time)`

Returns the second for *time*, in the range 0 to 59.

```
mysql> SELECT SECOND('10:05:03');
-> 3
```

- `SEC_TO_TIME(seconds)`

Returns the *seconds* argument, converted to hours, minutes, and seconds, as a value in 'HH:MM:SS' or HHMMSS format, depending on whether the function is used in a string or numeric context.

```
mysql> SELECT SEC_TO_TIME(2378);
-> '00:39:38'
mysql> SELECT SEC_TO_TIME(2378) + 0;
-> 3938
```

**Note:** You cannot use format "%X%V" to convert a year-week string to a date because the combination of a year and week does not uniquely identify a year and month if the week crosses a month boundary. To convert a year-week to a date, then you should also specify the weekday:

```
mysql> SELECT STR_TO_DATE('200442 Monday', '%X%V %W');
-> '2004-10-18'
```

- `STR_TO_DATE(str, format)`

This is the inverse of the `DATE_FORMAT()` function. It takes a string *str* and a format string *format*. `STR_TO_DATE()` returns a `DATETIME` value if the format string contains both date and time parts, or a `DATE` or `TIME` value if the string contains only date or time parts.

The date, time, or datetime values contained in *str* should be given in the format indicated by *format*. For the specifiers that can be used in *format*, see the `DATE_FORMAT()` function description. If *str* contains an illegal date, time, or datetime value, `STR_TO_DATE()` returns `NULL`.

Range checking on the parts of date values is as described in [Section 11.3.1, “The DATETIME, DATE, and TIMESTAMP Types”](#). This means, for example, that a date with a day part larger than the number of days in a month is allowable as long as the day part is in the range from 1 to 31. Also, “zero” dates or dates with part values of 0 are allowed.

```
mysql> SELECT STR_TO_DATE('00/00/0000', '%m/%d/%Y');
-> '0000-00-00'
mysql> SELECT STR_TO_DATE('04/31/2004', '%m/%d/%Y');
-> '2004-04-31'
```

`STR_TO_DATE()` is available as of MySQL 4.1.1.

- `SUBDATE(date, INTERVAL expr unit), SUBDATE(expr, days)`

When invoked with the `INTERVAL` form of the second argument, `SUBDATE()` is a synonym for `DATE_SUB()`. For information on the `INTERVAL unit` argument, see the discussion for `DATE_ADD()`.

```
mysql> SELECT DATE_SUB('1998-01-02', INTERVAL 31 DAY);
-> '1997-12-02'
mysql> SELECT SUBDATE('1998-01-02', INTERVAL 31 DAY);
-> '1997-12-02'
```

As of MySQL 4.1.1, the second syntax is allowed, where *expr* is a date or datetime expression and *days* is the number of days to be subtracted from *expr*.



```
mysql> SELECT SUBDATE('1998-01-02 12:00:00', 31);
-> '1997-12-02 12:00:00'
```

- `SUBTIME(expr1, expr2)`

`SUBTIME()` returns *expr1* – *expr2* expressed as a value in the same format as *expr1*. *expr1* is a time or datetime expression, and *expr2* is a time expression.

```
mysql> SELECT SUBTIME('1997-12-31 23:59:59.999999', '1 1:1:1.000002');
-> '1997-12-30 22:58:58.999997'
mysql> SELECT SUBTIME('01:00:00.999999', '02:00:00.999998');
-> '-00:59:59.999999'
```

`SUBTIME()` was added in MySQL 4.1.1.

- `SYSDATE()`

Returns the current date and time as a value in 'YYYY-MM-DD HH:MM:SS' or YYYYMMDDH-HMMSS format, depending on whether the function is used in a string or numeric context.

- `TIME(expr)`

Extracts the time part of the time or datetime expression *expr* and returns it as a string.

```
mysql> SELECT TIME('2003-12-31 01:02:03');
-> '01:02:03'
mysql> SELECT TIME('2003-12-31 01:02:03.000123');
-> '01:02:03.000123'
```

`TIME()` is available as of MySQL 4.1.1.

- `TIMEDIFF(expr1, expr2)`

`TIMEDIFF()` returns *expr1* – *expr2* expressed as a time value. *expr1* and *expr2* are time or date-and-time expressions, but both must be of the same type.

```
mysql> SELECT TIMEDIFF('2000:01:01 00:00:00',
-> '2000:01:01 00:00:00.000001');
-> '-00:00:00.000001'
mysql> SELECT TIMEDIFF('1997-12-31 23:59:59.000001',
-> '1997-12-30 01:01:01.000002');
-> '46:58:57.999999'
```

`TIMEDIFF()` was added in MySQL 4.1.1.

- `TIMESTAMP(expr)`, `TIMESTAMP(expr1, expr2)`

With a single argument, this function returns the date or datetime expression *expr* as a datetime value. With two arguments, it adds the time expression *expr2* to the date or datetime expression *expr1* and returns the result as a datetime value.

```
mysql> SELECT TIMESTAMP('2003-12-31');
-> '2003-12-31 00:00:00'
mysql> SELECT TIMESTAMP('2003-12-31 12:00:00', '12:00:00');
-> '2004-01-01 00:00:00'
```

`TIMESTAMP()` is available as of MySQL 4.1.1.

- `TIME_FORMAT(time, format)`

This is used like the `DATE_FORMAT()` function, but the *format* string may contain format specifiers only for hours, minutes, and seconds. Other specifiers produce a `NULL` value or `0`.

If the *time* value contains an hour part that is greater than `23`, the `%H` and `%k` hour format specifiers produce a value larger than the usual range of `0..23`. The other hour format specifiers produce the hour value modulo 12.

```
mysql> SELECT TIME_FORMAT('100:00:00', '%H %k %h %I %l');
-> '100 100 04 04 4'
```

- `TIME_TO_SEC(time)`

Returns the *time* argument, converted to seconds.

```
mysql> SELECT TIME_TO_SEC('22:23:00');
-> 80580
mysql> SELECT TIME_TO_SEC('00:39:38');
-> 2378
```

- `TO_DAYS(date)`

Given a date *date*, returns a day number (the number of days since year 0).

```
mysql> SELECT TO_DAYS(950501);
-> 728779
mysql> SELECT TO_DAYS('1997-10-07');
-> 729669
```

`TO_DAYS()` is not intended for use with values that precede the advent of the Gregorian calendar (1582), because it does not take into account the days that were lost when the calendar was changed. For dates before 1582 (and possibly a later year in other locales), results from this function are not reliable. See [Section 12.6, “What Calendar Is Used By MySQL?”](#), for details.

Remember that MySQL converts two-digit year values in dates to four-digit form using the rules in [Section 11.3, “Date and Time Types”](#). For example, `'1997-10-07'` and `'97-10-07'` are seen as identical dates:

```
mysql> SELECT TO_DAYS('1997-10-07'), TO_DAYS('97-10-07');
-> 729669, 729669
```

- `UNIX_TIMESTAMP()`, `UNIX_TIMESTAMP(date)`

If called with no argument, returns a Unix timestamp (seconds since `'1970-01-01 00:00:00'` UTC) as an unsigned integer. If `UNIX_TIMESTAMP()` is called with a *date* argument, it returns the value of the argument as seconds since `'1970-01-01 00:00:00'` UTC. *date* may be a `DATE` string, a `DATETIME` string, a `TIMESTAMP`, or a number in the format `YYMMDD` or `YYYYMMDD`. The server interprets *date* as a value in the current time zone and converts it to an internal value in UTC. Clients can set their time zone as described in [Section 5.10.8, “MySQL Server Time Zone Support”](#).

```
mysql> SELECT UNIX_TIMESTAMP();
-> 882226357
mysql> SELECT UNIX_TIMESTAMP('1997-10-04 22:23:00');
-> 875996580
```

When `UNIX_TIMESTAMP` is used on a `TIMESTAMP` column, the function returns the internal timestamp value directly, with no implicit “string-to-Unix-timestamp” conversion. If you pass an out-of-range date to `UNIX_TIMESTAMP()`, it returns `0`, but please note that only basic range checking is performed (year from `1970` to `2037`, month from `01` to `12`, day from `01` from `31`).

Note: If you use `UNIX_TIMESTAMP()` and `FROM_UNIXTIME()` to convert between `TIMESTAMP` values and Unix timestamp values, the conversion is lossy because the mapping is not one-to-one in both directions. For example, due to conventions for local time zone changes, it is possible for two `UNIX_TIMESTAMP()` to map two `TIMESTAMP` values to the same Unix timestamp value. `FROM_UNIXTIME()` will map that value back to only one of the original `TIMESTAMP` values. Here is an example, using `TIMESTAMP` values in the `CET` time zone:

```
mysql> SELECT UNIX_TIMESTAMP('2005-03-27 03:00:00');
+-----+
| UNIX_TIMESTAMP('2005-03-27 03:00:00') |
+-----+
| 1111885200 |
+-----+
mysql> SELECT UNIX_TIMESTAMP('2005-03-27 02:00:00');
+-----+
| UNIX_TIMESTAMP('2005-03-27 02:00:00') |
+-----+
| 1111885200 |
+-----+
mysql> SELECT FROM_UNIXTIME(1111885200);
+-----+
| FROM_UNIXTIME(1111885200) |
+-----+
| 2005-03-27 03:00:00 |
+-----+
```

If you want to subtract `UNIX_TIMESTAMP()` columns, you might want to cast the result to signed integers. See [Section 12.8, “Cast Functions and Operators”](#).

- `UTC_DATE`, `UTC_DATE()`

Returns the current UTC date as a value in `'YYYY-MM-DD'` or `YYYYMMDD` format, depending on whether the function is used in a string or numeric context.

```
mysql> SELECT UTC_DATE(), UTC_DATE() + 0;
-> '2003-08-14', 20030814
```

`UTC_DATE()` is available as of MySQL 4.1.1.

- `UTC_TIME`, `UTC_TIME()`

Returns the current UTC time as a value in `'HH:MM:SS'` or `HHMMSS` format, depending on whether the function is used in a string or numeric context.

```
mysql> SELECT UTC_TIME(), UTC_TIME() + 0;
-> '18:07:53', 180753
```

`UTC_TIME()` is available as of MySQL 4.1.1.

- `UTC_TIMESTAMP`, `UTC_TIMESTAMP()`

Returns the current UTC date and time as a value in `'YYYY-MM-DD HH:MM:SS'` or `YYYYMM-DDHHMMSS` format, depending on whether the function is used in a string or numeric context.

```
mysql> SELECT UTC_TIMESTAMP(), UTC_TIMESTAMP() + 0;
-> '2003-08-14 18:08:04', 20030814180804
```

`UTC_TIMESTAMP()` is available as of MySQL 4.1.1.

- `WEEK(date[, mode])`

This function returns the week number for `date`. The two-argument form of `WEEK()` allows you to

specify whether the week starts on Sunday or Monday and whether the return value should be in the range from 0 to 53 or from 1 to 53. If the *mode* argument is omitted, the value of the `default_week_format` system variable is used (or 0 before MySQL 4.0.14). See [Section 5.2.2, “System Variables”](#).

The following table describes how the *mode* argument works.

Mode	First day of week	Range	Week 1 is the first week ...
0	Sunday	0-53	with a Sunday in this year
1	Monday	0-53	with more than 3 days this year
2	Sunday	1-53	with a Sunday in this year
3	Monday	1-53	with more than 3 days this year
4	Sunday	0-53	with more than 3 days this year
5	Monday	0-53	with a Monday in this year
6	Sunday	1-53	with more than 3 days this year
7	Monday	1-53	with a Monday in this year

A *mode* value of 3 can be used as of MySQL 4.0.5. Values of 4 and above can be used as of MySQL 4.0.17.

```
mysql> SELECT WEEK('1998-02-20');
-> 7
mysql> SELECT WEEK('1998-02-20',0);
-> 7
mysql> SELECT WEEK('1998-02-20',1);
-> 8
mysql> SELECT WEEK('1998-12-31',1);
-> 53
```

Note: In MySQL 4.0, `WEEK(date,0)` was changed to match the calendar in the USA. Before that, `WEEK()` was calculated incorrectly for dates in the USA. (In effect, `WEEK(date)` and `WEEK(date,0)` were incorrect for all cases.)

Note that if a date falls in the last week of the previous year, MySQL returns 0 if you do not use 2, 3, 6, or 7 as the optional *mode* argument:

```
mysql> SELECT YEAR('2000-01-01'), WEEK('2000-01-01',0);
-> 2000, 0
```

One might argue that MySQL should return 52 for the `WEEK()` function, because the given date actually occurs in the 52nd week of 1999. We decided to return 0 instead because we want the function to return “the week number in the given year.” This makes use of the `WEEK()` function reliable when combined with other functions that extract a date part from a date.

If you would prefer the result to be evaluated with respect to the year that contains the first day of the week for the given date, use 0, 2, 5, or 7 as the optional *mode* argument.

```
mysql> SELECT WEEK('2000-01-01',2);
-> 52
```

Alternatively, use the `YEARWEEK()` function:

```
mysql> SELECT YEARWEEK('2000-01-01');
-> 199952
mysql> SELECT MID(YEARWEEK('2000-01-01'),5,2);
```

```
-> '52'
```

- `WEEKDAY(date)`

Returns the weekday index for *date* (0 = Monday, 1 = Tuesday, ... 6 = Sunday).

```
mysql> SELECT WEEKDAY('1998-02-03 22:23:00');
-> 1
mysql> SELECT WEEKDAY('1997-11-05');
-> 2
```

- `WEEKOFYEAR(date)`

Returns the calendar week of the date as a number in the range from 1 to 53. `WEEKOFYEAR()` is a compatibility function that is equivalent to `WEEK(date, 3)`.

```
mysql> SELECT WEEKOFYEAR('1998-02-20');
-> 8
```

`WEEKOFYEAR()` is available as of MySQL 4.1.1.

- `YEAR(date)`

Returns the year for *date*, in the range 1000 to 9999, or 0 for the “zero” date.

```
mysql> SELECT YEAR('98-02-03');
-> 1998
```

- `YEARWEEK(date)`, `YEARWEEK(date, start)`

Returns year and week for a date. The *start* argument works exactly like the *start* argument to `WEEK()`. The year in the result may be different from the year in the date argument for the first and the last week of the year.

```
mysql> SELECT YEARWEEK('1987-01-01');
-> 198653
```

Note that the week number is different from what the `WEEK()` function would return (0) for optional arguments 0 or 1, as `WEEK()` then returns the week in the context of the given year.

`YEARWEEK()` was added in MySQL 3.23.8.

## 12.6. What Calendar Is Used By MySQL?

MySQL uses what is known as a *proleptic Gregorian calendar*.

Every country that has switched from the Julian to the Gregorian calendar has had to discard at least ten days during the switch. To see how this works, consider the month of October 1582, when the first Julian-to-Gregorian switch occurred:

Monday	Tuesday	Wednesday	Thursday	Friday	Saturday	Sunday
1	2	3	4	15	16	17
18	19	20	21	22	23	24
25	26	27	28	29	30	31

There are no dates between October 4 and October 15. This discontinuity is called the *cutover*. Any dates before the cutover are Julian, and any dates following the cutover are Gregorian. Dates during a cutover are non-existent.

A calendar applied to dates when it wasn't actually in use is called *proleptic*. Thus, if we assume there was never a cutover and Gregorian rules always rule, we have a proleptic Gregorian calendar. This is what is used by MySQL, as is required by standard SQL. For this reason, dates prior to the cutover stored as MySQL `DATE` or `DATETIME` values must be adjusted to compensate for the difference. It is important to realize that the cutover did not occur at the same time in all countries, and that the later it happened, the more days were lost. For example, in Great Britain, it took place in 1752, when Wednesday September 2 was followed by Thursday September 14. Russia remained on the Julian calendar until 1918, losing 13 days in the process, and what is popularly referred to as its “October Revolution” occurred in November according to the Gregorian calendar.

## 12.7. Full-Text Search Functions

```
MATCH (col1,col2,...) AGAINST (expr [search_modifier])
search_modifier: { IN BOOLEAN MODE | WITH QUERY EXPANSION }
```

As of MySQL 3.23.23, MySQL has support for full-text indexing and searching:

- A full-text index in MySQL is an index of type `FULLTEXT`.
- Full-text indexes can be used only with `MyISAM` tables, and can be created only for `CHAR`, `VARCHAR`, or `TEXT` columns.
- A `FULLTEXT` index definition can be given in the `CREATE TABLE` statement when a table is created, or added later using `ALTER TABLE` or `CREATE INDEX`.
- For large datasets, it is much faster to load your data into a table that has no `FULLTEXT` index and then create the index after that, than to load data into a table that has an existing `FULLTEXT` index.

Full-text searching is performed using `MATCH() ... AGAINST` syntax. `MATCH()` takes a comma-separated list that names the columns to be searched. `AGAINST` takes a string to search for, and an optional modifier that indicates what type of search to perform. The search string must be a literal string, not a variable or a column name. There are three types of full-text searches:

- A boolean search interprets the search string using the rules of a special query language. The string contains the words to search for. It can also contain operators that specify requirements such that a word must be present or absent in matching rows, or that it should be weighted higher or lower than usual. Common words such as “some” or “then” are stopwords and do not match if present in the search string. The `IN BOOLEAN MODE` modifier specifies a boolean search. For more information, see [Section 12.7.1, “Boolean Full-Text Searches”](#).
- A natural language search interprets the search string as a phrase in natural human language (a phrase in free text). There are no special operators. The stopword list applies. In addition, words that are present in more than 50% of the rows are considered common and do not match. Full-text searches are natural language searches if no modifier is given.
- A query expansion search is a modification of a natural language search. The search string is used to perform a natural language search. Then words from the most relevant rows returned by the search are added to the search string and the search is done again. The query returns the rows from the second search. The `WITH QUERY EXPANSION` modifier specifies a query expansion search. For more information, see [Section 12.7.2, “Full-Text Searches with Query Expansion”](#).

Constraints on full-text searching are listed in [Section 12.7.4, “Full-Text Restrictions”](#).

```
mysql> CREATE TABLE articles (
->   id INT UNSIGNED AUTO_INCREMENT NOT NULL PRIMARY KEY,
->   title VARCHAR(200),
->   body TEXT,
->   FULLTEXT (title,body)
-> );
Query OK, 0 rows affected (0.00 sec)

mysql> INSERT INTO articles (title,body) VALUES
-> ('MySQL Tutorial','DBMS stands for DataBase ...'),
-> ('How To Use MySQL Well','After you went through a ...'),
-> ('Optimizing MySQL','In this tutorial we will show ...'),
-> ('1001 MySQL Tricks','1. Never run mysqld as root. 2. ...'),
-> ('MySQL vs. YourSQL','In the following database comparison ...'),
-> ('MySQL Security','When configured properly, MySQL ...');
Query OK, 6 rows affected (0.00 sec)
Records: 6 Duplicates: 0 Warnings: 0

mysql> SELECT * FROM articles
-> WHERE MATCH (title,body) AGAINST ('database');
+-----+-----+-----+
| id | title | body |
+-----+-----+-----+
| 5 | MySQL vs. YourSQL | In the following database comparison ... |
| 1 | MySQL Tutorial | DBMS stands for DataBase ... |
+-----+-----+-----+
2 rows in set (0.00 sec)
```

The `MATCH()` function performs a natural language search for a string against a *text collection*. A collection is a set of one or more columns included in a `FULLTEXT` index. The search string is given as the argument to `AGAINST()`. For each row in the table, `MATCH()` returns a relevance value; that is, a similarity measure between the search string and the text in that row in the columns named in the `MATCH()` list.

By default, the search is performed in case-insensitive fashion. In MySQL 4.1 and up, you can make a full-text search by using a binary collation for the indexed columns. For example, a column that has a character set of `latin1` character set of can be assigned a collation of `latin1_bin` to make it case sensitive for full-text searches.

When `MATCH()` is used in a `WHERE` clause, as in the example shown earlier, the rows returned are automatically sorted with the highest relevance first. Relevance values are non-negative floating-point numbers. Zero relevance means no similarity. Relevance is computed based on the number of words in the row, the number of unique words in that row, the total number of words in the collection, and the number of documents (rows) that contain a particular word.

For natural-language full-text searches, it is a requirement that the columns named in the `MATCH()` function be the same columns included in some `FULLTEXT` index in your table. For the preceding query, note that the columns named in the `MATCH()` function (`title` and `body`) are the same as those named in the definition of the `article` table's `FULLTEXT` index. If you wanted to search the `title` or `body` separately, you would need to create separate `FULLTEXT` indexes for each column.

It is also possible to perform a boolean search or a search with query expansion. These search types are described in [Section 12.7.1, “Boolean Full-Text Searches”](#), and [Section 12.7.2, “Full-Text Searches with Query Expansion”](#).

The preceding example is a basic illustration that shows how to use the `MATCH()` function where rows are returned in order of decreasing relevance. The next example shows how to retrieve the relevance values explicitly. Returned rows are not ordered because the `SELECT` statement includes neither `WHERE` nor `ORDER BY` clauses:

```
mysql> SELECT id, MATCH (title,body) AGAINST ('Tutorial')
-> FROM articles;
+-----+-----+
| id | MATCH (title,body) AGAINST ('Tutorial') |
+-----+-----+
```

1	0.65545833110809	0
2		0
3	0.66266459226608	0
4		0
5		0
6		0

6 rows in set (0.00 sec)

The following example is more complex. The query returns the relevance values and it also sorts the rows in order of decreasing relevance. To achieve this result, you should specify `MATCH()` twice: once in the `SELECT` list and once in the `WHERE` clause. This causes no additional overhead, because the MySQL optimizer notices that the two `MATCH()` calls are identical and invokes the full-text search code only once.

```
mysql> SELECT id, body, MATCH (title,body) AGAINST
-> ('Security implications of running MySQL as root') AS score
-> FROM articles WHERE MATCH (title,body) AGAINST
-> ('Security implications of running MySQL as root');
```

id	body	score
4	1. Never run mysqld as root. 2. ...	1.5219271183014
6	When configured properly, MySQL ...	1.3114095926285

2 rows in set (0.00 sec)

The MySQL `FULLTEXT` implementation regards any sequence of true word characters (letters, digits, and underscores) as a word. That sequence may also contain apostrophes (''), but not more than one in a row. This means that `aaa'bbb` is regarded as one word, but `aaa' 'bbb` is regarded as two words. Apostrophes at the beginning or the end of a word are stripped by the `FULLTEXT` parser; `'aaa'bbb'` would be parsed as `aaa'bbb`.

The `FULLTEXT` parser determines where words start and end by looking for certain delimiter characters; for example, ' ' (space), ',' (comma), and '.' (period). If words are not separated by delimiters (as in, for example, Chinese), the `FULLTEXT` parser cannot determine where a word begins or ends. To be able to add words or other indexed terms in such languages to a `FULLTEXT` index, you must preprocess them so that they are separated by some arbitrary delimiter such as '" '.

Some words are ignored in full-text searches:

- Any word that is too short is ignored. The default minimum length of words that are found by full-text searches is four characters.
- Words in the stopword list are ignored. A stopword is a word such as “the” or “some” that is so common that it is considered to have zero semantic value. There is a built-in stopword list, but it can be overwritten by a user-defined list.

The default stopword list is given in [Section 12.7.3, “Full-Text Stopwords”](#). The default minimum word length and stopword list can be changed as described in [Section 12.7.5, “Fine-Tuning MySQL Full-Text Search”](#).

Every correct word in the collection and in the query is weighted according to its significance in the collection or query. Consequently, a word that is present in many documents has a lower weight (and may even have a zero weight), because it has lower semantic value in this particular collection. Conversely, if the word is rare, it receives a higher weight. The weights of the words are combined to compute the relevance of the row.

Such a technique works best with large collections (in fact, it was carefully tuned this way). For very small tables, word distribution does not adequately reflect their semantic value, and this model may sometimes produce bizarre results. For example, although the word “MySQL” is present in every row of



the `articles` table shown earlier, a search for the word produces no results:

```
mysql> SELECT * FROM articles
-> WHERE MATCH (title,body) AGAINST ('MySQL');
Empty set (0.00 sec)
```

The search result is empty because the word “MySQL” is present in at least 50% of the rows. As such, it is effectively treated as a stopword. For large datasets, this is the most desirable behavior: A natural language query should not return every second row from a 1GB table. For small datasets, it may be less desirable.

A word that matches half of the rows in a table is less likely to locate relevant documents. In fact, it most likely finds plenty of irrelevant documents. We all know this happens far too often when we are trying to find something on the Internet with a search engine. It is with this reasoning that rows containing the word are assigned a low semantic value for *the particular dataset in which they occur*. A given word may exceed the 50% threshold in one dataset but not another.

The 50% threshold has a significant implication when you first try full-text searching to see how it works: If you create a table and insert only one or two rows of text into it, every word in the text occurs in at least 50% of the rows. As a result, no search returns any results. Be sure to insert at least three rows, and preferably many more. Users who need to bypass the 50% limitation can use the boolean search mode; see [Section 12.7.1, “Boolean Full-Text Searches”](#).

## 12.7.1. Boolean Full-Text Searches

As of version 4.0.1, MySQL can perform boolean full-text searches using the `IN BOOLEAN MODE` modifier:

```
mysql> SELECT * FROM articles WHERE MATCH (title,body)
-> AGAINST ('+MySQL -YourSQL' IN BOOLEAN MODE);
+-----+-----+-----+
| id | title | body |
+-----+-----+-----+
| 1 | MySQL Tutorial | DBMS stands for DataBase ... |
| 2 | How To Use MySQL Well | After you went through a ... |
| 3 | Optimizing MySQL | In this tutorial we will show ... |
| 4 | 1001 MySQL Tricks | 1. Never run mysqld as root. 2. ... |
| 6 | MySQL Security | When configured properly, MySQL ... |
+-----+-----+-----+
```

The `+` and `-` operators indicate that a word is required to be present or absent, respectively, for a match to occur. Thus, this query retrieves all the rows that contain the word “MySQL” but that do *not* contain the word “YourSQL”.

Boolean full-text searches have these characteristics:

- They do not use the 50% threshold.
- They do not automatically sort rows in order of decreasing relevance. You can see this from the preceding query result: The row with the highest relevance is the one that contains “MySQL” twice, but it is listed last, not first.
- They can work even without a `FULLTEXT` index, although a search executed in this fashion would be quite slow.
- The minimum and maximum word length full-text parameters apply.
- The stopword list applies.

The boolean full-text search capability supports the following operators:

- +

A leading plus sign indicates that this word *must* be present in each row that is returned.

- -

A leading minus sign indicates that this word must *not* be present in any of the rows that are returned.

Note: The - operator acts only to exclude rows that are otherwise matched by other search terms. Thus, a boolean-mode search that contains only terms preceded by - returns an empty result. It does not return “all rows except those containing any of the excluded terms.”

- (no operator)

By default (when neither + nor - is specified) the word is optional, but the rows that contain it are rated higher. This mimics the behavior of `MATCH() . . . AGAINST()` without the `IN BOOLEAN MODE` modifier.

- > <

These two operators are used to change a word's contribution to the relevance value that is assigned to a row. The > operator increases the contribution and the < operator decreases it. See the example following this list.

- ( )

Parentheses group words into subexpressions. Parenthesized groups can be nested.

- ~

A leading tilde acts as a negation operator, causing the word's contribution to the row's relevance to be negative. This is useful for marking “noise” words. A row containing such a word is rated lower than others, but is not excluded altogether, as it would be with the - operator.

- \*

The asterisk serves as the truncation (or wildcard) operator. Unlike the other operators, it should be *appended* to the word to be affected. Words match if they begin with the word preceding the \* operator.

- "

A phrase that is enclosed within double quote (‘ ’) characters matches only rows that contain the phrase *literally, as it was typed*. The full-text engine splits the phrase into words, performs a search in the `FULLTEXT` index for the words. The engine then performs a substring search for the phrase in the records that are found, so the match must include non-word characters in the phrase. For example, `"test phrase"` does *not* match `test, phrase`.

If the phrase contains no words that are in the index, the result is empty. For example, if all words are either stopwords or shorter than the minimum length of indexed words, the result is empty.

The following examples demonstrate some search strings that use boolean full-text operators:

- `'apple banana'`

Find rows that contain at least one of the two words.

- `'+apple +juice'`

Find rows that contain both words.

- `'+apple macintosh'`

Find rows that contain the word “apple”, but rank rows higher if they also contain “macintosh”.

- `'+apple -macintosh'`

Find rows that contain the word “apple” but not “macintosh”.

- `'+apple ~macintosh'`

Find rows that contain the word “apple”, but if the row also contains the word “macintosh”, rate it lower than if row does not. This is “softer” than a search for `'+apple -macintosh'`, for which the presence of “macintosh” causes the row not to be returned at all.

- `'+apple +( >turnover <strudel)'`

Find rows that contain the words “apple” and “turnover”, or “apple” and “strudel” (in any order), but rank “apple turnover” higher than “apple strudel”.

- `'apple*'`

Find rows that contain words such as “apple”, “apples”, “applesauce”, or “applet”.

- `'"some words"'`

Find rows that contain the exact phrase “some words” (for example, rows that contain “some words of wisdom” but not “some noise words”). Note that the ‘”’ characters that enclose the phrase are operator characters that delimit the phrase. They are not the quotes that enclose the search string itself.

## 12.7.2. Full-Text Searches with Query Expansion

As of MySQL 4.1.1, full-text search supports query expansion (in particular, its variant “blind query expansion”). This is generally useful when a search phrase is too short, which often means that the user is relying on implied knowledge that the full-text search engine lacks. For example, a user searching for “database” may really mean that “MySQL”, “Oracle”, “DB2”, and “RDBMS” all are phrases that should match “databases” and should be returned, too. This is implied knowledge.

Blind query expansion (also known as automatic relevance feedback) is enabled by adding `WITH QUERY EXPANSION` following the search phrase. It works by performing the search twice, where the search phrase for the second search is the original search phrase concatenated with the few most highly relevant documents from the first search. Thus, if one of these documents contains the word “databases” and the word “MySQL”, the second search finds the documents that contain the word “MySQL” even if they do not contain the word “database”. The following example shows this difference:

```
mysql> SELECT * FROM articles
-> WHERE MATCH (title,body) AGAINST ('database');
+-----+-----+-----+
| id | title | body |
+-----+-----+-----+
| 5 | MySQL vs. YourSQL | In the following database comparison ... |
| 1 | MySQL Tutorial | DBMS stands for DataBase ... |
+-----+-----+-----+
2 rows in set (0.00 sec)
```

```
mysql> SELECT * FROM articles
-> WHERE MATCH (title,body)
-> AGAINST ('database' WITH QUERY EXPANSION);
```

id	title	body
1	MySQL Tutorial	DBMS stands for DataBase ...
5	MySQL vs. YourSQL	In the following database comparison ...
3	Optimizing MySQL	In this tutorial we will show ...

```
3 rows in set (0.00 sec)
```

Another example could be searching for books by Georges Simenon about Maigret, when a user is not sure how to spell “Maigret”. A search for “Megre and the reluctant witnesses” finds only “Maigret and the Reluctant Witnesses” without query expansion. A search with query expansion finds all books with the word “Maigret” on the second pass.

**Note:** Because blind query expansion tends to increase noise significantly by returning non-relevant documents, it is meaningful to use only when a search phrase is rather short.

### 12.7.3. Full-Text Stopwords

The following table shows the default list of full-text stopwords.

a's	able	about	above	according
accordingly	across	actually	after	afterwards
again	against	ain't	all	allow
allows	almost	alone	along	already
also	although	always	am	among
amongst	an	and	another	any
anybody	anyhow	anyone	anything	anyway
anyways	anywhere	apart	appear	appreciate
appropriate	are	aren't	around	as
aside	ask	asking	associated	at
available	away	awfully	be	became
because	become	becomes	becoming	been
before	beforehand	behind	being	believe
below	beside	besides	best	better
between	beyond	both	brief	but
by	c'mon	c's	came	can
can't	cannot	cant	cause	causes
certain	certainly	changes	clearly	co
com	come	comes	concerning	consequently
consider	considering	contain	containing	contains
corresponding	could	couldn't	course	currently
definitely	described	despite	did	didn't
different	do	does	doesn't	doing
don't	done	down	downwards	during
each	edu	eg	eight	either
else	elsewhere	enough	entirely	especially

Functions and Operators

et	etc	even	ever	every
everybody	everyone	everything	everywhere	ex
exactly	example	except	far	few
fifth	first	five	followed	following
follows	for	former	formerly	forth
four	from	further	furthermore	get
gets	getting	given	gives	go
goes	going	gone	got	gotten
greetings	had	hadn't	happens	hardly
has	hasn't	have	haven't	having
he	he's	hello	help	hence
her	here	here's	hereafter	hereby
herein	hereupon	hers	herself	hi
him	himself	his	hither	hopefully
how	howbeit	however	i'd	i'll
i'm	i've	ie	if	ignored
immediate	in	inasmuch	inc	indeed
indicate	indicated	indicates	inner	insofar
instead	into	inward	is	isn't
it	it'd	it'll	it's	its
itself	just	keep	keeps	kept
know	knows	known	last	lately
later	latter	latterly	least	less
lest	let	let's	like	liked
likely	little	look	looking	looks
ltd	mainly	many	may	maybe
me	mean	meanwhile	merely	might
more	moreover	most	mostly	much
must	my	myself	name	namely
nd	near	nearly	necessary	need
needs	neither	never	nevertheless	new
next	nine	no	nobody	non
none	noone	nor	normally	not
nothing	novel	now	nowhere	obviously
of	off	often	oh	ok
okay	old	on	once	one
ones	only	onto	or	other
others	otherwise	ought	our	ours
ourselves	out	outside	over	overall
own	particular	particularly	per	perhaps
placed	please	plus	possible	presumably
probably	provides	que	quite	qv

Functions and Operators

rather	rd	re	really	reasonably
regarding	regardless	regards	relatively	respectively
right	said	same	saw	say
saying	says	second	secondly	see
seeing	seem	seemed	seeming	seems
seen	self	selves	sensible	sent
serious	seriously	seven	several	shall
she	should	shouldn't	since	six
so	some	somebody	somehow	someone
something	sometime	sometimes	somewhat	somewhere
soon	sorry	specified	specify	specifying
still	sub	such	sup	sure
t's	take	taken	tell	tends
th	than	thank	thanks	thanx
that	that's	that's	the	their
theirs	them	themselves	then	thence
there	there's	thereafter	thereby	therefore
therein	theres	thereupon	these	they
they'd	they'll	they're	they've	think
third	this	thorough	thoroughly	those
though	three	through	throughout	thru
thus	to	together	too	took
toward	towards	tried	tries	truly
try	trying	twice	two	un
under	unfortunately	unless	unlikely	until
unto	up	upon	us	use
used	useful	uses	using	usually
value	various	very	via	viz
vs	want	wants	was	wasn't
way	we	we'd	we'll	we're
we've	welcome	well	went	were
weren't	what	what's	whatever	when
whence	whenever	where	where's	whereafter
whereas	whereby	wherein	whereupon	wherever
whether	which	while	whither	who
who's	whoever	whole	whom	whose
why	will	willing	wish	with
within	without	won't	wonder	would
would	wouldn't	yes	yet	you
you'd	you'll	you're	you've	your
yours	yourself	yourselves	zero	

## 12.7.4. Full-Text Restrictions

- Full-text searches are supported for [MyISAM](#) tables only.
- As of MySQL 4.1.1, full-text searches can be used with most multi-byte character sets. The exception is that for Unicode, the [utf8](#) character set can be used, but not the [ucs2](#) character set.
- Ideographic languages such as Chinese and Japanese do not have word delimiters. Therefore, the [FULLTEXT](#) parser *cannot determine where words begin and end in these and other such languages*. The implications of this and some workarounds for the problem are described in [Section 12.7, “Full-Text Search Functions”](#).
- As of MySQL 4.1, the use of multiple character sets within a single table is supported. However, all columns in a [FULLTEXT](#) index must use the same character set and collation.
- The [MATCH\(\)](#) column list must match exactly the column list in some [FULLTEXT](#) index definition for the table, unless this [MATCH\(\)](#) is [IN BOOLEAN MODE](#). Boolean-mode searches can be done on non-indexed columns, although they are likely to be slow.
- The argument to [AGAINST\(\)](#) must be a constant string.

## 12.7.5. Fine-Tuning MySQL Full-Text Search

MySQL's full-text search capability has few user-tunable parameters. You can exert more control over full-text searching behavior if you have a MySQL source distribution because some changes require source code modifications. See [Section 2.9, “MySQL Installation Using a Source Distribution”](#).

Note that full-text search is carefully tuned for the most effectiveness. Modifying the default behavior in most cases can actually decrease effectiveness. *Do not alter the MySQL sources unless you know what you are doing.*

Most full-text variables described in this section must be set at server startup time. A server restart is required to change them; they cannot be modified while the server is running.

Some variable changes require that you rebuild the [FULLTEXT](#) indexes in your tables. Instructions for doing this are given at the end of this section.

- The minimum and maximum lengths of words to be indexed are defined by the [ft\\_min\\_word\\_len](#) and [ft\\_max\\_word\\_len](#) system variables (available as of MySQL 4.0.0). See [Section 5.2.2, “System Variables”](#).) The default minimum value is four characters; the default maximum is version dependent. If you change either value, you must rebuild your [FULLTEXT](#) indexes. For example, if you want three-character words to be searchable, you can set the [ft\\_min\\_word\\_len](#) variable by putting the following lines in an option file:

```
[mysqld]
ft_min_word_len=3
```

Then you must restart the server and rebuild your [FULLTEXT](#) indexes. Note particularly the remarks regarding [mysamchk](#) in the instructions following this list.

- To override the default stopword list, set the [ft\\_stopword\\_file](#) system variable (available as of MySQL 4.0.10). See [Section 5.2.2, “System Variables”](#).) The variable value should be the pathname of the file containing the stopword list, or the empty string to disable stopword filtering. After changing the value of this variable or the contents of the stopword file, restart the server and rebuild your [FULLTEXT](#) indexes.

The stopword list is free-form. That is, you may use any non-alphanumeric character such as newline, space, or comma to separate stopwords. Exceptions are the underscore character ('\_') and a single apostrophe (' ') which are treated as part of a word. The character set of the stopword list is the server's default character set; see [Section 10.3.1, “Server Character Set and Collation”](#).

- The 50% threshold for natural language searches is determined by the particular weighting scheme chosen. To disable it, look for the following line in `myisam/ftdefs.h`:

```
#define GWS_IN_USE GWS_PROB
```

Change that line to this:

```
#define GWS_IN_USE GWS_FREQ
```

Then recompile MySQL. There is no need to rebuild the indexes in this case. **Note:** By making this change, you *severely* decrease MySQL's ability to provide adequate relevance values for the `MATCH()` function. If you really need to search for such common words, it would be better to search using `IN BOOLEAN MODE` instead, which does not observe the 50% threshold.

- To change the operators used for boolean full-text searches, set the `ft_boolean_syntax` system variable (available as of MySQL 4.0.1). The variable can be changed while the server is running, but you must have the `SUPER` privilege to do so. No rebuilding of indexes is necessary in this case. See [Section 5.2.2, “System Variables”](#), which describes the rules governing how to set this variable.
- If you want to change the set of characters that are considered word characters, you can do so in two ways. Suppose that you want to treat the hyphen character ('-') as a word character. Use either of these methods:
  - Modify the MySQL source: In `myisam/ftdefs.h`, see the `true_word_char()` and `misc_word_char()` macros. Add '-' to one of those macros and recompile MySQL.
  - Modify a character set file: This requires no recompilation. The `true_word_char()` macro uses a “character type” table to distinguish letters and numbers from other characters. You can edit the `<ctype><map>` contents in one of the character set XML files to specify that '-' is a “letter.” Then use the given character set for your `FULLTEXT` indexes.

After making the modification, you must rebuild the indexes for each table that contains any `FULLTEXT` indexes.

If you modify full-text variables that affect indexing (`ft_min_word_len`, `ft_max_word_len`, or `ft_stopword_file`), or if you change the stopword file itself, you must rebuild your `FULLTEXT` indexes after making the changes and restarting the server. To rebuild the indexes in this case, it is sufficient to do a `QUICK` repair operation:

```
mysql> REPAIR TABLE tbl_name QUICK;
```

Each table that contains any `FULLTEXT` index must be repaired as just shown. Otherwise, queries for the table may yield incorrect results, and modifications to the table will cause the server to see the table as corrupt and in need of repair.

With regard specifically to using the `IN BOOLEAN MODE` capability, if you upgrade from MySQL 3.23 to 4.0 or later, it is necessary to replace the index header as well. To do this, perform a `USE_FRM` repair operation:

```
mysql> REPAIR TABLE tbl_name USE_FRM;
```



This is necessary because boolean full-text searches require a flag in the index header that was not present in MySQL 3.23, and that is not added if you do only a `QUICK` repair. If you attempt a boolean full-text search without rebuilding the indexes this way, the search returns incorrect results.

Note that if you use `myisamchk` to perform an operation that modifies table indexes (such as repair or analyze), the `FULLTEXT` indexes are rebuilt using the *default* full-text parameter values for minimum word length, maximum word length, and stopword file unless you specify otherwise. This can result in queries failing.

The problem occurs because these parameters are known only by the server. They are not stored in `MyISAM` index files. To avoid the problem if you have modified the minimum or maximum word length or stopword file values used by the server, specify the same `ft_min_word_len`, `ft_max_word_len`, and `ft_stopword_file` values to `myisamchk` that you use for `mysqld`. For example, if you have set the minimum word length to 3, you can repair a table with `myisamchk` like this:

```
shell> myisamchk --recover --ft_min_word_len=3 tbl_name.MYI
```

To ensure that `myisamchk` and the server use the same values for full-text parameters, place each one in both the `[mysqld]` and `[myisamchk]` sections of an option file:

```
[mysqld]
ft_min_word_len=3

[myisamchk]
ft_min_word_len=3
```

An alternative to using `myisamchk` is to use the `REPAIR TABLE`, `ANALYZE TABLE`, `OPTIMIZE TABLE`, or `ALTER TABLE` statements. These statements are performed by the server, which knows the proper full-text parameter values to use.

## 12.8. Cast Functions and Operators

- `BINARY`

The `BINARY` operator casts the string following it to a binary string. This is an easy way to force a column comparison to be done byte by byte rather than character by character. This causes the comparison to be case sensitive even if the column isn't defined as `BINARY` or `BLOB`. `BINARY` also causes trailing spaces to be significant.

```
mysql> SELECT 'a' = 'A';
-> 1
mysql> SELECT BINARY 'a' = 'A';
-> 0
mysql> SELECT 'a' = 'a ';
-> 1
mysql> SELECT BINARY 'a' = 'a ';
-> 0
```

In a comparison, `BINARY` affects the entire operation; it can be given before either operand with the same result.

`BINARY` was added in MySQL 3.23.0. As of MySQL 4.0.2, `BINARY str` is a shorthand for `CAST(str AS BINARY)`.

Note that in some contexts, if you cast an indexed column to `BINARY`, MySQL is not able to use the index efficiently.

- `CAST(expr AS type)`, `CONVERT(expr, type)`, `CONVERT(expr USING transcoding_name)`

The `CAST()` and `CONVERT()` functions take a value of one type and produce a value of another type.

The *type* can be one of the following values:

- `BINARY` (and `BINARY[N]` as of MySQL 4.1.1)
- `CHAR` (and `CHAR[N]` as of MySQL 4.1.1)
- `DATE`
- `DATETIME`
- `SIGNED [INTEGER]`
- `TIME`
- `UNSIGNED [INTEGER]`

`BINARY` produces a string with the `BINARY` data type. See [Section 11.4.2, “The BINARY and VARBINARY Types”](#) for a description of how this affects comparisons. If the optional length *N* is given, `BINARY(N)` causes the cast to use no more than *N* bytes of the argument. Similarly, `CHAR[N]` causes the cast to use no more than *N* characters of the argument.

`CAST()` and `CONVERT()` are available as of MySQL 4.0.2. The `CHAR` conversion type is available as of 4.0.6. The `USING` form of `CONVERT()` is available as of 4.1.0.

`CAST()` and `CONVERT(... USING ...)` are standard SQL syntax. The non-`USING` form of `CONVERT()` is ODBC syntax.

`CONVERT()` with `USING` is used to convert data between different character sets. In MySQL, transcoding names are the same as the corresponding character set names. For example, this statement converts the string `'abc'` in the default character set to the corresponding string in the `utf8` character set:

```
SELECT CONVERT('abc' USING utf8);
```

If you want to compare a `BLOB` value or other binary string in case-insensitive fashion, you can do so as follows:

- Before MySQL 4.1.1, use the `UPPER()` function to convert the binary string to uppercase before performing the comparison:

```
SELECT 'A' LIKE UPPER(blob_col) FROM tbl_name;
```

If the comparison value is lowercase, convert the string value using `LOWER()` instead.

- For MySQL 4.1.1 and up, binary strings have no character set, and thus no concept of lettercase. To perform a case-insensitive comparison, use the `CONVERT()` function to convert the value to a non-binary string. If the character set of the result has a case-insensitive collation, the `LIKE` operation is not case sensitive:

```
SELECT 'A' LIKE CONVERT(blob_col USING latin1) FROM tbl_name;
```

To use a different character set, substitute its name for `latin1` in the preceding statement. To en-

sure that a case-insensitive collation is used, specify a `COLLATE` clause following the `CONVERT()` call.

`CONVERT()` can be used more generally for comparing strings that are represented in different character sets.

The cast functions are useful when you want to create a column with a specific type in a `CREATE . . . SELECT` statement:

```
CREATE TABLE new_table SELECT CAST('2000-01-01' AS DATE);
```

The functions also can be useful for sorting `ENUM` columns in lexical order. Normally, sorting of `ENUM` columns occurs using the internal numeric values. Casting the values to `CHAR` results in a lexical sort:

```
SELECT enum_col FROM tbl_name ORDER BY CAST(enum_col AS CHAR);
```

`CAST(str AS BINARY)` is the same thing as `BINARY str`. `CAST(expr AS CHAR)` treats the expression as a string with the default character set.

**Note:** In MySQL 4.0, a `CAST()` to `DATE`, `DATETIME`, or `TIME` only marks the column to be a specific type but does not change the value of the column.

As of MySQL 4.1.0, the value is converted to the correct column type when it is sent to the user (this is a feature of how the new protocol in 4.1 sends date information to the client):

```
mysql> SELECT CAST(NOW() AS DATE);
-> 2003-05-26
```

As of MySQL 4.1.1, `CAST()` also changes the result if you use it as part of a more complex expression such as `CONCAT('Date: ', CAST(NOW() AS DATE))`.

You should not use `CAST()` to extract data in different formats but instead use string functions like `LEFT()` or `EXTRACT()`. See [Section 12.5, “Date and Time Functions”](#).

To cast a string to a numeric value in numeric context, you normally do not have to do anything other than to use the string value as though it were a number:

```
mysql> SELECT 1+'1';
-> 2
```

If you use a number in string context, the number automatically is converted to a `BINARY` string.

```
mysql> SELECT CONCAT('hello you ',2);
-> 'hello you 2'
```

MySQL supports arithmetic with both signed and unsigned 64-bit values. If you are using numeric operators (such as `+` or `-`) and one of the operands is an unsigned integer, the result is unsigned. You can override this by using the `SIGNED` and `UNSIGNED` cast operators to cast the operation to a signed or unsigned 64-bit integer, respectively.

```
mysql> SELECT CAST(1-2 AS UNSIGNED)
-> 18446744073709551615
mysql> SELECT CAST(CAST(1-2 AS UNSIGNED) AS SIGNED);
-> -1
```

Note that if either operand is a floating-point value, the result is a floating-point value and is not affected by the preceding rule. (In this context, `DECIMAL` column values are regarded as floating-point values.)

```
mysql> SELECT CAST(1 AS UNSIGNED) - 2.0;
-> -1.0
```

If you are using a string in an arithmetic operation, this is converted to a floating-point number.

The handling of unsigned values was changed in MySQL 4.0 to be able to support `BIGINT` values properly. If you have some code that you want to run in both MySQL 4.0 and 3.23, you probably cannot use the `CAST()` function. You can use the following technique to get a signed result when subtracting two unsigned integer columns `ucol1` and `ucol2`:

```
mysql> SELECT (ucol1+0.0)-(ucol2+0.0) FROM ...;
```

The idea is that the columns are converted to floating-point values before the subtraction occurs.

If you have a problem with `UNSIGNED` columns in old MySQL applications when porting them to MySQL 4.0, you can use the `--sql-mode=NO_UNSIGNED_SUBTRACTION` option when starting `mysqld`. However, as long as you use this option, you are not able to make efficient use of the `BIGINT UNSIGNED` data type.

## 12.9. Other Functions

### 12.9.1. Bit Functions

MySQL uses `BIGINT` (64-bit) arithmetic for bit operations, so these operators have a maximum range of 64 bits.

- |

Bitwise OR:

```
mysql> SELECT 29 | 15;
-> 31
```

The result is an unsigned 64-bit integer.

- &

Bitwise AND:

```
mysql> SELECT 29 & 15;
-> 13
```

The result is an unsigned 64-bit integer.

- ^

Bitwise XOR:

```
mysql> SELECT 1 ^ 1;
-> 0
mysql> SELECT 1 ^ 0;
-> 1
mysql> SELECT 11 ^ 3;
-> 8
```

The result is an unsigned 64-bit integer.

Bitwise XOR was added in MySQL 4.0.2.

- <<

Shifts a longlong (`BIGINT`) number to the left.

```
mysql> SELECT 1 << 2;
-> 4
```

The result is an unsigned 64-bit integer.

- >>

Shifts a longlong (`BIGINT`) number to the right.

```
mysql> SELECT 4 >> 2;
-> 1
```

The result is an unsigned 64-bit integer.

- ~

Invert all bits.

```
mysql> SELECT 5 & ~1;
-> 4
```

The result is an unsigned 64-bit integer.

- `BIT_COUNT(N)`

Returns the number of bits that are set in the argument *N*.

```
mysql> SELECT BIT_COUNT(29);
-> 4
```

## 12.9.2. Encryption and Compression Functions

The functions in this section perform encryption and decryption, and compression and uncompression.

**Note:** The encryption and compression functions return binary strings. For many of these functions, the result might contain arbitrary byte values. If you want to store these results, use a `BLOB` column rather than a `CHAR` or `VARCHAR` column to avoid potential problems with trailing space removal that would change data values.

**Note:** Exploits for the MD5 and SHA-1 algorithms have become known. You may wish to consider using one of the other encryption functions described in this section instead.

- `AES_ENCRYPT(str, key_str)`, `AES_DECRYPT(encrypt_str, key_str)`

These functions allow encryption and decryption of data using the official AES (Advanced Encryption Standard) algorithm, previously known as “Rijndael.” Encoding with a 128-bit key length is used, but you can extend it up to 256 bits by modifying the source. We chose 128 bits because it is much faster and it is secure enough for most purposes.

`AES_ENCRYPT()` encrypts a string and returns a binary string. `AES_DECRYPT()` decrypts the encrypted string and returns the original string. The input arguments may be any length. If either argument is `NULL`, the result of this function is also `NULL`.

Because AES is a block-level algorithm, padding is used to encode uneven length strings and so the result string length may be calculated using this formula:

```
16 × (trunc(string_length / 16) + 1)
```

If `AES_DECRYPT()` detects invalid data or incorrect padding, it returns `NULL`. However, it is possible for `AES_DECRYPT()` to return a non-`NULL` value (possibly garbage) if the input data or the key is invalid.

You can use the AES functions to store data in an encrypted form by modifying your queries:

```
INSERT INTO t VALUES (1,AES_ENCRYPT('text','password'));
```

`AES_ENCRYPT()` and `AES_DECRYPT()` were added in MySQL 4.0.2, and can be considered the most cryptographically secure encryption functions available in MySQL.

- `COMPRESS(string_to_compress)`

Compresses a string and returns the result as a binary string. This function requires MySQL to have been compiled with a compression library such as `zlib`. Otherwise, the return value is always `NULL`. The compressed string can be uncompressed with `UNCOMPRESS()`.

```
mysql> SELECT LENGTH(COMPRESS(REPEAT('a',1000)));
-> 21
mysql> SELECT LENGTH(COMPRESS(''));
-> 0
mysql> SELECT LENGTH(COMPRESS('a'));
-> 13
mysql> SELECT LENGTH(COMPRESS(REPEAT('a',16)));
-> 15
```

The compressed string contents are stored the following way:

- Empty strings are stored as empty strings.
- Non-empty strings are stored as a four-byte length of the uncompressed string (low byte first), followed by the compressed string. If the string ends with space, an extra `.` character is added to avoid problems with endspace trimming should the result be stored in a `CHAR` or `VARCHAR` column. (Use of `CHAR` or `VARCHAR` to store compressed strings is not recommended. It is better to use a `BLOB` column instead.)

`COMPRESS()` was added in MySQL 4.1.1.

- `DECODE(encrypt_str,pass_str)`

Decrypts the encrypted string `encrypt_str` using `pass_str` as the password. `encrypt_str` should be a string returned from `ENCODE()`.

- `ENCODE(str,pass_str)`

Encrypt `str` using `pass_str` as the password. To decrypt the result, use `DECODE()`.

The result is a binary string of the same length as `str`.

The strength of the encryption is based on how good the random generator is. It should suffice for

short strings.

- `DES_DECRYPT(crypt_str[, key_str])`

Decrypts a string encrypted with `DES_ENCRYPT()`. If an error occurs, this function returns `NULL`.

Note that this function works only if MySQL has been configured with SSL support. See [Section 5.8.7, “Using Secure Connections”](#).

If no *key\_str* argument is given, `DES_DECRYPT()` examines the first byte of the encrypted string to determine the DES key number that was used to encrypt the original string, and then reads the key from the DES key file to decrypt the message. For this to work, the user must have the `SUPER` privilege. The key file can be specified with the `--des-key-file` server option.

If you pass this function a *key\_str* argument, that string is used as the key for decrypting the message.

If the *crypt\_str* argument does not appear to be an encrypted string, MySQL returns the given *crypt\_str*.

`DES_DECRYPT()` was added in MySQL 4.0.1.

- `DES_ENCRYPT(str[, {key_num|key_str}] )`

Encrypts the string with the given key using the Triple-DES algorithm.

Note that this function works only if MySQL has been configured with SSL support. See [Section 5.8.7, “Using Secure Connections”](#).

The encryption key to use is chosen based on the second argument to `DES_ENCRYPT()`, if one was given:

Argument	Description
No argument	The first key from the DES key file is used.
<i>key_num</i>	The given key number (0-9) from the DES key file is used.
<i>key_str</i>	The given key string is used to encrypt <i>str</i> .

The key file can be specified with the `--des-key-file` server option.

The return string is a binary string where the first character is `CHAR(128 | key_num)`. If an error occurs, `DES_ENCRYPT()` returns `NULL`.

The 128 is added to make it easier to recognize an encrypted key. If you use a string key, *key\_num* is 127.

The string length for the result is given by this formula:

$$new\_len = orig\_len + (8 - (orig\_len \% 8)) + 1$$

Each line in the DES key file has the following format:

```
key_num des_key_str
```

Each *key\_num* value must be a number in the range from 0 to 9. Lines in the file may be in any order. *des\_key\_str* is the string that is used to encrypt the message. There should be at least one

space between the number and the key. The first key is the default key that is used if you do not specify any key argument to `DES_ENCRYPT()`.

You can tell MySQL to read new key values from the key file with the `FLUSH DES_KEY_FILE` statement. This requires the `RELOAD` privilege.

One benefit of having a set of default keys is that it gives applications a way to check for the existence of encrypted column values, without giving the end user the right to decrypt those values.

```
mysql> SELECT customer_address FROM customer_table
> WHERE crypt_credit_card = DES_ENCRYPT('credit_card_number');
```

`DES_ENCRYPT()` was added in MySQL 4.0.1.

- `ENCRYPT(str[, salt])`

Encrypts `str` using the Unix `crypt()` system call and returns a binary string. The `salt` argument should be a string with two characters. (As of MySQL 3.22.16, `salt` may be longer than two characters.) If no `salt` argument is given, a random value is used.

```
mysql> SELECT ENCRYPT('hello');
-> 'VxuFAJXVARROc'
```

`ENCRYPT()` ignores all but the first eight characters of `str`, at least on some systems. This behavior is determined by the implementation of the underlying `crypt()` system call.

The use of `ENCRYPT()` with multi-byte character sets other than `utf8` is not recommended because the system call expects a string terminated by a zero byte.

If `crypt()` is not available on your system (as is the case with Windows), `ENCRYPT()` always returns `NULL`.

- `MD5(str)`

Calculates an MD5 128-bit checksum for the string. The value is returned as a binary string of 32 hex digits, or `NULL` if the argument was `NULL`. The return value can, for example, be used as a hash key.

```
mysql> SELECT MD5('testing');
-> 'ae2b1fca515949e5d54fb22b8ed95575'
```

This is the “RSA Data Security, Inc. MD5 Message-Digest Algorithm.”

If you want to convert the value to uppercase, see the description of binary string conversion given in the entry for the `BINARY` operator in [Section 12.8, “Cast Functions and Operators”](#).

See the note regarding the MD5 algorithm at the beginning this section.

`MD5()` was added in MySQL 3.23.2.

- `OLD_PASSWORD(str)`

`OLD_PASSWORD()` is available as of MySQL 4.1, when the implementation of `PASSWORD()` was changed to improve security. `OLD_PASSWORD()` returns the value of the pre-4.1 implementation of `PASSWORD()` as a binary string, and is intended to permit you to reset passwords for any pre-4.1 clients that need to connect to your version 4.1 MySQL server without locking them out. See [Section 5.7.9, “Password Hashing as of MySQL 4.1”](#).

- `PASSWORD(str)`



Calculates and returns a password string from the plaintext password *str* and returns a binary string, or `NULL` if the argument was `NULL`. This is the function that is used for encrypting MySQL passwords for storage in the `Password` column of the `user` grant table.

```
mysql> SELECT PASSWORD('badpwd');
-> '7f84554057dd964b'
```

`PASSWORD()` encryption is one-way (not reversible).

`PASSWORD()` does not perform password encryption in the same way that Unix passwords are encrypted. See `ENCRYPT()`.

**Note:** The `PASSWORD()` function is used by the authentication system in MySQL Server; you should *not* use it in your own applications. For that purpose, consider `MD5()` or `SHA1()` instead. Also see RFC 2195 for more information about handling passwords and authentication securely in your applications.

- `SHA1(str)`, `SHA(str)`

Calculates an SHA-1 160-bit checksum for the string, as described in RFC 3174 (Secure Hash Algorithm). The value is returned as a binary string of 40 hex digits, or `NULL` if the argument was `NULL`. One of the possible uses for this function is as a hash key. You can also use it as a cryptographic function for storing passwords. `SHA()` is synonymous with `SHA1()`.

```
mysql> SELECT SHA1('abc');
-> 'a9993e364706816aba3e25717850c26c9cd0d89d'
```

`SHA1()` was added in MySQL 4.0.2, and can be considered a cryptographically more secure equivalent of `MD5()`. However, see the note regarding the MD5 and SHA-1 algorithms at the beginning this section.

- `UNCOMPRESS(string_to_uncompress)`

Uncompresses a string compressed by the `COMPRESS()` function. If the argument is not a compressed value, the result is `NULL`. This function requires MySQL to have been compiled with a compression library such as `zlib`. Otherwise, the return value is always `NULL`.

```
mysql> SELECT UNCOMPRESS(COMPRESS('any string'));
-> 'any string'
mysql> SELECT UNCOMPRESS('any string');
-> NULL
```

`UNCOMPRESS()` was added in MySQL 4.1.1.

- `UNCOMPRESSED_LENGTH(compressed_string)`

Returns the length that the compressed string had before being compressed.

```
mysql> SELECT UNCOMPRESSED_LENGTH(COMPRESS(REPEAT('a',30)));
-> 30
```

`UNCOMPRESSED_LENGTH()` was added in MySQL 4.1.1.

## 12.9.3. Information Functions

- `BENCHMARK(count,expr)`

The `BENCHMARK()` function executes the expression *expr* repeatedly *count* times. It may be used to time how quickly MySQL processes the expression. The result value is always 0. The intended use is from within the `mysql` client, which reports query execution times:

```
mysql> SELECT BENCHMARK(1000000,ENCODE('hello','goodbye'));
+-----+
| BENCHMARK(1000000,ENCODE('hello','goodbye')) |
+-----+
|                                             0 |
+-----+
1 row in set (4.74 sec)
```

The time reported is elapsed time on the client end, not CPU time on the server end. It is advisable to execute `BENCHMARK()` several times, and to interpret the result with regard to how heavily loaded the server machine is.

- `CHARSET(str)`

Returns the character set of the string argument.

```
mysql> SELECT CHARSET('abc');
-> 'latin1'
mysql> SELECT CHARSET(CONVERT('abc' USING utf8));
-> 'utf8'
mysql> SELECT CHARSET(USER());
-> 'utf8'
```

`CHARSET()` was added in MySQL 4.1.0.

- `COERCIBILITY(str)`

Returns the collation coercibility value of the string argument.

```
mysql> SELECT COERCIBILITY('abc' COLLATE latin1_swedish_ci);
-> 0
mysql> SELECT COERCIBILITY(USER());
-> 3
mysql> SELECT COERCIBILITY('abc');
-> 4
```

The return values have the meanings shown in the following table. Lower values have higher precedence.

Coercibility	Meaning	Example
0	Explicit collation	Value with <code>COLLATE</code> clause
1	No collation	Concatenation of strings with different collations
2	Implicit collation	Column value
3	System constant	<code>USER()</code> return value
4	Coercible	Literal string
5	Ignorable	<code>NULL</code> or an expression derived from <code>NULL</code>

Before MySQL 4.1.11, the return values are shown in following table, and functions such as `USER()` have a coercibility of 2.

Coercibility	Meaning	Example
0	Explicit collation	Value with <code>COLLATE</code> clause
1	No collation	Concatenation of strings with different collations
2	Implicit collation	Column value
3	Coercible	Literal string

`COERCIBILITY()` was added in MySQL 4.1.1.

- `COLLATION(str)`

Returns the collation of the string argument.

```
mysql> SELECT COLLATION('abc');
      -> 'latin1_swedish_ci'
mysql> SELECT COLLATION(_utf8'abc');
      -> 'utf8_general_ci'
```

`COLLATION()` was added in MySQL 4.1.0.

- `CONNECTION_ID()`

Returns the connection ID (thread ID) for the connection. Every connection has an ID that is unique among the set of currently connected clients.

```
mysql> SELECT CONNECTION_ID();
      -> 23786
```

`CONNECTION_ID()` was added in MySQL 3.23.14.

- `CURRENT_USER, CURRENT_USER()`

Returns the username and hostname combination for the MySQL account that the server used to authenticate the current client. This account determines your access privileges.

The value of `CURRENT_USER()` can differ from the value of `USER()`.

```
mysql> SELECT USER();
      -> 'davida@localhost'
mysql> SELECT * FROM mysql.user;
ERROR 1044: Access denied for user ''@'localhost' to
database 'mysql'
mysql> SELECT CURRENT_USER();
      -> '@localhost'
```

The example illustrates that although the client specified a username of `davida` (as indicated by the value of the `USER()` function), the server authenticated the client using an anonymous user account (as seen by the empty username part of the `CURRENT_USER()` value). One way this might occur is that there is no account listed in the grant tables for `davida`.

`CURRENT_USER()` was added in MySQL 4.0.6. As of MySQL 4.1.0, the string uses the `utf8` character set.

- `DATABASE()`

Returns the default (current) database name. As of MySQL 4.1, the string uses the `utf8` character

set. If there is no default database, `DATABASE ( )` returns `NULL` as of MySQL 4.1.1, and the empty string before that.

```
mysql> SELECT DATABASE();
-> 'test'
```

- `FOUND_ROWS ( )`

A `SELECT` statement may include a `LIMIT` clause to restrict the number of rows the server returns to the client. In some cases, it is desirable to know how many rows the statement would have returned without the `LIMIT`, but without running the statement again. To obtain this row count, include a `SQL_CALC_FOUND_ROWS` option in the `SELECT` statement, and then invoke `FOUND_ROWS ( )` afterward:

```
mysql> SELECT SQL_CALC_FOUND_ROWS * FROM tbl_name
-> WHERE id > 100 LIMIT 10;
mysql> SELECT FOUND_ROWS();
```

The second `SELECT` returns a number indicating how many rows the first `SELECT` would have returned had it been written without the `LIMIT` clause. (If the preceding `SELECT` statement does not include the `SQL_CALC_FOUND_ROWS` option, then `FOUND_ROWS ( )` may return a different result when `LIMIT` is used than when it is not.)

The row count available through `FOUND_ROWS ( )` is transient and not intended to be available past the statement following the `SELECT SQL_CALC_FOUND_ROWS` statement. If you need to refer to the value later, save it:

```
mysql> SELECT SQL_CALC_FOUND_ROWS * FROM ... ;
mysql> SET @rows = FOUND_ROWS();
```

If you are using `SELECT SQL_CALC_FOUND_ROWS`, MySQL must calculate how many rows are in the full result set. However, this is faster than running the query again without `LIMIT`, because the result set need not be sent to the client.

`SQL_CALC_FOUND_ROWS` and `FOUND_ROWS ( )` can be useful in situations when you want to restrict the number of rows that a query returns, but also determine the number of rows in the full result set without running the query again. An example is a Web script that presents a paged display containing links to the pages that show other sections of a search result. Using `FOUND_ROWS ( )` allows you to determine how many other pages are needed for the rest of the result.

The use of `SQL_CALC_FOUND_ROWS` and `FOUND_ROWS ( )` is more complex for `UNION` statements than for simple `SELECT` statements, because `LIMIT` may occur at multiple places in a `UNION`. It may be applied to individual `SELECT` statements in the `UNION`, or global to the `UNION` result as a whole.

The intent of `SQL_CALC_FOUND_ROWS` for `UNION` is that it should return the row count that would be returned without a global `LIMIT`. The conditions for use of `SQL_CALC_FOUND_ROWS` with `UNION` are:

- The `SQL_CALC_FOUND_ROWS` keyword must appear in the first `SELECT` of the `UNION`.
- The value of `FOUND_ROWS ( )` is exact only if `UNION ALL` is used. If `UNION` without `ALL` is used, duplicate removal occurs and the value of `FOUND_ROWS ( )` is only approximate.
- If no `LIMIT` is present in the `UNION`, `SQL_CALC_FOUND_ROWS` is ignored and returns the number of rows in the temporary table that is created to process the `UNION`.

`SQL_CALC_FOUND_ROWS` and `FOUND_ROWS ( )` are available starting at MySQL 4.0.0.

- `LAST_INSERT_ID()`, `LAST_INSERT_ID(expr)`

`LAST_INSERT_ID()` (with no argument) returns the *first* automatically generated value that was set for an `AUTO_INCREMENT` column by the *most recently executed* `INSERT` or `UPDATE` statement to affect such a column. For example, after inserting a row that generates an `AUTO_INCREMENT` value, you can get the value like this:

```
mysql> SELECT LAST_INSERT_ID();
-> 195
```

The currently executing statement does not affect the value of `LAST_INSERT_ID()`. Suppose that you generate an `AUTO_INCREMENT` value with one statement, and then refer to `LAST_INSERT_ID()` in a multiple-row `INSERT` statement that inserts rows into a table with its own `AUTO_INCREMENT` column. The value of `LAST_INSERT_ID()` will remain stable in the second statement; its value for the second and later rows is not affected by the earlier row insertions. (However, if you mix references to `LAST_INSERT_ID()` and `LAST_INSERT_ID(expr)`, the effect is undefined.)

The ID that was generated is maintained in the server on a *per-connection basis*. This means that the value returned by the function to a given client is the first `AUTO_INCREMENT` value generated for most recent statement affecting an `AUTO_INCREMENT` column *by that client*. This value cannot be affected by other clients, even if they generate `AUTO_INCREMENT` values of their own. This behavior ensures that each client can retrieve its own ID without concern for the activity of other clients, and without the need for locks or transactions.

The value of `LAST_INSERT_ID()` is not changed if you set the `AUTO_INCREMENT` column of a row to a non-“magic” value (that is, a value that is not `NULL` and not 0).

**Important:** If you insert multiple rows using a single `INSERT` statement, `LAST_INSERT_ID()` returns the value generated for the *first* inserted row *only*. The reason for this is to make it possible to reproduce easily the same `INSERT` statement against some other server.

For example:

```
mysql> USE test;
Database changed
mysql> CREATE TABLE t (
->   id INT AUTO_INCREMENT NOT NULL PRIMARY KEY,
->   name VARCHAR(10) NOT NULL
-> );
Query OK, 0 rows affected (0.09 sec)

mysql> INSERT INTO t VALUES (NULL, 'Bob');
Query OK, 1 row affected (0.01 sec)

mysql> SELECT * FROM t;
+----+-----+
| id | name |
+----+-----+
|  1 | Bob  |
+----+-----+
1 row in set (0.01 sec)

mysql> SELECT LAST_INSERT_ID();
+-----+
| LAST_INSERT_ID() |
+-----+
|                  1 |
+-----+
1 row in set (0.00 sec)

mysql> INSERT INTO t VALUES
-> (NULL, 'Mary'), (NULL, 'Jane'), (NULL, 'Lisa');
Query OK, 3 rows affected (0.00 sec)
Records: 3 Duplicates: 0 Warnings: 0

mysql> SELECT * FROM t;
+----+-----+
| id | name |
+----+-----+
```

```

+----+-----+
| id | name |
+----+-----+
|  1 | Bob  |
|  2 | Mary |
|  3 | Jane |
|  4 | Lisa |
+----+-----+
4 rows in set (0.01 sec)

mysql> SELECT LAST_INSERT_ID();
+-----+
| LAST_INSERT_ID() |
+-----+
|                2 |
+-----+
1 row in set (0.00 sec)

```

Although the second `INSERT` statement inserted three new rows into `t`, the ID generated for the first of these rows was 2, and it is this value that is returned by `LAST_INSERT_ID()` for the following `SELECT` statement.

If you use `INSERT IGNORE` and the row is ignored, the `AUTO_INCREMENT` counter is not incremented and `LAST_INSERT_ID()` returns 0, which reflects that no row was inserted. (Before MySQL 4.1, the `AUTO_INCREMENT` counter is still incremented and `LAST_INSERT_ID()` returns the new value.)

If `expr` is given as an argument to `LAST_INSERT_ID()`, the value of the argument is returned by the function and is remembered as the next value to be returned by `LAST_INSERT_ID()`. This can be used to simulate sequences:

1. Create a table to hold the sequence counter and initialize it:

```
mysql> CREATE TABLE sequence (id INT NOT NULL);
mysql> INSERT INTO sequence VALUES (0);
```

2. Use the table to generate sequence numbers like this:

```
mysql> UPDATE sequence SET id=LAST_INSERT_ID(id+1);
mysql> SELECT LAST_INSERT_ID();
```

The `UPDATE` statement increments the sequence counter and causes the next call to `LAST_INSERT_ID()` to return the updated value. The `SELECT` statement retrieves that value. The `mysql_insert_id()` C API function can also be used to get the value. See [Section 17.2.3.35](#), “`mysql_insert_id()`”.

You can generate sequences without calling `LAST_INSERT_ID()`, but the utility of using the function this way is that the ID value is maintained in the server as the last automatically generated value. It is multi-user safe because multiple clients can issue the `UPDATE` statement and get their own sequence value with the `SELECT` statement (or `mysql_insert_id()`), without affecting or being affected by other clients that generate their own sequence values.

Note that `mysql_insert_id()` is only updated after `INSERT` and `UPDATE` statements, so you cannot use the C API function to retrieve the value for `LAST_INSERT_ID(expr)` after executing other SQL statements like `SELECT` or `SET`.

- `SESSION_USER()`

`SESSION_USER()` is a synonym for `USER()`.

- `SYSTEM_USER()`

`SYSTEM_USER()` is a synonym for `USER()`.

- `USER()`

Returns the current MySQL username and hostname.

```
mysql> SELECT USER();
-> 'davida@localhost'
```

The value indicates the username you specified when connecting to the server, and the client host from which you connected. The value can be different from that of `CURRENT_USER()`.

Prior to MySQL 3.22.11, the function value does not include the client hostname. You can extract only the username part, regardless of whether the value includes a hostname part, like this:

```
mysql> SELECT SUBSTRING_INDEX(USER(), '@', 1);
-> 'davida'
```

As of MySQL 4.1, `USER()` returns a value in the `utf8` character set, so you should also make sure that the `'@'` string literal is interpreted in that character set:

```
mysql> SELECT SUBSTRING_INDEX(USER(), _utf8'@', 1);
-> 'davida'
```

- `VERSION()`

Returns a string that indicates the MySQL server version. As of MySQL 4.1, the string has the `utf8` character set.

```
mysql> SELECT VERSION();
-> '4.1.22-standard'
```

Note that if your version string ends with `-log` this means that logging is enabled.

## 12.9.4. Miscellaneous Functions

- `DEFAULT(col_name)`

Returns the default value for a table column.

```
mysql> UPDATE t SET i = DEFAULT(i)+1 WHERE id < 100;
```

`DEFAULT()` was added in MySQL 4.1.0.

- `FORMAT(X,D)`

Formats the number `X` to a format like `'#,###,###.##'`, rounded to `D` decimal places, and returns the result as a string. For details, see [Section 12.3, “String Functions”](#).

- `GET_LOCK(str, timeout)`

Tries to obtain a lock with a name given by the string `str`, using a timeout of `timeout` seconds. Returns `1` if the lock was obtained successfully, `0` if the attempt timed out (for example, because another client has previously locked the name), or `NULL` if an error occurred (such as running out of memory or the thread was killed with `mysqladmin kill`). If you have a lock obtained with `GET_LOCK()`, it is released when you execute `RELEASE_LOCK()`, execute a new `GET_LOCK()`, or your connection terminates (either normally or abnormally). Locks obtained with `GET_LOCK()` do not interact with transactions. That is, committing a transaction does not release any such locks

obtained during the transaction.

This function can be used to implement application locks or to simulate record locks. Names are locked on a server-wide basis. If a name has been locked by one client, `GET_LOCK()` blocks any request by another client for a lock with the same name. This allows clients that agree on a given lock name to use the name to perform cooperative advisory locking. But be aware that it also allows a client that is not among the set of cooperating clients to lock a name, either inadvertently or deliberately, and thus prevent any of the cooperating clients from locking that name. One way to reduce the likelihood of this is to use lock names that are database-specific or application-specific. For example, use lock names of the form `db_name.str` or `app_name.str`.

```
mysql> SELECT GET_LOCK('lock1',10);
-> 1
mysql> SELECT IS_FREE_LOCK('lock2');
-> 1
mysql> SELECT GET_LOCK('lock2',10);
-> 1
mysql> SELECT RELEASE_LOCK('lock2');
-> 1
mysql> SELECT RELEASE_LOCK('lock1');
-> NULL
```

The second `RELEASE_LOCK()` call returns `NULL` because the lock 'lock1' was automatically released by the second `GET_LOCK()` call.

Note: If a client attempts to acquire a lock that is already held by another client, it blocks according to the `timeout` argument. If the blocked client terminates, its thread does not die until the lock request times out. This is a known bug.

- `INET_ATON(expr)`

Given the dotted-quad representation of a network address as a string, returns an integer that represents the numeric value of the address. Addresses may be 4- or 8-byte addresses.

```
mysql> SELECT INET_ATON('209.207.224.40');
-> 3520061480
```

The generated number is always in network byte order. For the example just shown, the number is calculated as  $209 \times 256^3 + 207 \times 256^2 + 224 \times 256 + 40$ .

As of MySQL 4.1.2, `INET_ATON()` also understands short-form IP addresses:

```
mysql> SELECT INET_ATON('127.0.0.1'), INET_ATON('127.1');
-> 2130706433, 2130706433
```

**Note:** When storing values generated by `INET_ATON()`, it is recommended that you use an `INT UNSIGNED` column. If you use a (signed) `INT` column, values corresponding to IP addresses for which the first octet is greater than 127 cannot be stored correctly. See [Section 11.2, “Numeric Types”](#).

`INET_ATON()` was added in MySQL 3.23.15.

- `INET_NTOA(expr)`

Given a numeric network address (4 or 8 byte), returns the dotted-quad representation of the address as a string.

```
mysql> SELECT INET_NTOA(3520061480);
-> '209.207.224.40'
```



`INET_NTOA()` was added in MySQL 3.23.15.

- `IS_FREE_LOCK(str)`

Checks whether the lock named `str` is free to use (that is, not locked). Returns `1` if the lock is free (no one is using the lock), `0` if the lock is in use, and `NULL` if an error occurs (such as an incorrect argument).

`IS_FREE_LOCK()` was added in MySQL 4.0.2.

- `IS_USED_LOCK(str)`

Checks whether the lock named `str` is in use (that is, locked). If so, it returns the connection identifier of the client that holds the lock. Otherwise, it returns `NULL`.

`IS_USED_LOCK()` was added in MySQL 4.1.0.

- `MASTER_POS_WAIT(log_name, log_pos[, timeout])`

This function is useful for control of master/slave synchronization. It blocks until the slave has read and applied all updates up to the specified position in the master log. The return value is the number of log events the slave had to wait for to advance to the specified position. The function returns `NULL` if the slave SQL thread is not started, the slave's master information is not initialized, the arguments are incorrect, or an error occurs. It returns `-1` if the timeout has been exceeded. If the slave SQL thread stops while `MASTER_POS_WAIT()` is waiting, the function returns `NULL`. If the slave is past the specified position, the function returns immediately.

If a `timeout` value is specified, `MASTER_POS_WAIT()` stops waiting when `timeout` seconds have elapsed. `timeout` must be greater than 0; a zero or negative `timeout` means no timeout.

`MASTER_POS_WAIT()` was added in MySQL 3.23.32. The `timeout` argument was added in 4.0.10.

- `RELEASE_LOCK(str)`

Releases the lock named by the string `str` that was obtained with `GET_LOCK()`. Returns `1` if the lock was released, `0` if the lock was not established by this thread (in which case the lock is not released), and `NULL` if the named lock did not exist. The lock does not exist if it was never obtained by a call to `GET_LOCK()` or if it has previously been released.

The `DO` statement is convenient to use with `RELEASE_LOCK()`. See [Section 13.2.2, “DO Syntax”](#).

- `UUID()`

Returns a Universal Unique Identifier (UUID) generated according to “DCE 1.1: Remote Procedure Call” (Appendix A) CAE (Common Applications Environment) Specifications published by The Open Group in October 1997 (Document Number C706, <http://www.opengroup.org/public/pubs/catalog/c706.htm>).

A UUID is designed as a number that is globally unique in space and time. Two calls to `UUID()` are expected to generate two different values, even if these calls are performed on two separate computers that are not connected to each other.

A UUID is a 128-bit number represented by a string of five hexadecimal numbers in `aaaaaaa-bbbb-cccc-dddd-eeeeeeeeeeee` format:

- The first three numbers are generated from a timestamp.

- The fourth number preserves temporal uniqueness in case the timestamp value loses monotonicity (for example, due to daylight saving time).
- The fifth number is an IEEE 802 node number that provides spatial uniqueness. A random number is substituted if the latter is not available (for example, because the host computer has no Ethernet card, or we do not know how to find the hardware address of an interface on your operating system). In this case, spatial uniqueness cannot be guaranteed. Nevertheless, a collision should have *very* low probability.

Currently, the MAC address of an interface is taken into account only on FreeBSD and Linux. On other operating systems, MySQL uses a randomly generated 48-bit number.

```
mysql> SELECT UUID();
-> '6ccd780c-baba-1026-9564-0040f4311e29'
```

Note that `UUID()` does not yet work with replication.

`UUID()` was added in MySQL 4.1.2.

- `VALUES(col_name)`

In an `INSERT ... ON DUPLICATE KEY UPDATE` statement, you can use the `VALUES(col_name)` function in the `UPDATE` clause to refer to column values from the `INSERT` portion of the statement. In other words, `VALUES(col_name)` in the `UPDATE` clause refers to the value of `col_name` that would be inserted, had no duplicate-key conflict occurred. This function is especially useful in multiple-row inserts. The `VALUES()` function is meaningful only in `INSERT ... ON DUPLICATE KEY UPDATE` statements and returns `NULL` otherwise. [Section 13.2.4.3, “INSERT ... ON DUPLICATE KEY UPDATE Syntax”](#).

```
mysql> INSERT INTO table (a,b,c) VALUES (1,2,3),(4,5,6)
-> ON DUPLICATE KEY UPDATE c=VALUES(a)+VALUES(b);
```

`VALUES()` was added in MySQL 4.1.1.

## 12.10. Functions and Modifiers for Use with `GROUP BY` Clauses

### 12.10.1. `GROUP BY` (Aggregate) Functions

This section describes group (aggregate) functions that operate on sets of values. Unless otherwise stated, group functions ignore `NULL` values.

If you use a group function in a statement containing no `GROUP BY` clause, it is equivalent to grouping on all rows.

For numeric arguments, the variance, standard deviation, `SUM()`, and `AVG()` functions return a `DOUBLE` value.

The `SUM()` and `AVG()` aggregate functions do not work with temporal values. (They convert the values to numbers, which loses the part after the first non-numeric character.) To work around this problem, you can convert to numeric units, perform the aggregate operation, and convert back to a temporal value. Examples:

```
SELECT SEC_TO_TIME(SUM(TIME_TO_SEC(time_col))) FROM tbl_name;
SELECT FROM_DAYS(SUM(TO_DAYS(date_col))) FROM tbl_name;
```

- Returns the average value of *expr*.

`AVG ( )` returns `NULL` if there were no matching rows.

```
mysql> SELECT student_name, AVG(test_score)
-> FROM student
-> GROUP BY student_name;
```

- `BIT_AND (expr)`

Returns the bitwise `AND` of all bits in *expr*. The calculation is performed with 64-bit (`BIGINT`) precision.

As of MySQL 4.0.17, this function returns `18446744073709551615` if there were no matching rows. (This is the value of an unsigned `BIGINT` value with all bits set to 1.) Before 4.0.17, the function returns `-1` if there were no matching rows.

- `BIT_OR (expr)`

Returns the bitwise `OR` of all bits in *expr*. The calculation is performed with 64-bit (`BIGINT`) precision.

This function returns `0` if there were no matching rows.

- `BIT_XOR (expr)`

Returns the bitwise `XOR` of all bits in *expr*. The calculation is performed with 64-bit (`BIGINT`) precision.

This function returns `0` if there were no matching rows.

This function is available as of MySQL 4.1.1.

- `COUNT (expr)`

Returns a count of the number of non-`NULL` values in the rows retrieved by a `SELECT` statement. The result is a `BIGINT` value.

`COUNT ( )` returns `0` if there were no matching rows.

```
mysql> SELECT student.student_name, COUNT(*)
-> FROM student, course
-> WHERE student.student_id=course.student_id
-> GROUP BY student_name;
```

`COUNT ( * )` is somewhat different in that it returns a count of the number of rows retrieved, whether or not they contain `NULL` values.

`COUNT ( * )` is optimized to return very quickly if the `SELECT` retrieves from one table, no other columns are retrieved, and there is no `WHERE` clause. For example:

```
mysql> SELECT COUNT(*) FROM student;
```

This optimization applies only to `MyISAM` and `ISAM` tables only, because an exact row count is stored for these storage engines and can be accessed very quickly. For transactional storage engines such as `InnoDB` and `BDB`, storing an exact row count is more problematic because multiple transactions may be occurring, each of which may affect the count.

- `COUNT(DISTINCT expr, [expr...])`

Returns a count of the number of different non-NULL values.

`COUNT(DISTINCT)` returns 0 if there were no matching rows.

```
mysql> SELECT COUNT(DISTINCT results) FROM student;
```

In MySQL, you can obtain the number of distinct expression combinations that do not contain NULL by giving a list of expressions. In standard SQL, you would have to do a concatenation of all expressions inside `COUNT(DISTINCT ...)`.

`COUNT(DISTINCT ...)` was added in MySQL 3.23.2.

- `GROUP_CONCAT(expr)`

This function returns a string result with the concatenated non-NULL values from a group. It returns NULL if there are no non-NULL values. The full syntax is as follows:

```
GROUP_CONCAT([DISTINCT] expr [,expr ...]
             [ORDER BY {unsigned_integer | col_name | expr}
             [ASC | DESC] [,col_name ...])
             [SEPARATOR str_val])
```

```
mysql> SELECT student_name,
->        GROUP_CONCAT(test_score)
->        FROM student
->        GROUP BY student_name;
```

Or:

```
mysql> SELECT student_name,
->        GROUP_CONCAT(DISTINCT test_score
->                      ORDER BY test_score DESC SEPARATOR ' ')
->        FROM student
->        GROUP BY student_name;
```

In MySQL, you can get the concatenated values of expression combinations. You can eliminate duplicate values by using `DISTINCT`. If you want to sort values in the result, you should use `ORDER BY` clause. To sort in reverse order, add the `DESC` (descending) keyword to the name of the column you are sorting by in the `ORDER BY` clause. The default is ascending order; this may be specified explicitly using the `ASC` keyword. `SEPARATOR` is followed by the string value that should be inserted between values of result. The default is a comma (','). You can eliminate the separator altogether by specifying `SEPARATOR ''`.

You can set a maximum allowed length with the `group_concat_max_len` system variable. (The default value is 1024.) The syntax to do this at runtime is as follows, where `val` is an unsigned integer:

```
SET [SESSION | GLOBAL] group_concat_max_len = val;
```

If a maximum length has been set, the result is truncated to this maximum length.

Beginning with MySQL 4.1.19, the type returned by `GROUP_CONCAT()` is always `VARCHAR` unless `group_concat_max_len` is greater than 512, in which case, it returns a `BLOB`. (Previously, it returned a `BLOB` with `group_concat_max_len` greater than 512 only if the query included an `ORDER BY` clause.)

`GROUP_CONCAT()` was added in MySQL 4.1.

Note: Before MySQL 4.1.6, there are some small limitations with `GROUP_CONCAT()` for `BLOB` and `TEXT` values when it comes to using `DISTINCT` together with `ORDER BY`. To work around this limitation, use `MID(expr, 1, 255)` instead.

See also `CONCAT()` and `CONCAT_WS()`: [Section 12.3, “String Functions”](#).

- `MIN(expr), MAX(expr)`

Returns the minimum or maximum value of `expr`. `MIN()` and `MAX()` may take a string argument; in such cases they return the minimum or maximum string value. See [Section 7.4.5, “How MySQL Uses Indexes”](#).

`MIN()` and `MAX()` return `NULL` if there were no matching rows.

```
mysql> SELECT student_name, MIN(test_score), MAX(test_score)
->      FROM student
->      GROUP BY student_name;
```

For `MIN()`, `MAX()`, and other aggregate functions, MySQL currently compares `ENUM` and `SET` columns by their string value rather than by the string's relative position in the set. This differs from how `ORDER BY` compares them. This is expected to be rectified in a future MySQL release.

- `STD(expr) STDDEV(expr)`

Returns the population standard deviation of `expr`. This is an extension to standard SQL. The `STDDEV()` form of this function is provided for compatibility with Oracle.

These functions return `NULL` if there were no matching rows.

- `SUM(expr)`

Returns the sum of `expr`. If the return set has no rows, `SUM()` returns `NULL`.

`SUM()` returns `NULL` if there were no matching rows.

- `VARIANCE(expr)`

Returns the population standard variance of `expr`. This is an extension to standard SQL, available in MySQL 4.1 or later.

`VARIANCE()` returns `NULL` if there were no matching rows.

## 12.10.2. GROUP BY Modifiers

As of MySQL 4.1.1, the `GROUP BY` clause allows a `WITH ROLLUP` modifier that causes extra rows to be added to the summary output. These rows represent higher-level (or super-aggregate) summary operations. `ROLLUP` thus allows you to answer questions at multiple levels of analysis with a single query. It can be used, for example, to provide support for OLAP (Online Analytical Processing) operations.

Suppose that a table named `sales` has `year`, `country`, `product`, and `profit` columns for recording sales profitability:

```
CREATE TABLE sales
(
  year      INT NOT NULL,
  country   VARCHAR(20) NOT NULL,
  product   VARCHAR(32) NOT NULL,
  profit    INT
);
```

The table's contents can be summarized per year with a simple `GROUP BY` like this:

```
mysql> SELECT year, SUM(profit) FROM sales GROUP BY year;
```

year	SUM(profit)
2000	4525
2001	3010

This output shows the total profit for each year, but if you also want to determine the total profit summed over all years, you must add up the individual values yourself or run an additional query.

Or you can use `ROLLUP`, which provides both levels of analysis with a single query. Adding a `WITH ROLLUP` modifier to the `GROUP BY` clause causes the query to produce another row that shows the grand total over all year values:

```
mysql> SELECT year, SUM(profit) FROM sales GROUP BY year WITH ROLLUP;
```

year	SUM(profit)
2000	4525
2001	3010
NULL	7535

The grand total super-aggregate line is identified by the value `NULL` in the `year` column.

`ROLLUP` has a more complex effect when there are multiple `GROUP BY` columns. In this case, each time there is a “break” (change in value) in any but the last grouping column, the query produces an extra super-aggregate summary row.

For example, without `ROLLUP`, a summary on the `sales` table based on `year`, `country`, and `product` might look like this:

```
mysql> SELECT year, country, product, SUM(profit)
-> FROM sales
-> GROUP BY year, country, product;
```

year	country	product	SUM(profit)
2000	Finland	Computer	1500
2000	Finland	Phone	100
2000	India	Calculator	150
2000	India	Computer	1200
2000	USA	Calculator	75
2000	USA	Computer	1500
2001	Finland	Phone	10
2001	USA	Calculator	50
2001	USA	Computer	2700
2001	USA	TV	250

The output indicates summary values only at the year/country/product level of analysis. When `ROLLUP` is added, the query produces several extra rows:

```
mysql> SELECT year, country, product, SUM(profit)
-> FROM sales
-> GROUP BY year, country, product WITH ROLLUP;
```

year	country	product	SUM(profit)
2000	Finland	Computer	1500
2000	Finland	Phone	100
2000	Finland	NULL	1600
2000	India	Calculator	150
2000	India	Computer	1200

2000	India	NULL	1350
2000	USA	Calculator	75
2000	USA	Computer	1500
2000	USA	NULL	1575
2000	NULL	NULL	4525
2001	Finland	Phone	10
2001	Finland	NULL	10
2001	USA	Calculator	50
2001	USA	Computer	2700
2001	USA	TV	250
2001	USA	NULL	3000
2001	NULL	NULL	3010
NULL	NULL	NULL	7535

For this query, adding `ROLLUP` causes the output to include summary information at four levels of analysis, not just one. Here's how to interpret the `ROLLUP` output:

- Following each set of product rows for a given year and country, an extra summary row is produced showing the total for all products. These rows have the `product` column set to `NULL`.
- Following each set of rows for a given year, an extra summary row is produced showing the total for all countries and products. These rows have the `country` and `products` columns set to `NULL`.
- Finally, following all other rows, an extra summary row is produced showing the grand total for all years, countries, and products. This row has the `year`, `country`, and `products` columns set to `NULL`.

### Other Considerations When using `ROLLUP`

The following items list some behaviors specific to the MySQL implementation of `ROLLUP`:

When you use `ROLLUP`, you cannot also use an `ORDER BY` clause to sort the results. In other words, `ROLLUP` and `ORDER BY` are mutually exclusive. However, you still have some control over sort order. `GROUP BY` in MySQL sorts results, and you can use explicit `ASC` and `DESC` keywords with columns named in the `GROUP BY` list to specify sort order for individual columns. (The higher-level summary rows added by `ROLLUP` still appear after the rows from which they are calculated, regardless of the sort order.)

`LIMIT` can be used to restrict the number of rows returned to the client. `LIMIT` is applied after `ROLLUP`, so the limit applies against the extra rows added by `ROLLUP`. For example:

```
mysql> SELECT year, country, product, SUM(profit)
-> FROM sales
-> GROUP BY year, country, product WITH ROLLUP
-> LIMIT 5;
```

year	country	product	SUM(profit)
2000	Finland	Computer	1500
2000	Finland	Phone	100
2000	Finland	NULL	1600
2000	India	Calculator	150
2000	India	Computer	1200

Using `LIMIT` with `ROLLUP` may produce results that are more difficult to interpret, because you have less context for understanding the super-aggregate rows.

The `NULL` indicators in each super-aggregate row are produced when the row is sent to the client. The server looks at the columns named in the `GROUP BY` clause following the leftmost one that has changed value. For any column in the result set with a name that is a lexical match to any of those names, its value is set to `NULL`. (If you specify grouping columns by column number, the server identifies which columns to set to `NULL` by number.)

Because the `NULL` values in the super-aggregate rows are placed into the result set at such a late stage in query processing, you cannot test them as `NULL` values within the query itself. For example, you cannot add `HAVING product IS NULL` to the query to eliminate from the output all but the super-aggregate rows.

On the other hand, the `NULL` values do appear as `NULL` on the client side and can be tested as such using any MySQL client programming interface.

### 12.10.3. GROUP BY and HAVING with Hidden Fields

MySQL extends the use of `GROUP BY` so that you can use non-aggregated columns or calculations in the `SELECT` list that do not appear in the `GROUP BY` clause. You can use this feature to get better performance by avoiding unnecessary column sorting and grouping. For example, you do not need to group on `customer.name` in the following query:

```
SELECT order.custid, customer.name, MAX(payments)
FROM order, customer
WHERE order.custid = customer.custid
GROUP BY order.custid;
```

In standard SQL, you would have to add `customer.name` to the `GROUP BY` clause. In MySQL, the name is redundant.

Do *not* use this feature if the columns you omit from the `GROUP BY` part are not constant in the group. The server is free to return any value from the group, so the results are indeterminate unless all values are the same.

A similar MySQL extension applies to the `HAVING` clause. The SQL standard does not allow the `HAVING` clause to name any column that is not found in the `GROUP BY` clause if it is not enclosed in an aggregate function. MySQL allows the use of such columns to simplify calculations. This extension assumes that the non-grouped columns will have the same group-wise values. Otherwise, the result is indeterminate.

If the `ONLY_FULL_GROUP_BY` SQL mode is enabled, the MySQL extension to `GROUP BY` does not apply to the `SELECT`. That is, columns not named in the `GROUP BY` clause cannot be used in the `SELECT` list if not used in an aggregate function.

The select list extension also applies to `ORDER BY`. That is, you can use non-aggregated columns or calculations in the `ORDER BY` clause that do not appear in the `GROUP BY` clause. This extension does not apply if the `ONLY_FULL_GROUP_BY` SQL mode is enabled.

In some cases, you can use `MIN()` and `MAX()` to obtain a specific column value even if it isn't unique. The following gives the value of `column` from the row containing the smallest value in the `sort` column:

```
SUBSTR(MIN(CONCAT(RPAD(sort,6,' '),column)),7)
```

See [Section 3.6.4, “The Rows Holding the Group-wise Maximum of a Certain Field”](#).

Note that if you are using MySQL 3.22 (or earlier) or if you are trying to follow standard SQL, you cannot use expressions in `GROUP BY` or `ORDER BY` clauses. You can work around this limitation by using an alias for the expression:

```
mysql> SELECT id,FLOOR(value/100) AS val
-> FROM tbl_name
-> GROUP BY id, val ORDER BY val;
```

In MySQL 3.23 and up, aliases are unnecessary. You can use expressions in `GROUP BY` and `ORDER`



BY clauses. For example:

```
mysql> SELECT id, FLOOR(value/100) FROM tbl_name ORDER BY RAND();
```

Note that if you are using MySQL 3.22 (or earlier) or if you are trying to follow standard SQL, you can't use expressions in `GROUP BY` clauses. You can work around this limitation by using an alias for the expression:

```
SELECT id, FLOOR(value/100) AS val
FROM tbl_name
GROUP BY id, val;
```

In MySQL 3.23 and up, aliases are unnecessary and MySQL does allow expressions in `GROUP BY` clauses. For example:

```
SELECT id, FLOOR(value/100)
FROM tbl_name
GROUP BY id, FLOOR(value/100);
```

Before MySQL 3.23, MySQL also requires use of aliases to refer to expressions in `ORDER BY` clauses.

---

# Chapter 13. SQL Statement Syntax

This chapter describes the syntax for the SQL statements supported in MySQL versions 4.1 and earlier.

## 13.1. Data Definition Statements

### 13.1.1. ALTER DATABASE Syntax

```
ALTER DATABASE [db_name]  
    alter_specification [alter_specification] ...  
  
alter_specification:  
    [DEFAULT] CHARACTER SET charset_name  
    | [DEFAULT] COLLATE collation_name
```

**ALTER DATABASE** enables you to change the overall characteristics of a database. These characteristics are stored in the `db.opt` file in the database directory. To use **ALTER DATABASE**, you need the **ALTER** privilege on the database.

The **CHARACTER SET** clause changes the default database character set. The **COLLATE** clause changes the default database collation. [Chapter 10, Character Set Support](#), discusses character set and collation names.

**ALTER DATABASE** was added in MySQL 4.1.1. Beginning with MySQL 4.1.8, the database name can be omitted, in which case the statement applies to the default database.

### 13.1.2. ALTER TABLE Syntax

```
ALTER [IGNORE] TABLE tbl_name  
    alter_specification [, alter_specification] ...  
  
alter_specification:  
    ADD [COLUMN] column_definition [FIRST | AFTER col_name ]  
    | ADD [COLUMN] (column_definition,...)  
    | ADD {INDEX|KEY} [index_name] [index_type] (index_col_name,...)  
    | ADD [CONSTRAINT [symbol]]  
        PRIMARY KEY [index_type] (index_col_name,...)  
    | ADD [CONSTRAINT [symbol]]  
        UNIQUE [INDEX|KEY] [index_name] [index_type] (index_col_name,...)  
    | ADD [FULLTEXT|SPATIAL] [INDEX|KEY] [index_name] (index_col_name,...)  
    | ADD [CONSTRAINT [symbol]]  
        FOREIGN KEY [index_name] (index_col_name,...)  
        [reference_definition]  
    | ALTER [COLUMN] col_name {SET DEFAULT literal | DROP DEFAULT}  
    | CHANGE [COLUMN] old_col_name column_definition  
        [FIRST|AFTER col_name]  
    | MODIFY [COLUMN] column_definition [FIRST | AFTER col_name]  
    | DROP [COLUMN] col_name  
    | DROP PRIMARY KEY  
    | DROP {INDEX|KEY} index_name  
    | DROP FOREIGN KEY fk_symbol  
    | DISABLE KEYS  
    | ENABLE KEYS  
    | RENAME [TO] new_tbl_name  
    | ORDER BY col_name  
    | CONVERT TO CHARACTER SET charset_name [COLLATE collation_name]  
    | [DEFAULT] CHARACTER SET charset_name [COLLATE collation_name]  
    | DISCARD TABLESPACE  
    | IMPORT TABLESPACE  
    | table_option ...  
  
index_col_name:  
    col_name [(length)] [ASC | DESC]  
  
index_type:  
    USING {BTREE | HASH}
```

`ALTER TABLE` enables you to change the structure of an existing table. For example, you can add or delete columns, create or destroy indexes, change the type of existing columns, or rename columns or the table itself. You can also change the comment for the table and type of the table.

The syntax for many of the allowable alterations is similar to clauses of the `CREATE TABLE` statement. This includes *table\_option* modifications, for options such as `ENGINE`, `AUTO_INCREMENT`, and `AVG_ROW_LENGTH`. (However, `ALTER TABLE` ignores the `DATA DIRECTORY` and `INDEX DIRECTORY` table options.) [Section 13.1.5, “CREATE TABLE Syntax”](#), lists all table options.

Some operations may result in warnings if attempted on a table for which the storage engine does not support the operation. In MySQL 4.1 and up, these warnings can be displayed with `SHOW WARNINGS`. See [Section 13.5.4.21, “SHOW WARNINGS Syntax”](#).

If you use `ALTER TABLE` to change a column specification but `DESCRIBE tbl_name` indicates that your column was not changed, it is possible that MySQL ignored your modification for one of the reasons described in [Section 13.1.5.1, “Silent Column Specification Changes”](#). For example, if you try to change a `VARCHAR` column to `CHAR`, MySQL still uses `VARCHAR` if the table contains other variable-length columns.

In most cases, `ALTER TABLE` works by making a temporary copy of the original table. The alteration is performed on the copy, and then the original table is deleted and the new one is renamed. While `ALTER TABLE` is executing, the original table is readable by other clients. Updates and writes to the table are stalled until the new table is ready, and then are automatically redirected to the new table without any failed updates.

If you use `ALTER TABLE tbl_name RENAME TO new_tbl_name` without any other options, MySQL simply renames any files that correspond to the table `tbl_name`. There is no need to create a temporary table. (You can also use the `RENAME TABLE` statement to rename tables. See [Section 13.1.9, “RENAME TABLE Syntax”](#).)

If you use any option to `ALTER TABLE` other than `RENAME`, MySQL always creates a temporary table, even if the data wouldn't strictly need to be copied (such as when you change the name of a column). For `MyISAM` tables, you can speed up the index re-creation operation (which is the slowest part of the alteration process) by setting the `myisam_sort_buffer_size` system variable to a high value.

- To use `ALTER TABLE`, you need `ALTER`, `INSERT`, and `CREATE` privileges for the table.
- `IGNORE` is a MySQL extension to standard SQL. It controls how `ALTER TABLE` works if there are duplicates on unique keys in the new table or if warnings occur when strict mode is enabled. If `IGNORE` is not specified, the copy is aborted and rolled back if duplicate-key errors occur. If `IGNORE` is specified, only the first row is used of rows with duplicates on a unique key, The other conflicting rows are deleted. Incorrect values are truncated to the closest matching acceptable value.
- You can issue multiple `ADD`, `ALTER`, `DROP`, and `CHANGE` clauses in a single `ALTER TABLE` statement, separated by commas. This is a MySQL extension to standard SQL, which allows only one of each clause per `ALTER TABLE` statement. For example, to drop multiple columns in a single statement, do this:

```
ALTER TABLE t2 DROP COLUMN c, DROP COLUMN d;
```

- `CHANGE col_name`, `DROP col_name`, and `DROP INDEX` are MySQL extensions to standard SQL.
- `MODIFY` is an Oracle extension to `ALTER TABLE`.
- The word `COLUMN` is optional and can be omitted.

- *column\_definition* clauses use the same syntax for `ADD` and `CHANGE` as for `CREATE TABLE`. Note that this syntax includes the column name, not just its data type. See [Section 13.1.5, “CREATE TABLE Syntax”](#).
- You can rename a column using a `CHANGE old_col_name column_definition` clause. To do so, specify the old and new column names and the type that the column currently has. For example, to rename an `INTEGER` column from `a` to `b`, you can do this:

```
ALTER TABLE t1 CHANGE a b INTEGER;
```

If you want to change a column's type but not the name, `CHANGE` syntax still requires an old and new column name, even if they are the same. For example:

```
ALTER TABLE t1 CHANGE b b BIGINT NOT NULL;
```

However, as of MySQL 3.22.16a, you can also use `MODIFY` to change a column's type without renaming it:

```
ALTER TABLE t1 MODIFY b BIGINT NOT NULL;
```

- If you use `CHANGE` or `MODIFY` to shorten a column for which an index exists on the column, and the resulting column length is less than the index length, MySQL shortens the index automatically.
- When you change a data type using `CHANGE` or `MODIFY`, MySQL tries to convert existing column values to the new type as well as possible.
- In MySQL 3.22 or later, to add a column at a specific position within a table row, use `FIRST` or `AFTER col_name`. The default is to add the column last. From MySQL 4.0.1 on, you can also use `FIRST` and `AFTER` in `CHANGE` or `MODIFY` operations.
- `ALTER ... SET DEFAULT` or `ALTER ... DROP DEFAULT` specify a new default value for a column or remove the old default value, respectively. If the old default is removed and the column can be `NULL`, the new default is `NULL`. If the column cannot be `NULL`, MySQL assigns a default value, as described in [Section 11.1.4, “Data Type Default Values”](#).
- `DROP INDEX` removes an index. This is a MySQL extension to standard SQL. See [Section 13.1.7, “DROP INDEX Syntax”](#).
- If columns are dropped from a table, the columns are also removed from any index of which they are a part. If all columns that make up an index are dropped, the index is dropped as well.
- If a table contains only one column, the column cannot be dropped. If what you intend is to remove the table, use `DROP TABLE` instead.
- `DROP PRIMARY KEY` drops the primary index. (Prior to MySQL 4.1.2, if no primary index exists, `DROP PRIMARY KEY` drops the first `UNIQUE` index in the table. MySQL marks the first `UNIQUE` key as the `PRIMARY KEY` if no `PRIMARY KEY` was specified explicitly.)

If you add a `UNIQUE INDEX` or `PRIMARY KEY` to a table, it is stored before any non-unique index so that MySQL can detect duplicate keys as early as possible.

- From MySQL 4.1.0 on, some storage engines allow you to specify an index type when creating an index. The syntax for the *index\_type* specifier is `USING type_name`. For details about `USING`, see [Section 13.1.4, “CREATE INDEX Syntax”](#).
- `ORDER BY` enables you to create the new table with the rows in a specific order. Note that the table does not remain in this order after inserts and deletes. This option is useful primarily when you know that you are mostly to query the rows in a certain order most of the time. By using this option after

major changes to the table, you might be able to get higher performance. In some cases, it might make sorting easier for MySQL if the table is in order by the column that you want to order it by later.

- If you use `ALTER TABLE` on a `MyISAM` table, all non-unique indexes are created in a separate batch (as for `REPAIR TABLE`). This should make `ALTER TABLE` much faster when you have many indexes.

As of MySQL 4.0, this feature can be activated explicitly. `ALTER TABLE ... DISABLE KEYS` tells MySQL to stop updating non-unique indexes for a `MyISAM` table. `ALTER TABLE ... ENABLE KEYS` then should be used to re-create missing indexes. MySQL does this with a special algorithm that is much faster than inserting keys one by one, so disabling keys before performing bulk insert operations should give a considerable speedup. Using `ALTER TABLE ... DISABLE KEYS` requires the `INDEX` privilege in addition to the privileges mentioned earlier.

- The `FOREIGN KEY` and `REFERENCES` clauses are supported by the `InnoDB` storage engine, which implements `ADD [CONSTRAINT [symbol]] FOREIGN KEY (...) REFERENCES ... (...)`. See [Section 14.2.7.4, “FOREIGN KEY Constraints”](#). For other storage engines, the clauses are parsed but ignored. The `CHECK` clause is parsed but ignored by all storage engines. See [Section 13.1.5, “CREATE TABLE Syntax”](#). The reason for accepting but ignoring syntax clauses is for compatibility, to make it easier to port code from other SQL servers, and to run applications that create tables with references. See [Section 1.9.5, “MySQL Differences from Standard SQL”](#).

You cannot add a foreign key and drop a foreign key in separate clauses of a single `ALTER TABLE` statement. You must use separate statements.

- Starting from MySQL 4.0.13, `InnoDB` supports the use of `ALTER TABLE` to drop foreign keys:

```
ALTER TABLE tbl_name DROP FOREIGN KEY fk_symbol;
```

You cannot add a foreign key and drop a foreign key in separate clauses of a single `ALTER TABLE` statement. You must use separate statements.

For more information, see [Section 14.2.7.4, “FOREIGN KEY Constraints”](#).

- Pending `INSERT DELAYED` statements are lost if a table is write locked and `ALTER TABLE` is used to modify the table structure.
- From MySQL 4.1.2 on, if you want to change the table default character set and all character columns (`CHAR`, `VARCHAR`, `TEXT`) to a new character set, use a statement like this:

```
ALTER TABLE tbl_name CONVERT TO CHARACTER SET charset_name;
```

This is useful, for example, after upgrading from MySQL 4.0.x to 4.1.x. See [Section 10.9, “Upgrading Character Sets from MySQL 4.0”](#).

**Warning:** The preceding operation converts column values between the character sets. This is *not* what you want if you have a column in one character set (like `latin1`) but the stored values actually use some other, incompatible character set (like `utf8`). In this case, you have to do the following for each such column:

```
ALTER TABLE t1 CHANGE c1 c1 BLOB;
ALTER TABLE t1 CHANGE c1 c1 TEXT CHARACTER SET utf8;
```

The reason this works is that there is no conversion when you convert to or from `BLOB` columns.

If you specify `CONVERT TO CHARACTER SET binary`, the `CHAR`, `VARCHAR`, and `TEXT` columns are converted to their corresponding binary string types (`BINARY`, `VARBINARY`, `BLOB`).

This means that the columns no longer will have a character set and a subsequent `CONVERT TO` operation will not apply to them.

To change only the *default* character set for a table, use this statement:

```
ALTER TABLE tbl_name DEFAULT CHARACTER SET charset_name;
```

The word `DEFAULT` is optional. The default character set is the character set that is used if you do not specify the character set for a new column which you add to a table (for example, with `ALTER TABLE ... ADD column`).

**Warning:** From MySQL 4.1.2 and up, `ALTER TABLE ... DEFAULT CHARACTER SET` and `ALTER TABLE ... CHARACTER SET` are equivalent and change only the default table character set. In MySQL 4.1 releases before 4.1.2, `ALTER TABLE ... DEFAULT CHARACTER SET` changes the default character set, but `ALTER TABLE ... CHARACTER SET` (without `DEFAULT`) changes the default character set *and also converts all columns to the new character set*.

- For an `InnoDB` table that is created with its own tablespace in an `.ibd` file, that file can be discarded and imported. To discard the `.ibd` file, use this statement:

```
ALTER TABLE tbl_name DISCARD TABLESPACE;
```

This deletes the current `.ibd` file, so be sure that you have a backup first. Attempting to access the table while the tablespace file is discarded results in an error.

To import the backup `.ibd` file back into the table, copy it into the database directory, and then issue this statement:

```
ALTER TABLE tbl_name IMPORT TABLESPACE;
```

See [Section 14.2.4.1, “Using Per-Table Tablespaces”](#).

With the `mysql_info()` C API function, you can find out how many rows were copied, and (when `IGNORE` is used) how many rows were deleted due to duplication of unique key values. See [Section 17.2.3.33, “mysql\\_info\(\)”](#).

Here are some examples that show uses of `ALTER TABLE`. Begin with a table `t1` that is created as shown here:

```
CREATE TABLE t1 (a INTEGER,b CHAR(10));
```

To rename the table from `t1` to `t2`:

```
ALTER TABLE t1 RENAME t2;
```

To change column `a` from `INTEGER` to `TINYINT NOT NULL` (leaving the name the same), and to change column `b` from `CHAR(10)` to `CHAR(20)` as well as renaming it from `b` to `c`:

```
ALTER TABLE t2 MODIFY a TINYINT NOT NULL, CHANGE b c CHAR(20);
```

To add a new `TIMESTAMP` column named `d`:

```
ALTER TABLE t2 ADD d TIMESTAMP;
```

To add indexes on column `d` and on column `a`:

```
ALTER TABLE t2 ADD INDEX (d), ADD INDEX (a);
```

To remove column `c`:

```
ALTER TABLE t2 DROP COLUMN c;
```

To add a new `AUTO_INCREMENT` integer column named `c`:

```
ALTER TABLE t2 ADD c INT UNSIGNED NOT NULL AUTO_INCREMENT,
  ADD PRIMARY KEY (c);
```

Note that we indexed `c` (as a `PRIMARY KEY`), because `AUTO_INCREMENT` columns must be indexed, and also that we declare `c` as `NOT NULL`, because primary key columns cannot be `NULL`.

When you add an `AUTO_INCREMENT` column, column values are filled in with sequence numbers for you automatically. For `MyISAM` tables, you can set the first sequence number by executing `SET INSERT_ID=value` before `ALTER TABLE` or by using the `AUTO_INCREMENT=value` table option. See [Section 13.5.3, “SET Syntax”](#).

From MySQL 4.1.2, you can use the `ALTER TABLE ... AUTO_INCREMENT=value` table option for `InnoDB` tables to set the sequence number for new rows if the value is greater than the maximum value in the `AUTO_INCREMENT` column. *If the value is less than the current maximum value in the column, no error message is given and the current sequence value is not changed.*

With `MyISAM` tables, if you do not change the `AUTO_INCREMENT` column, the sequence number is not affected. If you drop an `AUTO_INCREMENT` column and then add another `AUTO_INCREMENT` column, the numbers are resequenced beginning with 1.

When replication is used, adding an `AUTO_INCREMENT` column to a table might not produce the same ordering of the rows on the slave and the master. This occurs because the order in which the rows are numbered depends on the specific storage engine used for the table and the order in which the rows were inserted. If it is important to have the same order on the master and slave, the rows must be ordered before assigning an `AUTO_INCREMENT` number. Assuming that you want to add an `AUTO_INCREMENT` column to the table `t1`, the following statements produce a new table `t2` identical to `t1` but with an `AUTO_INCREMENT` column:

```
CREATE TABLE t2 (id INT AUTO_INCREMENT PRIMARY KEY)
SELECT * FROM t1 ORDER BY col1, col2;
```

This assumes that the table `t1` has columns `col1` and `col2`.

This set of statements will also produce a new table `t2` identical to `t1`, with the addition of an `AUTO_INCREMENT` column:

```
CREATE TABLE t2 LIKE t1;
ALTER TABLE T2 ADD id INT AUTO_INCREMENT PRIMARY KEY;
INSERT INTO t2 SELECT * FROM t1 ORDER BY col1, col2;
```

**Important:** To guarantee the same ordering on both master and slave, *all* columns of `t1` must be referenced in the `ORDER BY` clause.

Regardless of the method used to create and populate the copy having the `AUTO_INCREMENT` column, the final step is to drop the original table and then rename the copy:

```
DROP t1;
ALTER TABLE t2 RENAME t1;
```

See also [Section A.7.1, “Problems with ALTER TABLE”](#).

### 13.1.3. CREATE DATABASE Syntax

```
CREATE DATABASE [IF NOT EXISTS] db_name
  [create_specification [create_specification] ...]

create_specification:
  [DEFAULT] CHARACTER SET charset_name
  | [DEFAULT] COLLATE collation_name
```

**CREATE DATABASE** creates a database with the given name. To use this statement, you need the **CREATE** privilege for the database.

An error occurs if the database exists and you did not specify **IF NOT EXISTS**.

As of MySQL 4.1.1, *create\_specification* options specify database characteristics. Database characteristics are stored in the `db.opt` file in the database directory. The **CHARACTER SET** clause specifies the default database character set. The **COLLATE** clause specifies the default database collation. [Chapter 10, Character Set Support](#), discusses character set and collation names.

A database in MySQL is implemented as a directory containing files that correspond to tables in the database. Because there are no tables in a database when it is initially created, the **CREATE DATABASE** statement only creates a directory under the MySQL data directory (and the `db.opt` file, for MySQL 4.1.1 and up). Rules for allowable database names are given in [Section 9.2, “Database, Table, Index, Column, and Alias Names”](#).

If you manually create a directory under the data directory (for example, with `mkdir`), the server considers it a database directory and it shows up in the output of **SHOW DATABASES**.

You can also use the `mysqladmin` program to create databases. See [Section 8.9, “mysqladmin — Client for Administering a MySQL Server”](#).

### 13.1.4. CREATE INDEX Syntax

```
CREATE [UNIQUE|FULLTEXT|SPATIAL] INDEX index_name
  [index_type]
  ON tbl_name (index_col_name, ...)

index_col_name:
  col_name [(length)] [ASC | DESC]

index_type:
  USING {BTREE | HASH}
```

In MySQL 3.22 or later, **CREATE INDEX** is mapped to an **ALTER TABLE** statement to create indexes. See [Section 13.1.2, “ALTER TABLE Syntax”](#). The **CREATE INDEX** statement does not do anything prior to MySQL 3.22. For more information about indexes, see [Section 7.4.5, “How MySQL Uses Indexes”](#).

Normally, you create all indexes on a table at the time the table itself is created with **CREATE TABLE**. See [Section 13.1.5, “CREATE TABLE Syntax”](#). **CREATE INDEX** enables you to add indexes to existing tables.

A column list of the form `(col1,col2,...)` creates a multiple-column index. Index values are formed by concatenating the values of the given columns.

For **CHAR**, **VARCHAR**, **BINARY**, and **VARBINARY** columns, indexes can be created that use only the leading part of column values, using `col_name(length)` syntax to specify an index prefix length. **BLOB** and **TEXT** columns also can be indexed, but a prefix length *must* be given. Prefix lengths are given in characters for non-binary string types and in bytes for binary string types. That is, index entries consist of the first *length* characters of each column value for **CHAR**, **VARCHAR**, and **TEXT** columns, and the first *length* bytes of each column value for **BINARY**, **VARBINARY**, and **BLOB** columns.



The statement shown here creates an index using the first 10 characters of the `name` column:

```
CREATE INDEX part_of_name ON customer (name(10));
```

If names in the column usually differ in the first 10 characters, this index should not be much slower than an index created from the entire `name` column. Also, using partial columns for indexes can make the index file much smaller, which could save a lot of disk space and might also speed up `INSERT` operations.

Prefixes can be up to 1000 bytes long (767 bytes for `InnoDB` tables). (Before MySQL 4.1.2, the limit is 255 bytes for all tables.) Note that prefix limits are measured in bytes, whereas the prefix length in `CREATE INDEX` statements is interpreted as number of characters for non-binary data types (`CHAR`, `VARCHAR`, `TEXT`). Take this into account when specifying a prefix length for a column that uses a multi-byte character set.

A `UNIQUE` index creates a constraint such that all values in the index must be distinct. An error occurs if you try to add a new row with a key value that matches an existing row. This constraint does not apply to `NULL` values except for the `BDB` storage engine. For other engines, a `UNIQUE` index allows multiple `NULL` values for columns that can contain `NULL`.

`FULLTEXT` indexes are supported only for `MyISAM` tables and can include only `CHAR`, `VARCHAR`, and `TEXT` columns. Indexing always happens over the entire column; partial indexing is not supported and any prefix length is ignored if specified. See [Section 12.7, “Full-Text Search Functions”](#), for details of operation. `FULLTEXT` indexes are available in MySQL 3.23.23 or later.

`SPATIAL` indexes are supported only for `MyISAM` tables and can include only spatial columns that are defined as `NOT NULL`. `SPATIAL` indexes are available in MySQL 4.1 or later. [Chapter 16, \*Spatial Extensions\*](#), describes the spatial data types.

You can add an index on a column that can have `NULL` values only if you are using MySQL 3.23.2 or newer and are using the `MyISAM`, `InnoDB`, or `BDB` storage engine. This is also true for `MEMORY` tables as of MySQL 4.0.2. You can only add an index on a `BLOB` or `TEXT` column if you are using MySQL 3.23.2 or newer and are using the `MyISAM` or `BDB` storage engine, or MySQL 4.0.14 or newer and the `InnoDB` storage engine.

An `index_col_name` specification can end with `ASC` or `DESC`. These keywords are allowed for future extensions for specifying ascending or descending index value storage. Currently, they are parsed but ignored; index values are always stored in ascending order.

From MySQL 4.1.0 on, some storage engines allow you to specify an index type when creating an index. The allowable index type values supported by different storage engines are shown in the following table. Where multiple index types are listed, the first one is the default when no index type specifier is given.

Storage Engine	Allowable Index Types
<code>MyISAM</code>	<code>BTREE</code>
<code>InnoDB</code>	<code>BTREE</code>
<code>MEMORY/HEAP</code>	<code>HASH, BTREE</code>

If you specify an index type that is not legal for a given storage engine, but there is another index type available that the engine can use without affecting query results, the engine uses the available type.

Examples:

```
CREATE TABLE lookup (id INT) ENGINE = MEMORY;
CREATE INDEX id_index USING BTREE ON lookup (id);
```

`TYPE type_name` is recognized as a synonym for `USING type_name`. However, `USING` is the preferred form.

### 13.1.5. CREATE TABLE Syntax

```
CREATE [TEMPORARY] TABLE [IF NOT EXISTS] tbl_name
  (create_definition,...)
  [table_option ...]
```

Or:

```
CREATE [TEMPORARY] TABLE [IF NOT EXISTS] tbl_name
  [(create_definition,...)]
  [table_option ...]
  select_statement
```

Or:

```
CREATE [TEMPORARY] TABLE [IF NOT EXISTS] tbl_name
  { LIKE old_tbl_name | (LIKE old_tbl_name) }

create_definition:
  column_definition
  | [CONSTRAINT [symbol]] PRIMARY KEY [index_type] (index_col_name,...)
  | {INDEX|KEY} [index_name] [index_type] (index_col_name,...)
  | [CONSTRAINT [symbol]] UNIQUE [INDEX|KEY]
    [index_name] [index_type] (index_col_name,...)
  | {FULLTEXT|SPATIAL} [INDEX|KEY] [index_name] (index_col_name,...)
  | [CONSTRAINT [symbol]] FOREIGN KEY
    [index_name] (index_col_name,...) [reference_definition]
  | CHECK (expr)

column_definition:
  col_name data_type [NOT NULL | NULL] [DEFAULT default_value]
  [AUTO_INCREMENT] [UNIQUE [KEY] | [PRIMARY] KEY]
  [COMMENT 'string'] [reference_definition]

data_type:
  TINYINT[(length)] [UNSIGNED] [ZEROFILL]
  SMALLINT[(length)] [UNSIGNED] [ZEROFILL]
  MEDIUMINT[(length)] [UNSIGNED] [ZEROFILL]
  INT[(length)] [UNSIGNED] [ZEROFILL]
  INTEGER[(length)] [UNSIGNED] [ZEROFILL]
  BIGINT[(length)] [UNSIGNED] [ZEROFILL]
  REAL[(length,decimals)] [UNSIGNED] [ZEROFILL]
  DOUBLE[(length,decimals)] [UNSIGNED] [ZEROFILL]
  FLOAT[(length,decimals)] [UNSIGNED] [ZEROFILL]
  DECIMAL(length,decimals) [UNSIGNED] [ZEROFILL]
  NUMERIC(length,decimals) [UNSIGNED] [ZEROFILL]
  DATE
  TIME
  TIMESTAMP
  DATETIME
  YEAR
  CHAR(length)
  [CHARACTER SET charset_name] [COLLATE collation_name]
  VARCHAR(length)
  [CHARACTER SET charset_name] [COLLATE collation_name]
  BINARY(length)
  VARBINARY(length)
  TINYBLOB
  BLOB
  MEDIUMBLOB
  LONGBLOB
  TINYTEXT [BINARY]
  [CHARACTER SET charset_name] [COLLATE collation_name]
  TEXT [BINARY]
  [CHARACTER SET charset_name] [COLLATE collation_name]
  MEDIUMTEXT [BINARY]
  [CHARACTER SET charset_name] [COLLATE collation_name]
  LONGTEXT [BINARY]
  [CHARACTER SET charset_name] [COLLATE collation_name]
  ENUM(value1,value2,value3,...)
  [CHARACTER SET charset_name] [COLLATE collation_name]
```

```

| SET(value1,value2,value3,...)
| [CHARACTER SET charset_name] [COLLATE collation_name]
| spatial_type

index_col_name:
  col_name [(length)] [ASC | DESC]

index_type:
  USING {BTREE | HASH}

reference_definition:
  REFERENCES tbl_name [(index_col_name,...)]
  [MATCH FULL | MATCH PARTIAL | MATCH SIMPLE]
  [ON DELETE reference_option]
  [ON UPDATE reference_option]

reference_option:
  RESTRICT | CASCADE | SET NULL | NO ACTION

table_option:
  {ENGINE|TYPE} = engine_name
  AUTO_INCREMENT = value
  AVG_ROW_LENGTH = value
  [DEFAULT] CHARACTER SET charset_name
  CHECKSUM = {0 | 1}
  COLLATE collation_name
  COMMENT = 'string'
  DATA DIRECTORY = 'absolute path to directory'
  DELAY_KEY_WRITE = {0 | 1}
  INDEX DIRECTORY = 'absolute path to directory'
  INSERT_METHOD = { NO | FIRST | LAST }
  MAX_ROWS = value
  MIN_ROWS = value
  PACK_KEYS = {0 | 1 | DEFAULT}
  PASSWORD = 'string'
  RAID_TYPE = { 1 | STRIPED | RAID0 }
  RAID_CHUNKS = value
  RAID_CHUNKSIZE = value
  ROW_FORMAT = {DEFAULT|DYNAMIC|FIXED|COMPRESSED|REDUNDANT|COMPACT}
  UNION = (tbl_name[,tbl_name]...)

select_statement:
  [IGNORE | REPLACE] [AS] SELECT ... (Some legal select statement)

```

**CREATE TABLE** creates a table with the given name. You must have the **CREATE** privilege for the table.

Rules for allowable table names are given in [Section 9.2, “Database, Table, Index, Column, and Alias Names”](#). By default, the table is created in the default database. An error occurs if the table exists, if there is no default database, or if the database does not exist.

In MySQL 3.22 or later, the table name can be specified as *db\_name.tbl\_name* to create the table in a specific database. This works regardless of whether there is a default database, assuming that the database exists. If you use quoted identifiers, quote the database and table names separately. For example, write ``mydb`.`mytbl``, not ``mydb.mytbl``.

From MySQL 3.23 on, you can use the **TEMPORARY** keyword when creating a table. A **TEMPORARY** table is visible only to the current connection, and is dropped automatically when the connection is closed. This means that two different connections can use the same temporary table name without conflicting with each other or with an existing non-**TEMPORARY** table of the same name. (The existing table is hidden until the temporary table is dropped.) From MySQL 4.0.2 on, to create temporary tables, you must have the **CREATE TEMPORARY TABLES** privilege.

In MySQL 3.23 or later, the keywords **IF NOT EXISTS** prevent an error from occurring if the table exists. However, there is no verification that the existing table has a structure identical to that indicated by the **CREATE TABLE** statement. *Note:* If you use **IF NOT EXISTS** in a **CREATE TABLE ... SELECT** statement, any rows selected by the **SELECT** part are inserted regardless of whether the table already exists.

MySQL represents each table by an `.frm` table format (definition) file in the database directory. The

storage engine for the table might create other files as well. In the case of `MyISAM` tables, the storage engine creates data and index files. Thus, for each `MyISAM` table `tbl_name`, there are three disk files:

File	Purpose
<code>tbl_name.frm</code>	Table format (definition) file
<code>tbl_name.MYD</code>	Data file
<code>tbl_name.MYI</code>	Index file

Chapter 14, *Storage Engines and Table Types*, describes what files each storage engine creates to represent tables.

`data_type` represents the data type in a column definition. `spatial_type` represents a spatial data type. For general information on the properties of data types other than the spatial types, see Chapter 11, *Data Types*. For information about spatial data types, see Chapter 16, *Spatial Extensions*.

Some attributes do not apply to all data types. `AUTO_INCREMENT` applies only to integer types. `DEFAULT` does not apply to the `BLOB` or `TEXT` types.

- If neither `NULL` nor `NOT NULL` is specified, the column is treated as though `NULL` had been specified.
- An integer column can have the additional attribute `AUTO_INCREMENT`. When you insert a value of `NULL` (recommended) or `0` into an indexed `AUTO_INCREMENT` column, the column is set to the next sequence value. Typically this is `value+1`, where `value` is the largest value for the column currently in the table. `AUTO_INCREMENT` sequences begin with `1`.

To retrieve an `AUTO_INCREMENT` value after inserting a row, use the `LAST_INSERT_ID()` SQL function or the `mysql_insert_id()` C API function. See Section 12.9.3, “Information Functions”, and Section 17.2.3.35, “`mysql_insert_id()`”.

As of MySQL 4.1.1, if the `NO_AUTO_VALUE_ON_ZERO` SQL mode is enabled, you can store `0` in `AUTO_INCREMENT` columns as `0` without generating a new sequence value. See Section 5.2.5, “SQL Modes”.

**Note:** There can be only one `AUTO_INCREMENT` column per table, it must be indexed, and it cannot have a `DEFAULT` value. As of MySQL 3.23, an `AUTO_INCREMENT` column works properly only if it contains only positive values. Inserting a negative number is regarded as inserting a very large positive number. This is done to avoid precision problems when numbers “wrap” over from positive to negative and also to ensure that you do not accidentally get an `AUTO_INCREMENT` column that contains `0`.

For `MyISAM` and `BDB` tables, you can specify an `AUTO_INCREMENT` secondary column in a multiple-column key. See Section 3.6.9, “Using `AUTO_INCREMENT`”.

To make MySQL compatible with some ODBC applications, you can find the `AUTO_INCREMENT` value for the last inserted row with the following query:

```
SELECT * FROM tbl_name WHERE auto_col IS NULL
```

For information about `InnoDB` and `AUTO_INCREMENT`, see Section 14.2.7.3, “How `AUTO_INCREMENT` Columns Work in `InnoDB`”.

- From MySQL 4.1.0 on, the attribute `SERIAL` is an alias for `BIGINT UNSIGNED NOT NULL AUTO_INCREMENT UNIQUE`.
- As of MySQL 4.1, character data types (`CHAR`, `VARCHAR`, `TEXT`) can include `CHARACTER SET`

and `COLLATE` attributes to specify the character set and collation for the column. For details, see [Chapter 10, Character Set Support](#). `CHARSET` is a synonym for `CHARACTER SET`. Example:

```
CREATE TABLE t (c CHAR(20) CHARACTER SET utf8 COLLATE utf8_bin);
```

Also as of 4.1, MySQL interprets length specifications in character column definitions in characters. (Earlier versions interpret them in bytes.) Lengths for `BINARY` and `VARBINARY` are in bytes.

- `NULL` values are handled differently for `TIMESTAMP` columns than for other column types. Before MySQL 4.1.6, you cannot store a literal `NULL` in a `TIMESTAMP` column; setting the column to `NULL` sets it to the current date and time. Because `TIMESTAMP` columns behave this way, the `NULL` and `NOT NULL` attributes do not apply in the normal way and are ignored if you specify them. On the other hand, to make it easier for MySQL clients to use `TIMESTAMP` columns, the server reports that such columns can be assigned `NULL` values (which is true), even though `TIMESTAMP` never actually contains a `NULL` value. You can see this when you use `DESCRIBE tbl_name` to get a description of your table.

Note that setting a `TIMESTAMP` column to `0` is not the same as setting it to `NULL`, because `0` is a valid `TIMESTAMP` value.

- The `DEFAULT` clause specifies a default value for a column. With one exception, the default value must be a constant; it cannot be a function or an expression. This means, for example, that you cannot set the default for a date column to be the value of a function such as `NOW()` or `CURRENT_DATE`. The exception is that you can specify `CURRENT_TIMESTAMP` as the default for a `TIMESTAMP` column as of MySQL 4.1.2. See [Section 11.3.1.2, “TIMESTAMP Properties as of MySQL 4.1”](#).

If a column definition includes no explicit `DEFAULT` value, MySQL determines the default value as described in [Section 11.1.4, “Data Type Default Values”](#).

`BLOB` and `TEXT` columns cannot be assigned a default value.

- A comment for a column can be specified with the `COMMENT` option. The comment is displayed by the `SHOW CREATE TABLE` and `SHOW FULL COLUMNS` statements. This option is operational as of MySQL 4.1. (It is allowed but ignored in earlier versions.)
- `KEY` is normally a synonym for `INDEX`. From MySQL 4.1, the key attribute `PRIMARY KEY` can also be specified as just `KEY` when given in a column definition. This was implemented for compatibility with other database systems.
- A `UNIQUE` index creates a constraint such that all values in the index must be distinct. An error occurs if you try to add a new row with a key value that matches an existing row. This constraint does not apply to `NULL` values except for the `BDB` storage engine. For other engines, a `UNIQUE` index allows multiple `NULL` values for columns that can contain `NULL`.
- A `PRIMARY KEY` is a unique index where all key columns must be defined as `NOT NULL`. If they are not explicitly declared as `NOT NULL`, MySQL declares them so implicitly (and silently). A table can have only one `PRIMARY KEY`. If you do not have a `PRIMARY KEY` and an application asks for the `PRIMARY KEY` in your tables, MySQL returns the first `UNIQUE` index that has no `NULL` columns as the `PRIMARY KEY`.

In `InnoDB` tables, having a long `PRIMARY KEY` wastes a lot of space. (See [Section 14.2.14, “InnoDB Table and Index Structures”](#).)

- In the created table, a `PRIMARY KEY` is placed first, followed by all `UNIQUE` indexes, and then the non-unique indexes. This helps the MySQL optimizer to prioritize which index to use and also more quickly to detect duplicated `UNIQUE` keys.

- A **PRIMARY KEY** can be a multiple-column index. However, you cannot create a multiple-column index using the **PRIMARY KEY** key attribute in a column specification. Doing so only marks that single column as primary. You must use a separate **PRIMARY KEY(*index\_col\_name*, ...)** clause.
- If a **PRIMARY KEY** or **UNIQUE** index consists of only one column that has an integer type, you can also refer to the column as **\_rowid** in **SELECT** statements (new in MySQL 3.23.11).
- In MySQL, the name of a **PRIMARY KEY** is **PRIMARY**. For other indexes, if you do not assign a name, the index is assigned the same name as the first indexed column, with an optional suffix (**\_2**, **\_3**, ...) to make it unique. You can see index names for a table using **SHOW INDEX FROM *tbl\_name***. See [Section 13.5.4.11, “SHOW INDEX Syntax”](#).
- From MySQL 4.1.0 on, some storage engines allow you to specify an index type when creating an index. The syntax for the *index\_type* specifier is **USING *type\_name***.

Example:

```
CREATE TABLE lookup
(id INT, INDEX USING BTREE (id))
ENGINE = MEMORY;
```

For details about **USING**, see [Section 13.1.4, “CREATE INDEX Syntax”](#).

For more information about indexes, see [Section 7.4.5, “How MySQL Uses Indexes”](#).

- Only the **MyISAM**, **InnoDB**, **BDB**, and (as of MySQL 4.0.2) **MEMORY** storage engines support indexes on columns that can have **NULL** values. In other cases, you must declare indexed columns as **NOT NULL** or an error results.
- For **CHAR**, **VARCHAR**, **BINARY**, and **VARBINARY** columns, indexes can be created that use only the leading part of column values, using *col\_name(length)* syntax to specify an index prefix length. **BLOB** and **TEXT** columns also can be indexed, but a prefix length *must* be given. Prefix lengths are given in characters for non-binary string types and in bytes for binary string types. That is, index entries consist of the first *length* characters of each column value for **CHAR**, **VARCHAR**, and **TEXT** columns, and the first *length* bytes of each column value for **BINARY**, **VARBINARY**, and **BLOB** columns. Indexing only a prefix of column values like this can make the index file much smaller. See [Section 7.4.3, “Column Indexes”](#).

Only the **MyISAM** and (as of MySQL 4.0.14) **InnoDB** storage engines support indexing on **BLOB** and **TEXT** columns. For example:

```
CREATE TABLE test (blob_col BLOB, INDEX(blob_col(10)));
```

Prefixes can be up to 1000 bytes long (767 bytes for **InnoDB** tables). (Before MySQL 4.1.2, the limit is 255 bytes for all tables.) Note that prefix limits are measured in bytes, whereas the prefix length in **CREATE TABLE** statements is interpreted as number of characters for non-binary data types (**CHAR**, **VARCHAR**, **TEXT**). Take this into account when specifying a prefix length for a column that uses a multi-byte character set.

- An *index\_col\_name* specification can end with **ASC** or **DESC**. These keywords are allowed for future extensions for specifying ascending or descending index value storage. Currently, they are parsed but ignored; index values are always stored in ascending order.
- When you use **ORDER BY** or **GROUP BY** on a **TEXT** or **BLOB** column in a **SELECT**, the server sorts values using only the initial number of bytes indicated by the *max\_sort\_length* system variable. See [Section 11.4.3, “The BLOB and TEXT Types”](#).

- In MySQL 3.23.23 or later, you can create special `FULLTEXT` indexes, which are used for full-text searches. Only the `MyISAM` table type supports `FULLTEXT` indexes. They can be created only from `CHAR`, `VARCHAR`, and `TEXT` columns. Indexing always happens over the entire column; partial indexing is not supported and any prefix length is ignored if specified. See [Section 12.7, “Full-Text Search Functions”](#), for details of operation.
- In MySQL 4.1 or later, you can create `SPATIAL` indexes on spatial data types. Spatial types are supported only for `MyISAM` tables and indexed columns must be declared as `NOT NULL`. See [Chapter 16, \*Spatial Extensions\*](#).
- In MySQL 3.23.44 or later, `InnoDB` tables support checking of foreign key constraints. See [Section 14.2, “The InnoDB Storage Engine”](#). Note that the `FOREIGN KEY` syntax in `InnoDB` is more restrictive than the syntax presented for the `CREATE TABLE` statement at the beginning of this section: The columns of the referenced table must always be explicitly named. `InnoDB` supports both `ON DELETE` and `ON UPDATE` actions on foreign keys as of MySQL 3.23.50 and 4.0.8, respectively. For the precise syntax, see [Section 14.2.7.4, “FOREIGN KEY Constraints”](#).

For other storage engines, MySQL Server parses and ignores the `FOREIGN KEY` and `REFERENCES` syntax in `CREATE TABLE` statements. The `CHECK` clause is parsed but ignored by all storage engines. See [Section 1.9.5.5, “Foreign Keys”](#).

- For `MyISAM` and `ISAM` tables, each `NULL` column takes one bit extra, rounded up to the nearest byte. The maximum row length in bytes can be calculated as follows:

```
row length = 1
             + (sum of column lengths)
             + (number of NULL columns + delete_flag + 7)/8
             + (number of variable-length columns)
```

`delete_flag` is 1 for tables with static row format. Static tables use a bit in the row record for a flag that indicates whether the row has been deleted. `delete_flag` is 0 for dynamic tables because the flag is stored in the dynamic row header.

These calculations do not apply for `InnoDB` tables, for which storage size is no different for `NULL` columns than for `NOT NULL` columns.

The `table_option` part of the `CREATE TABLE` syntax can be used in MySQL 3.23 and above. The `=` that separates an option name and its value is optional as of MySQL 4.1.

The `ENGINE` and `TYPE` options specify the storage engine for the table. `ENGINE` was added in MySQL 4.0.18 (for 4.0) and 4.1.2 (for 4.1). It is the preferred option name as of those versions, and `TYPE` has become deprecated. `TYPE` is supported throughout the 4.x series, but likely will be removed in the future.

The `ENGINE` and `TYPE` table options take the storage engine names shown in the following table.

Storage Engine	Description
<code>ARCHIVE</code>	The archiving storage engine. See <a href="#">Section 14.7, “The ARCHIVE Storage Engine”</a> .
<code>BDB</code>	Transaction-safe tables with page locking. Also known as <code>BerkeleyDB</code> . See <a href="#">Section 14.5, “The BDB (BerkeleyDB) Storage Engine”</a> .
<code>CSV</code>	Tables that store rows in comma-separated values format. See <a href="#">Section 14.8, “The CSV Storage Engine”</a> .
<code>EXAMPLE</code>	An example engine. See <a href="#">Section 14.6, “The EXAMPLE Storage Engine”</a> .
<code>HEAP</code>	The data for this table is stored only in memory. See <a href="#">Section 14.4, “The MEMORY (HEAP) Storage Engine”</a> .

ISAM	The original MySQL storage engine. See <a href="#">Section 14.10, “The ISAM Storage Engine”</a> .
InnoDB	Transaction-safe tables with row locking and foreign keys. See <a href="#">Section 14.2, “The InnoDB Storage Engine”</a> .
MEMORY	An alias for HEAP. (Actually, as of MySQL 4.1, MEMORY is the preferred term.)
MERGE	A collection of MyISAM tables used as one table. Also known as MRG_MyISAM. See <a href="#">Section 14.3, “The MERGE Storage Engine”</a> .
MyISAM	The binary portable storage engine that is the improved replacement for ISAM. See <a href="#">Section 14.1, “The MyISAM Storage Engine”</a> .
NDBCLUSTER	Clustered, fault-tolerant, memory-based tables. Also known as NDB. See <a href="#">Chapter 15, MySQL Cluster</a> .

If a storage engine is specified that is not available, MySQL uses the default engine instead. Normally, this is MyISAM. For example, if a table definition includes the `ENGINE=BDB` option but the MySQL server does not support BDB tables, the table is created as a MyISAM table. This makes it possible to have a replication setup where you have transactional tables on the master but tables created on the slave are non-transactional (to get more speed). In MySQL 4.1.1, a warning occurs if the storage engine specification is not honored.

The other table options are used to optimize the behavior of the table. In most cases, you do not have to specify any of them. These options apply to all storage engines unless otherwise indicated. Options that do not apply to a given storage engine may be accepted and remembered as part of the table definition. Such options then apply if you later use `ALTER TABLE` to convert the table to use a different storage engine.

- `AUTO_INCREMENT`

The initial `AUTO_INCREMENT` value for the table. This works for MyISAM only, for MEMORY as of MySQL 4.1.0, and for InnoDB as of MySQL 4.1.2. To set the first auto-increment value for engines that do not support the `AUTO_INCREMENT` table option, insert a “dummy” row with a value one less than the desired value after creating the table, and then delete the dummy row.

For engines that support the `AUTO_INCREMENT` table option in `CREATE TABLE` statements, you can also use `ALTER TABLE tbl_name AUTO_INCREMENT = N` to reset the `AUTO_INCREMENT` value.

- `AVG_ROW_LENGTH`

An approximation of the average row length for your table. You need to set this only for large tables with variable-size rows.

When you create a MyISAM table, MySQL uses the product of the `MAX_ROWS` and `AVG_ROW_LENGTH` options to decide how big the resulting table is. If you do not specify either option, the maximum size for a table is 4GB. (If your operating system does not support files that large, table sizes are constrained by the operating system limit.) If you want to keep down the pointer sizes to make the index smaller and faster and you do not really need big files, you can decrease the default pointer size by setting the `myisam_data_pointer_size` system variable, which was added in MySQL 4.1.2. (See [Section 5.2.2, “System Variables”](#).) If you want all your tables to be able to grow above the default limit and are willing to have your tables slightly slower and larger than necessary, you may increase the default pointer size by setting this variable.

- `[DEFAULT] CHARACTER SET`



Specify a default character set for the table. `CHARSET` is a synonym for `CHARACTER SET`.

- `CHECKSUM`

Set this to 1 if you want MySQL to maintain a live checksum for all rows (that is, a checksum that MySQL updates automatically as the table changes). This makes the table a little slower to update, but also makes it easier to find corrupted tables. The `CHECKSUM TABLE` statement reports the checksum. (`MyISAM` only.)

- `COLLATE`

Specify a default collation for the table.

- `COMMENT`

A comment for the table, up to 60 characters long.

- `DATA DIRECTORY, INDEX DIRECTORY`

By using `DATA DIRECTORY='directory'` or `INDEX DIRECTORY='directory'` you can specify where the `MyISAM` storage engine should put a table's data file and index file. The directory must be the full pathname to the directory, not a relative path.

These options work only for `MyISAM` tables from MySQL 4.0 on, when you are not using the `--skip-symbolic-links` option. Your operating system must also have a working, thread-safe `realpath()` call. See [Section 7.6.1.2, “Using Symbolic Links for Tables on Unix”](#), for more complete information.

- `DELAY_KEY_WRITE`

Set this to 1 if you want to delay key updates for the table until the table is closed. See the description of the `delay_key_write` system variable in [Section 5.2.2, “System Variables”](#). (`MyISAM` only.)

- `INSERT_METHOD`

If you want to insert data into a `MERGE` table, you must specify with `INSERT_METHOD` the table into which the row should be inserted. `INSERT_METHOD` is an option useful for `MERGE` tables only. Use a value of `FIRST` or `LAST` to have inserts go to the first or last table, or a value of `NO` to prevent inserts. This option was introduced in MySQL 4.0.0. See [Section 14.3, “The MERGE Storage Engine”](#).

- `MAX_ROWS`

The maximum number of rows you plan to store in the table. This is not a hard limit, but rather a hint to the storage engine that the table must be able to store at least this many rows.

- `MIN_ROWS`

The minimum number of rows you plan to store in the table.

- `PACK_KEYS`

`PACK_KEYS` takes effect only with `MyISAM` tables. Set this option to 1 if you want to have smaller indexes. This usually makes updates slower and reads faster. Setting the option to 0 disables all packing of keys. Setting it to `DEFAULT` tells the storage engine to pack only long `CHAR` or `VARCHAR` columns.

If you do not use `PACK_KEYS`, the default is to pack only strings, but not numbers. If you use `PACK_KEYS=1`, numbers are packed as well.

When packing binary number keys, MySQL uses prefix compression:

- Every key needs one extra byte to indicate how many bytes of the previous key are the same for the next key.
- The pointer to the row is stored in high-byte-first order directly after the key, to improve compression.

This means that if you have many equal keys on two consecutive rows, all following “same” keys usually only take two bytes (including the pointer to the row). Compare this to the ordinary case where the following keys takes `storage_size_for_key + pointer_size` (where the pointer size is usually 4). Conversely, you get a significant benefit from prefix compression only if you have many numbers that are the same. If all keys are totally different, you use one byte more per key, if the key is not a key that can have `NULL` values. (In this case, the packed key length is stored in the same byte that is used to mark if a key is `NULL`.)

- `PASSWORD`

Encrypt the `.frm` file with a password. This option does nothing in the standard MySQL version.

- `RAID_TYPE`

The `RAID_TYPE` option can help you to exceed the 2GB/4GB limit for the `MyISAM` data file (not the index file) on operating systems that do not support big files. This option is unnecessary and not recommended for filesystems that support big files.

You can get more speed from the I/O bottleneck by putting `RAID` directories on different physical disks. The only allowed `RAID_TYPE` is `STRIPED`. `1` and `RAID0` are aliases for `STRIPED`.

If you specify the `RAID_TYPE` option for a `MyISAM` table, specify the `RAID_CHUNKS` and `RAID_CHUNKSIZE` options as well. The maximum `RAID_CHUNKS` value is 255. `MyISAM` creates `RAID_CHUNKS` subdirectories named `00`, `01`, `02`, ... `09`, `0a`, `0b`, ... in the database directory. In each of these directories, `MyISAM` creates a file `tbl_name.MYD`. When writing data to the data file, the `RAID` handler maps the first `RAID_CHUNKSIZE*1024` bytes to the first file, the next `RAID_CHUNKSIZE*1024` bytes to the next file, and so on.

`RAID_TYPE` works on any operating system, as long as you have built MySQL with the `-with-raid` option to `configure`. To determine whether a server supports `RAID` tables, use `SHOW VARIABLES LIKE 'have_raid'` to see whether the variable value is `YES`.

- `ROW_FORMAT`

Defines how the rows should be stored. Currently, this option works only with `MyISAM` tables. The option value can be `FIXED` or `DYNAMIC` for static or variable-length row format. `myisampack` sets the type to `COMPRESSED`. See [Section 14.1.3, “MyISAM Table Storage Formats”](#).

- `UNION`

`UNION` is used when you want to access a collection of identical `MyISAM` tables as one. This works only with `MERGE` tables. See [Section 14.3, “The MERGE Storage Engine”](#).

In MySQL 4.1, you must have `SELECT`, `UPDATE`, and `DELETE` privileges for the tables you map to a `MERGE` table. (*Note:* Originally, all tables used had to be in the same database as the `MERGE` table itself. This restriction has been lifted as of MySQL 4.1.1.

As of MySQL 3.23, you can create one table from another by adding a `SELECT` statement at the end of the `CREATE TABLE` statement:

```
CREATE TABLE new_tbl SELECT * FROM orig_tbl;
```

MySQL creates new columns for all elements in the `SELECT`. For example:

```
mysql> CREATE TABLE test (a INT NOT NULL AUTO_INCREMENT,
-> PRIMARY KEY (a), KEY(b))
-> TYPE=MyISAM SELECT b,c FROM test2;
```

This creates a `MyISAM` table with three columns, `a`, `b`, and `c`. Notice that the columns from the `SELECT` statement are appended to the right side of the table, not overlapped onto it. Take the following example:

```
mysql> SELECT * FROM foo;
+----+
| n |
+----+
| 1 |
+----+

mysql> CREATE TABLE bar (m INT) SELECT n FROM foo;
Query OK, 1 row affected (0.02 sec)
Records: 1 Duplicates: 0 Warnings: 0

mysql> SELECT * FROM bar;
+-----+-----+
| m | n |
+-----+-----+
| NULL | 1 |
+-----+-----+
1 row in set (0.00 sec)
```

For each row in table `foo`, a row is inserted in `bar` with the values from `foo` and default values for the new columns.

In a table resulting from `CREATE TABLE ... SELECT`, columns named only in the `CREATE TABLE` part come first. Columns named in both parts or only in the `SELECT` part come after that. The data type of `SELECT` columns can be overridden by also specifying the column in the `CREATE TABLE` part.

If any errors occur while copying the data to the table, it is automatically dropped and not created.

`CREATE TABLE ... SELECT` does not automatically create any indexes for you. This is done intentionally to make the statement as flexible as possible. If you want to have indexes in the created table, you should specify these before the `SELECT` statement:

```
mysql> CREATE TABLE bar (UNIQUE (n)) SELECT n FROM foo;
```

Some conversion of data types might occur. For example, the `AUTO_INCREMENT` attribute is not preserved, and `VARCHAR` columns can become `CHAR` columns.

When creating a table with `CREATE ... SELECT`, make sure to alias any function calls or expressions in the query. If you do not, the `CREATE` statement might fail or result in undesirable column names.

```
CREATE TABLE artists_and_works
SELECT artist.name, COUNT(work.artist_id) AS number_of_works
FROM artist LEFT JOIN work ON artist.id = work.artist_id
GROUP BY artist.id;
```

As of MySQL 4.1, you can explicitly specify the data type for a generated column:

```
CREATE TABLE foo (a TINYINT NOT NULL) SELECT b+1 AS a FROM bar;
```

---

In MySQL 4.1, you can also use [LIKE](#) to create an empty table based on the definition of another table, including any column attributes and indexes the original table has:

```
CREATE TABLE new_tbl LIKE orig_tbl;
```

[CREATE TABLE ... LIKE](#) does not preserve any [DATA DIRECTORY](#) or [INDEX DIRECTORY](#) table options that were specified for the original table, or any foreign key definitions.

You can precede the [SELECT](#) by [IGNORE](#) or [REPLACE](#) to indicate how to handle rows that duplicate unique key values. With [IGNORE](#), new rows that duplicate an existing row on a unique key value are discarded. With [REPLACE](#), new rows replace rows that have the same unique key value. If neither [IGNORE](#) nor [REPLACE](#) is specified, duplicate unique key values result in an error.

To ensure that the update log or binary log can be used to re-create the original tables, MySQL does not allow concurrent inserts for [CREATE TABLE ... SELECT](#) statements.

### 13.1.5.1. Silent Column Specification Changes

In some cases, MySQL silently changes column specifications from those given in a [CREATE TABLE](#) or [ALTER TABLE](#) statement. These might be changes to a data type, to attributes associated with a data type, or to an index specification.

- [VARCHAR](#) columns with a length less than four are changed to [CHAR](#).
- If any column in a table has a variable length, the entire row becomes variable-length as a result. Therefore, if a table contains any variable-length columns ([VARCHAR](#), [TEXT](#), or [BLOB](#)), all [CHAR](#) columns longer than three characters are changed to [VARCHAR](#) columns. This does not affect how you use the columns in any way; in MySQL, [VARCHAR](#) is just a different way to store characters. MySQL performs this conversion because it saves space and makes table operations faster. See [Chapter 14, Storage Engines and Table Types](#).
- From MySQL 4.1.0 onward, a [CHAR](#) or [VARCHAR](#) column with a length specification greater than 255 is converted to the smallest [TEXT](#) type that can hold values of the given length. For example, [VARCHAR\(500\)](#) is converted to [TEXT](#), and [VARCHAR\(200000\)](#) is converted to [MEDIUMTEXT](#). Note that this conversion results in a change in behavior with regard to treatment of trailing spaces.

Similar conversions occur for [BINARY](#) and [VARBINARY](#), except that they are converted to a [BLOB](#) type.

- For a specification of [DECIMAL\(M,D\)](#), if *M* is not larger than *D*, it is adjusted upward. For example, [DECIMAL\(10,10\)](#) becomes [DECIMAL\(11,10\)](#).

Other silent column specification changes include changes to attribute or index specifications:

- [TIMESTAMP](#) display sizes are discarded from MySQL 4.1 on, due to changes made to the [TIMESTAMP](#) data type in that version. Before MySQL 4.1, [TIMESTAMP](#) display sizes must be even and in the range from 2 to 14. If you specify a display size of 0 or greater than 14, the size is coerced to 14. Odd-valued sizes in the range from 1 to 13 are coerced to the next higher even number.
- Before MySQL 4.1.6, you cannot store a literal [NULL](#) in a [TIMESTAMP](#) column; setting it to [NULL](#) sets it to the current date and time. Because [TIMESTAMP](#) columns behave this way, the [NULL](#) and [NOT NULL](#) attributes do not apply in the normal way and are ignored if you specify them. [DESCRIBE tbl\\_name](#) always reports that a [TIMESTAMP](#) column can be assigned [NULL](#) values.

- Columns that are part of a `PRIMARY KEY` are made `NOT NULL` even if not declared that way.
- Starting from MySQL 3.23.51, trailing spaces are automatically deleted from `ENUM` and `SET` member values when the table is created.
- MySQL maps certain data types used by other SQL database vendors to MySQL types. See [Section 11.7, “Using Data Types from Other Database Engines”](#).
- If you include a `USING` clause to specify an index type that is not legal for a given storage engine, but there is another index type available that the engine can use without affecting query results, the engine uses the available type.

To see whether MySQL used a data type other than the one you specified, issue a `DESCRIBE` or `SHOW CREATE TABLE` statement after creating or altering the table.

Certain other data type changes can occur if you compress a table using `myisampack`. See [Section 14.1.3.3, “Compressed Table Characteristics”](#).

## 13.1.6. DROP DATABASE Syntax

```
DROP DATABASE [IF EXISTS] db_name
```

`DROP DATABASE` drops all tables in the database and deletes the database. Be *very* careful with this statement! To use `DROP DATABASE`, you need the `DROP` privilege on the database.

In MySQL 3.22 or later, you can use the keywords `IF EXISTS` to prevent an error from occurring if the database does not exist.

If you use `DROP DATABASE` on a symbolically linked database, both the link and the original database are deleted.

As of MySQL 4.1.2, `DROP DATABASE` returns the number of tables that were removed. This corresponds to the number of `.frm` files removed.

The `DROP DATABASE` statement removes from the given database directory those files and directories that MySQL itself may create during normal operation:

- All files with these extensions:

<code>.BAK</code>	<code>.DAT</code>	<code>.HSH</code>	<code>.ISD</code>
<code>.ISM</code>	<code>.MRG</code>	<code>.MYD</code>	<code>.MYI</code>
<code>.db</code>	<code>.frm</code>	<code>.ibd</code>	<code>.ndb</code>

- All subdirectories with names that consist of two hex digits `00-ff`. These are subdirectories used for `RAID` tables. (These directories are not removed in versions of MySQL after 4.1, where support for `RAID` tables is removed. You should convert any existing `RAID` tables and remove these directories manually before upgrading to later MySQL versions.)
- The `db.opt` file, if it exists.

If other files or directories remain in the database directory after MySQL removes those just listed, the database directory cannot be removed. In this case, you must remove any remaining files or directories manually and issue the `DROP DATABASE` statement again.

You can also drop databases with `mysqladmin`. See [Section 8.9, “mysqladmin — Client for Administering a MySQL Server”](#).

### 13.1.7. DROP INDEX Syntax

```
DROP INDEX index_name ON tbl_name
```

`DROP INDEX` drops the index named *index\_name* from the table *tbl\_name*. In MySQL 3.22 or later, `DROP INDEX` is mapped to an `ALTER TABLE` statement to drop the index. See [Section 13.1.2, “ALTER TABLE Syntax”](#). `DROP INDEX` does not do anything prior to MySQL 3.22.

### 13.1.8. DROP TABLE Syntax

```
DROP [TEMPORARY] TABLE [IF EXISTS]
tbl_name [, tbl_name] ...
[RESTRICT | CASCADE]
```

`DROP TABLE` removes one or more tables. You must have the `DROP` privilege for each table. All table data and the table definition are *removed*, so *be careful* with this statement! If any of the tables named in the argument list do not exist, MySQL returns an error indicating by name which non-existing tables it was unable to drop, but it also drops all of the tables in the list that do exist.

In MySQL 3.22 or later, you can use the keywords `IF EXISTS` to prevent an error from occurring for tables that do not exist. As of MySQL 4.1, a `NOTE` is generated for each non-existent table when using `IF EXISTS`. See [Section 13.5.4.21, “SHOW WARNINGS Syntax”](#).

`RESTRICT` and `CASCADE` are allowed to make porting easier. For the moment, they do nothing.

**Note:** `DROP TABLE` automatically commits the current active transaction, unless you are using MySQL 4.1 or higher and the `TEMPORARY` keyword.

The `TEMPORARY` keyword is ignored in MySQL 4.0. As of 4.1, it has the following effect:

- The statement drops only `TEMPORARY` tables.
- The statement does not end an ongoing transaction.
- No access rights are checked. (A `TEMPORARY` table is visible only to the client that created it, so no check is necessary.)

Using `TEMPORARY` is a good way to ensure that you do not accidentally drop a non-`TEMPORARY` table.

### 13.1.9. RENAME TABLE Syntax

```
RENAME TABLE tbl_name TO new_tbl_name
[, tbl_name2 TO new_tbl_name2] ...
```

This statement renames one or more tables. It was added in MySQL 3.23.23.

The rename operation is done atomically, which means that no other thread can access any of the tables while the rename is running. For example, if you have an existing table `old_table`, you can create another table `new_table` that has the same structure but is empty, and then replace the existing table with the empty one as follows (assuming that `backup_table` does not already exist):

```
CREATE TABLE new_table (...);
RENAME TABLE old_table TO backup_table, new_table TO old_table;
```

If the statement renames more than one table, renaming operations are done from left to right. If you want to swap two table names, you can do so like this (assuming that `tmp_table` does not already exist):

```
RENAME TABLE old_table TO tmp_table,
             new_table TO old_table,
             tmp_table TO new_table;
```

As long as two databases are on the same filesystem, you can use `RENAME TABLE` to move a table from one database to another:

```
RENAME TABLE current_db.tbl_name TO other_db.tbl_name;
```

When you execute `RENAME`, you cannot have any locked tables or active transactions. You must also have the `ALTER` and `DROP` privileges on the original table, and the `CREATE` and `INSERT` privileges on the new table.

If MySQL encounters any errors in a multiple-table rename, it does a reverse rename for all renamed tables to return everything to its original state.

## 13.2. Data Manipulation Statements

### 13.2.1. DELETE Syntax

Single-table syntax:

```
DELETE [LOW_PRIORITY] [QUICK] [IGNORE] FROM tbl_name
      [WHERE where_condition]
      [ORDER BY ...]
      [LIMIT row_count]
```

Multiple-table syntax:

```
DELETE [LOW_PRIORITY] [QUICK] [IGNORE]
      tbl_name[*] [, tbl_name[*]] ...
FROM table_references
[WHERE where_condition]
```

Or:

```
DELETE [LOW_PRIORITY] [QUICK] [IGNORE]
      FROM tbl_name[*] [, tbl_name[*]] ...
      USING table_references
      [WHERE where_condition]
```

For the single-table syntax, the `DELETE` statement deletes rows from `tbl_name` and returns the number of rows deleted. The `WHERE` clause, if given, specifies the conditions that identify which rows to delete. With no `WHERE` clause, all rows are deleted. If the `ORDER BY` clause is specified, the rows are deleted in the order that is specified. The `LIMIT` clause places a limit on the number of rows that can be deleted.

For the multiple-table syntax, `DELETE` deletes from each `tbl_name` the rows that satisfy the conditions. In this case, `ORDER BY` and `LIMIT` cannot be used.

`where_condition` is an expression that evaluates to true for each row to be deleted. It is specified as described in [Section 13.2.7, “SELECT Syntax”](#).

As stated, a `DELETE` statement with no `WHERE` clause deletes all rows. A faster way to do this, when you do not want to know the number of deleted rows, is to use `TRUNCATE TABLE`. See [Section 13.2.9, “TRUNCATE Syntax”](#).

In MySQL 3.23, `DELETE` without a `WHERE` clause returns zero as the number of affected rows.

In MySQL 3.23, if you really want to know how many rows are deleted when you are deleting all rows, and are willing to suffer a speed penalty, you can use a `DELETE` statement that includes a `WHERE` clause with an expression that is true for every row. For example:

```
mysql> DELETE FROM tbl_name WHERE 1>0;
```

This is much slower than `TRUNCATE tbl_name`, because it deletes rows one at a time.

If you delete the row containing the maximum value for an `AUTO_INCREMENT` column, the value is reused later for an `ISAM` or `BDB` table, but not for a `MyISAM` or `InnoDB` table. If you delete all rows in the table with `DELETE FROM tbl_name` (without a `WHERE` clause) in `AUTOCOMMIT` mode, the sequence starts over for all storage engines except `InnoDB` and (as of MySQL 4.0) `MyISAM`. There are some exceptions to this behavior for `InnoDB` tables, as discussed in [Section 14.2.7.3, “How AUTO\\_INCREMENT Columns Work in InnoDB”](#).

For `MyISAM` and `BDB` tables, you can specify an `AUTO_INCREMENT` secondary column in a multiple-column key. In this case, reuse of values deleted from the top of the sequence occurs even for `MyISAM` tables. See [Section 3.6.9, “Using AUTO\\_INCREMENT”](#).

The `DELETE` statement supports the following modifiers:

- If you specify `LOW_PRIORITY`, the server delays execution of the `DELETE` until no other clients are reading from the table.
- For `MyISAM` tables, if you use the `QUICK` keyword, the storage engine does not merge index leaves during delete, which may speed up some kinds of delete operations.
- The `IGNORE` keyword causes MySQL to ignore all errors during the process of deleting rows. (Errors encountered during the parsing stage are processed in the usual manner.) Errors that are ignored due to the use of `IGNORE` are returned as warnings. This option first appeared in MySQL 4.1.1.

The speed of delete operations may also be affected by factors discussed in [Section 7.2.15, “Speed of DELETE Statements”](#).

In `MyISAM` tables, deleted rows are maintained in a linked list and subsequent `INSERT` operations reuse old row positions. To reclaim unused space and reduce file sizes, use the `OPTIMIZE TABLE` statement or the `myisamchk` utility to reorganize tables. `OPTIMIZE TABLE` is easier, but `myisamchk` is faster. See [Section 13.5.2.5, “OPTIMIZE TABLE Syntax”](#), and [Section 8.3, “myisamchk — MyISAM Table-Maintenance Utility”](#).

The `QUICK` modifier affects whether index leaves are merged for delete operations. `DELETE QUICK` is most useful for applications where index values for deleted rows are replaced by similar index values from rows inserted later. In this case, the holes left by deleted values are reused.

`DELETE QUICK` is not useful when deleted values lead to undef-filled index blocks spanning a range of index values for which new inserts occur again. In this case, use of `QUICK` can lead to wasted space in the index that remains unreclaimed. Here is an example of such a scenario:

1. Create a table that contains an indexed `AUTO_INCREMENT` column.



2. Insert many rows into the table. Each insert results in an index value that is added to the high end of the index.
3. Delete a block of rows at the low end of the column range using `DELETE QUICK`.

In this scenario, the index blocks associated with the deleted index values become undef-filled but are not merged with other index blocks due to the use of `QUICK`. They remain undef-filled when new inserts occur, because new rows do not have index values in the deleted range. Furthermore, they remain undef-filled even if you later use `DELETE` without `QUICK`, unless some of the deleted index values happen to lie in index blocks within or adjacent to the undef-filled blocks. To reclaim unused index space under these circumstances, use `OPTIMIZE TABLE`.

If you are going to delete many rows from a table, it might be faster to use `DELETE QUICK` followed by `OPTIMIZE TABLE`. This rebuilds the index rather than performing many index block merge operations.

The MySQL-specific `LIMIT row_count` option to `DELETE` tells the server the maximum number of rows to be deleted before control is returned to the client. This can be used to ensure that a given `DELETE` statement does not take too much time. You can simply repeat the `DELETE` statement until the number of affected rows is less than the `LIMIT` value.

If the `DELETE` statement includes an `ORDER BY` clause, the rows are deleted in the order specified by the clause. This is really useful only in conjunction with `LIMIT`. For example, the following statement finds rows matching the `WHERE` clause, sorts them by `timestamp_column`, and deletes the first (oldest) one:

```
DELETE FROM somelog WHERE user = 'jcole'  
ORDER BY timestamp_column LIMIT 1;
```

`ORDER BY` can be used with `DELETE` beginning with MySQL 4.0.0.

From MySQL 4.0, you can specify multiple tables in the `DELETE` statement to delete rows from one or more tables depending on a particular condition in multiple tables. However, you cannot use `ORDER BY` or `LIMIT` in a multiple-table `DELETE`.

You can specify multiple tables in a `DELETE` statement to delete rows from one or more tables depending on the particular condition in the `WHERE` clause. However, you cannot use `ORDER BY` or `LIMIT` in a multiple-table `DELETE`. The `table_references` clause lists the tables involved in the join. Its syntax is described in [Section 13.2.7.1, “JOIN Syntax”](#).

The first multiple-table `DELETE` syntax is supported starting from MySQL 4.0.0. The second is supported starting from MySQL 4.0.2.

For the first multiple-table syntax, only matching rows from the tables listed before the `FROM` clause are deleted. For the second multiple-table syntax, only matching rows from the tables listed in the `FROM` clause (before the `USING` clause) are deleted. The effect is that you can delete rows from many tables at the same time and have additional tables that are used only for searching:

```
DELETE t1, t2 FROM t1, t2, t3 WHERE t1.id=t2.id AND t2.id=t3.id;
```

Or:

```
DELETE FROM t1, t2 USING t1, t2, t3 WHERE t1.id=t2.id AND t2.id=t3.id;
```

These statements use all three tables when searching for rows to delete, but delete matching rows only from tables `t1` and `t2`.

The preceding examples show inner joins that use the comma operator, but multiple-table **DELETE** statements can use any type of join allowed in **SELECT** statements, such as **LEFT JOIN**.

The syntax allows `.*` after the table names for compatibility with **Access**.

If you use a multiple-table **DELETE** statement involving **InnoDB** tables for which there are foreign key constraints, the MySQL optimizer might process tables in an order that differs from that of their parent/child relationship. In this case, the statement fails and rolls back. Instead, you should delete from a single table and rely on the **ON DELETE** capabilities that **InnoDB** provides to cause the other tables to be modified accordingly.

**Note:** The syntax for multiple-table **DELETE** statements that use table aliases changed between MySQL 4.0 and 4.1. In MySQL 4.0, you should use the true table name to refer to any table from which rows should be deleted:

```
DELETE test FROM test AS t1, test2 WHERE ...
```

In MySQL 4.1, you must use the alias:

```
DELETE t1 FROM test AS t1, test2 WHERE ...
```

We did not make this change in 4.0 to avoid breaking any old 4.0 applications that were using the old syntax. However, if you use such **DELETE** statements and are using replication, the change in syntax means that a 4.0 master cannot replicate to 4.1 (or higher) slaves.

Cross-database deletes are supported for multiple-table deletes, but in this case, you must refer to the tables without using aliases. For example:

```
DELETE test1.tmp1, test2.tmp2 FROM test1.tmp1, test2.tmp2 WHERE ...
```

Currently, you cannot delete from a table and select from the same table in a subquery.

## 13.2.2. **DO** Syntax

```
DO expr [, expr] ...
```

**DO** executes the expressions but does not return any results. In most respects, **DO** is shorthand for **SELECT *expr*, ...**, but has the advantage that it is slightly faster when you do not care about the result.

**DO** is useful primarily with functions that have side effects, such as **RELEASE\_LOCK()**.

**DO** was added in MySQL 3.23.47.

## 13.2.3. **HANDLER** Syntax

```
HANDLER tbl_name OPEN [ AS alias ]
HANDLER tbl_name READ index_name { = | >= | <= | < } (value1,value2,...)
  [ WHERE where_condition ] [LIMIT ... ]
HANDLER tbl_name READ index_name { FIRST | NEXT | PREV | LAST }
  [ WHERE where_condition ] [LIMIT ... ]
HANDLER tbl_name READ { FIRST | NEXT }
  [ WHERE where_condition ] [LIMIT ... ]
HANDLER tbl_name CLOSE
```

The **HANDLER** statement provides direct access to table storage engine interfaces. It is available for **MyISAM** tables as MySQL 4.0.0 and **InnoDB** tables as of MySQL 4.0.3.

The `HANDLER ... OPEN` statement opens a table, making it accessible via subsequent `HANDLER ... READ` statements. This table object is not shared by other threads and is not closed until the thread calls `HANDLER ... CLOSE` or the thread terminates. If you open the table using an alias, further references to the open table with other `HANDLER` statements must use the alias rather than the table name.

The first `HANDLER ... READ` syntax fetches a row where the index specified satisfies the given values and the `WHERE` condition is met. If you have a multiple-column index, specify the index column values as a comma-separated list. Either specify values for all the columns in the index, or specify values for a leftmost prefix of the index columns. Suppose that an index `my_idx` includes three columns named `col_a`, `col_b`, and `col_c`, in that order. The `HANDLER` statement can specify values for all three columns in the index, or for the columns in a leftmost prefix. For example:

```
HANDLER ... READ my_idx = (col_a_val,col_b_val,col_c_val) ...
HANDLER ... READ my_idx = (col_a_val,col_b_val) ...
HANDLER ... READ my_idx = (col_a_val) ...
```

To employ the `HANDLER` interface to refer to a table's `PRIMARY KEY`, use the quoted identifier ``PRIMARY``:

```
HANDLER tbl_name READ `PRIMARY` ...
```

The second `HANDLER ... READ` syntax fetches a row from the table in index order that matches the `WHERE` condition.

The third `HANDLER ... READ` syntax fetches a row from the table in natural row order that matches the `WHERE` condition. It is faster than `HANDLER tbl_name READ index_name` when a full table scan is desired. Natural row order is the order in which rows are stored in a `MyISAM` table data file. This statement works for `InnoDB` tables as well, but there is no such concept because there is no separate data file.

Without a `LIMIT` clause, all forms of `HANDLER ... READ` fetch a single row if one is available. To return a specific number of rows, include a `LIMIT` clause. It has the same syntax as for the `SELECT` statement. See [Section 13.2.7, “SELECT Syntax”](#).

`HANDLER ... CLOSE` closes a table that was opened with `HANDLER ... OPEN`.

`HANDLER` is a somewhat low-level statement. For example, it does not provide consistency. That is, `HANDLER ... OPEN` does *not* take a snapshot of the table, and does *not* lock the table. This means that after a `HANDLER ... OPEN` statement is issued, table data can be modified (by the current thread or other threads) and these modifications might be only partially visible to `HANDLER ... NEXT` or `HANDLER ... PREV` scans.

There are several reasons to use the `HANDLER` interface instead of normal `SELECT` statements:

- `HANDLER` is faster than `SELECT`:
  - A designated storage engine handler object is allocated for the `HANDLER ... OPEN`. The object is reused for subsequent `HANDLER` statements for that table; it need not be reinitialized for each one.
  - There is less parsing involved.
  - There is no optimizer or query-checking overhead.
  - The table does not have to be locked between two handler requests.
  - The handler interface does not have to provide a consistent look of the data (for example, dirty reads are allowed), so the storage engine can use optimizations that `SELECT` does not normally

allow.

- For applications that use a low-level [ISAM](#)-like interface, [HANDLER](#) makes it much easier to port them to MySQL.
- [HANDLER](#) enables you to traverse a database in a manner that is difficult (or even impossible) to accomplish with [SELECT](#). The [HANDLER](#) interface is a more natural way to look at data when working with applications that provide an interactive user interface to the database.

## 13.2.4. INSERT Syntax

```
INSERT [LOW_PRIORITY | DELAYED | HIGH_PRIORITY] [IGNORE]
[INTO] tbl_name [(col_name,...)]
VALUES ({expr | DEFAULT},...),(...),...
[ ON DUPLICATE KEY UPDATE col_name=expr, ... ]
```

Or:

```
INSERT [LOW_PRIORITY | DELAYED | HIGH_PRIORITY] [IGNORE]
[INTO] tbl_name
SET col_name={expr | DEFAULT}, ...
[ ON DUPLICATE KEY UPDATE col_name=expr, ... ]
```

Or:

```
INSERT [LOW_PRIORITY | HIGH_PRIORITY] [IGNORE]
[INTO] tbl_name [(col_name,...)]
SELECT ...
[ ON DUPLICATE KEY UPDATE col_name=expr, ... ]
```

[INSERT](#) inserts new rows into an existing table. The [INSERT ... VALUES](#) and [INSERT ... SET](#) forms of the statement insert rows based on explicitly specified values. The [INSERT ... SELECT](#) form inserts rows selected from another table or tables. The [INSERT ... VALUES](#) form with multiple value lists is supported in MySQL 3.22.5 or later. The [INSERT ... SET](#) syntax is supported in MySQL 3.22.10 or later. [INSERT ... SELECT](#) is discussed further in [Section 13.2.4.1, “INSERT ... SELECT Syntax”](#).

You can use [REPLACE](#) instead of [INSERT](#) to overwrite old rows. [REPLACE](#) is the counterpart to [INSERT IGNORE](#) in the treatment of new rows that contain unique key values that duplicate old rows: The new rows are used to replace the old rows rather than being discarded. See [Section 13.2.6, “REPLACE Syntax”](#).

*tbl\_name* is the table into which rows should be inserted. The columns for which the statement provides values can be specified as follows:

- You can provide a comma-separated list of column names following the table name. In this case, a value for each named column must be provided by the [VALUES](#) list or the [SELECT](#) statement.
- If you do not specify a list of column names for [INSERT ... VALUES](#) or [INSERT ... SELECT](#), values for every column in the table must be provided by the [VALUES](#) list or the [SELECT](#) statement. If you do not know the order of the columns in the table, use [DESCRIBE \*tbl\\_name\*](#) to find out.
- The [SET](#) clause indicates the column names explicitly.

Column values can be given in several ways:

- Normally, any column not explicitly given a value is set to its default (explicit or implicit) value. For example, if you specify a column list that does not name all the columns in the table, unnamed columns are set to their default values. Default value assignment is described in [Section 11.1.4](#), “Data Type Default Values”, and [Section 1.9.6.2](#), “Constraints on Invalid Data”.
- You can use the keyword `DEFAULT` to explicitly set a column to its default value. (New in MySQL 4.0.3.) This makes it easier to write `INSERT` statements that assign values to all but a few columns, because it enables you to avoid writing an incomplete `VALUES` list that does not include a value for each column in the table. Otherwise, you would have to write out the list of column names corresponding to each value in the `VALUES` list.

As of MySQL 4.1.0, you can use `DEFAULT(col_name)` as a more general form that can be used in expressions to produce a given column's default value.

- If both the column list and the `VALUES` list are empty, `INSERT` creates a row with each column set to its default value:

```
INSERT INTO tbl_name () VALUES();
```

- You can specify an expression `expr` to provide a column value. This might involve type conversion if the type of the expression does not match the type of the column, and conversion of a given value can result in different inserted values depending on the data type. For example, inserting the string `'1999.0e-2'` into an `INT`, `FLOAT`, `DECIMAL(10,6)`, or `YEAR` column results in the values `1999`, `19.9921`, `19.992100`, and `1999` being inserted, respectively. The reason the value stored in the `INT` and `YEAR` columns is `1999` is that the string-to-integer conversion looks only at as much of the initial part of the string as may be considered a valid integer or year. For the floating-point and fixed-point columns, the string-to-floating-point conversion considers the entire string a valid floating-point value.

An expression `expr` can refer to any column that was set earlier in a value list. For example, you can do this because the value for `col2` refers to `col1`, which has previously been assigned:

```
INSERT INTO tbl_name (col1,col2) VALUES(15,col1*2);
```

But the following is not legal, because the value for `col1` refers to `col2`, which is assigned after `col1`:

```
INSERT INTO tbl_name (col1,col2) VALUES(col2*2,15);
```

One exception involves columns that contain `AUTO_INCREMENT` values. Because the `AUTO_INCREMENT` value is generated after other value assignments, any reference to an `AUTO_INCREMENT` column in the assignment returns a 0.

`INSERT` statements that use `VALUES` syntax can insert multiple rows. To do this, include multiple lists of column values, each enclosed within parentheses and separated by commas. Example:

```
INSERT INTO tbl_name (a,b,c) VALUES(1,2,3),(4,5,6),(7,8,9);
```

The values list for each row must be enclosed within parentheses. The following statement is illegal because the number of values in the list does not match the number of column names:

```
INSERT INTO tbl_name (a,b,c) VALUES(1,2,3,4,5,6,7,8,9);
```

The rows-affected value for an `INSERT` can be obtained using the `mysql_affected_rows()` C API function. See [Section 17.2.3.1](#), “`mysql_affected_rows()`”.

If you use an `INSERT ... VALUES` statement with multiple value lists or `INSERT ... SELECT`, the statement returns an information string in this format:

```
Records: 100 Duplicates: 0 Warnings: 0
```

`Records` indicates the number of rows processed by the statement. (This is not necessarily the number of rows actually inserted because `Duplicates` can be non-zero.) `Duplicates` indicates the number of rows that could not be inserted because they would duplicate some existing unique index value. `Warnings` indicates the number of attempts to insert column values that were problematic in some way. Warnings can occur under any of the following conditions:

- Inserting `NULL` into a column that has been declared `NOT NULL`. For multiple-row `INSERT` statements or `INSERT INTO ... SELECT` statements, the column is set to the implicit default value for the column data type. This is `0` for numeric types, the empty string ( `' '` ) for string types, and the “zero” value for date and time types. `INSERT INTO ... SELECT` statements are handled the same way as multiple-row inserts because the server does not examine the result set from the `SELECT` to see whether it returns a single row. (For a single-row `INSERT`, no warning occurs when `NULL` is inserted into a `NOT NULL` column. Instead, the statement fails with an error.)
- Setting a numeric column to a value that lies outside the column's range. The value is clipped to the closest endpoint of the range.
- Assigning a value such as `'10.34 a'` to a numeric column. The trailing non-numeric text is stripped off and the remaining numeric part is inserted. If the string value has no leading numeric part, the column is set to `0`.
- Inserting a string into a string column (`CHAR`, `VARCHAR`, `TEXT`, or `BLOB`) that exceeds the column's maximum length. The value is truncated to the column's maximum length.
- Inserting a value into a date or time column that is illegal for the data type. The column is set to the appropriate zero value for the type.

If you are using the C API, the information string can be obtained by invoking the `mysql_info()` function. See [Section 17.2.3.33](#), “`mysql_info()`”.

If `INSERT` inserts a row into a table that has an `AUTO_INCREMENT` column, you can find the value used for that column by using the SQL `LAST_INSERT_ID()` function. From within the C API, use the `mysql_insert_id()` function. However, you should note that the two functions do not always behave identically. The behavior of `INSERT` statements with respect to `AUTO_INCREMENT` columns is discussed further in [Section 12.9.3](#), “Information Functions”, and [Section 17.2.3.35](#), “`mysql_insert_id()`”.

The `INSERT` statement supports the following modifiers:

- If you use the `DELAYED` keyword, the server puts the row or rows to be inserted into a buffer, and the client issuing the `INSERT DELAYED` statement can then continue immediately. If the table is in use, the server holds the rows. When the table is free, the server begins inserting rows, checking periodically to see whether there are any new read requests for the table. If there are, the delayed row queue is suspended until the table becomes free again. See [Section 13.2.4.2](#), “`INSERT DELAYED Syntax`”. `DELAYED` was added in MySQL 3.22.5.

`DELAYED` is ignored with `INSERT ... SELECT` or `INSERT ... ON DUPLICATE KEY UPDATE`.

- If you use the `LOW_PRIORITY` keyword, execution of the `INSERT` is delayed until no other clients

are reading from the table. This includes other clients that began reading while existing clients are reading, and while the `INSERT LOW_PRIORITY` statement is waiting. It is possible, therefore, for a client that issues an `INSERT LOW_PRIORITY` statement to wait for a very long time (or even forever) in a read-heavy environment. (This is in contrast to `INSERT DELAYED`, which lets the client continue at once. Note that `LOW_PRIORITY` should normally not be used with `MyISAM` tables because doing so disables concurrent inserts. See [Section 7.3.3, “Concurrent Inserts”](#).

`LOW_PRIORITY` was added in MySQL 3.22.5.

- If you specify `HIGH_PRIORITY`, it overrides the effect of the `--low-priority-updates` option if the server was started with that option. It also causes concurrent inserts not to be used. `HIGH_PRIORITY` was added in MySQL 3.23.11.
- If you use the `IGNORE` keyword, errors that occur while executing the `INSERT` statement are treated as warnings instead. For example, without `IGNORE`, a row that duplicates an existing `UNIQUE` index or `PRIMARY KEY` value in the table causes a duplicate-key error and the statement is aborted. With `IGNORE`, the row still is not inserted, but no error is issued. Data conversions that would trigger errors abort the statement if `IGNORE` is not specified. With `IGNORE`, invalid values are adjusted to the closest values and inserted; warnings are produced but the statement does not abort. You can determine with the `mysql_info()` C API function how many rows were actually inserted into the table.
- If you specify `ON DUPLICATE KEY UPDATE`, and a row is inserted that would cause a duplicate value in a `UNIQUE` index or `PRIMARY KEY`, an `UPDATE` of the old row is performed. See [Section 13.2.4.3, “INSERT ... ON DUPLICATE KEY UPDATE Syntax”](#). `ON DUPLICATE KEY UPDATE` was added in MySQL 4.1.0.

### 13.2.4.1. INSERT ... SELECT Syntax

```
INSERT [LOW_PRIORITY | HIGH_PRIORITY] [IGNORE]
  [INTO] tbl_name [(col_name,...)]
  SELECT ...
  [ ON DUPLICATE KEY UPDATE col_name=expr, ... ]
```

With `INSERT ... SELECT`, you can quickly insert many rows into a table from one or many tables. For example:

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

The following conditions hold for a `INSERT ... SELECT` statements:

- Prior to MySQL 4.0.1, `INSERT ... SELECT` implicitly operates in `IGNORE` mode. As of MySQL 4.0.1, specify `IGNORE` explicitly to ignore rows that would cause duplicate-key violations.
- `DELAYED` is ignored with `INSERT ... SELECT`.
- Prior to MySQL 4.0.14, the target table of the `INSERT` statement cannot appear in the `FROM` clause of the `SELECT` part of the query. This limitation is lifted in 4.0.14. In this case, MySQL creates a temporary table to hold the rows from the `SELECT` and then inserts those rows into the target table.
- `AUTO_INCREMENT` columns work as usual.
- To ensure that the binary log can be used to re-create the original tables, MySQL does not allow concurrent inserts for `INSERT ... SELECT` statements.
- Currently, you cannot insert into a table and select from the same table in a subquery.

In the values part of `ON DUPLICATE KEY UPDATE`, you can refer to columns in other tables, as long as you do not use `GROUP BY` in the `SELECT` part. One side effect is that you must qualify non-unique column names in the values part.

### 13.2.4.2. `INSERT DELAYED` Syntax

```
INSERT DELAYED ...
```

The `DELAYED` option for the `INSERT` statement is a MySQL extension to standard SQL that is very useful if you have clients that cannot or need not wait for the `INSERT` to complete. This is a common situation when you use MySQL for logging and you also periodically run `SELECT` and `UPDATE` statements that take a long time to complete. `DELAYED` was introduced in MySQL 3.22.15.

When a client uses `INSERT DELAYED`, it gets an okay from the server at once, and the row is queued to be inserted when the table is not in use by any other thread.

Another major benefit of using `INSERT DELAYED` is that inserts from many clients are bundled together and written in one block. This is much faster than performing many separate inserts.

Note that `INSERT DELAYED` is slower than a normal `INSERT` if the table is not otherwise in use. There is also the additional overhead for the server to handle a separate thread for each table for which there are delayed rows. This means that you should use `INSERT DELAYED` only when you are really sure that you need it.

The queued rows are held only in memory until they are inserted into the table. This means that if you terminate `mysqld` forcibly (for example, with `kill -9`) or if `mysqld` dies unexpectedly, *any queued rows that have not been written to disk are lost*.

There are some constraints on the use of `DELAYED`:

- `INSERT DELAYED` works only with `ISAM`, `MyISAM`, and (beginning with MySQL 4.1) `MEMORY` tables. See [Section 14.1, “The MyISAM Storage Engine”](#), [Section 14.4, “The MEMORY \(HEAP\) Storage Engine”](#), and [Section 14.7, “The ARCHIVE Storage Engine”](#).

For `MyISAM` tables, if there are no free blocks in the middle of the data file, concurrent `SELECT` and `INSERT` statements are supported. Under these circumstances, you very seldom need to use `INSERT DELAYED` with `MyISAM`.

- `INSERT DELAYED` should be used only for `INSERT` statements that specify value lists. This is enforced as of MySQL 4.0.18. The server ignores `DELAYED` for `INSERT ... SELECT` or `INSERT ... ON DUPLICATE KEY UPDATE` statements.
- Because the `INSERT DELAYED` statement returns immediately, before the rows are inserted, you cannot use `LAST_INSERT_ID( )` to get the `AUTO_INCREMENT` value that the statement might generate.
- `DELAYED` rows are not visible to `SELECT` statements until they actually have been inserted.
- `DELAYED` is ignored on slave replication servers because it could cause the slave to have different data than the master.
- Pending `INSERT DELAYED` statements are lost if a table is write locked and `ALTER TABLE` is used to modify the table structure.

The following describes in detail what happens when you use the `DELAYED` option to `INSERT` or `REPLACE`. In this description, the “thread” is the thread that received an `INSERT DELAYED` statement



and “handler” is the thread that handles all `INSERT DELAYED` statements for a particular table.

- When a thread executes a `DELAYED` statement for a table, a handler thread is created to process all `DELAYED` statements for the table, if no such handler already exists.
- The thread checks whether the handler has previously acquired a `DELAYED` lock; if not, it tells the handler thread to do so. The `DELAYED` lock can be obtained even if other threads have a `READ` or `WRITE` lock on the table. However, the handler waits for all `ALTER TABLE` locks or `FLUSH TABLES` statements to finish, to ensure that the table structure is up to date.
- The thread executes the `INSERT` statement, but instead of writing the row to the table, it puts a copy of the final row into a queue that is managed by the handler thread. Any syntax errors are noticed by the thread and reported to the client program.
- The client cannot obtain from the server the number of duplicate rows or the `AUTO_INCREMENT` value for the resulting row, because the `INSERT` returns before the insert operation has been completed. (If you use the C API, the `mysql_info()` function does not return anything meaningful, for the same reason.)
- The binary log is updated by the handler thread when the row is inserted into the table. In case of multiple-row inserts, the binary log is updated when the first row is inserted.
- Each time that `delayed_insert_limit` rows are written, the handler checks whether any `SELECT` statements are still pending. If so, it allows these to execute before continuing.
- When the handler has no more rows in its queue, the table is unlocked. If no new `INSERT DELAYED` statements are received within `delayed_insert_timeout` seconds, the handler terminates.
- If more than `delayed_queue_size` rows are pending in a specific handler queue, the thread requesting `INSERT DELAYED` waits until there is room in the queue. This is done to ensure that `mysqld` does not use all memory for the delayed memory queue.
- The handler thread shows up in the MySQL process list with `delayed_insert` in the `Command` column. It is killed if you execute a `FLUSH TABLES` statement or kill it with `KILL thread_id`. However, before exiting, it first stores all queued rows into the table. During this time it does not accept any new `INSERT` statements from other threads. If you execute an `INSERT DELAYED` statement after this, a new handler thread is created.

Note that this means that `INSERT DELAYED` statements have higher priority than normal `INSERT` statements if there is an `INSERT DELAYED` handler running. Other update statements have to wait until the `INSERT DELAYED` queue is empty, someone terminates the handler thread (with `KILL thread_id`), or someone executes a `FLUSH TABLES`.

- The following status variables provide information about `INSERT DELAYED` statements:

Status Variable	Meaning
<code>Delayed_insert_threads</code>	Number of handler threads
<code>Delayed_writes</code>	Number of rows written with <code>INSERT DELAYED</code>
<code>Not_flushed_delayed_rows</code>	Number of rows waiting to be written

You can view these variables by issuing a `SHOW STATUS` statement or by executing a `mysqladmin extended-status` command.

### 13.2.4.3. `INSERT ... ON DUPLICATE KEY UPDATE` Syntax

If you specify `ON DUPLICATE KEY UPDATE` (added in MySQL 4.1.0), and a row is inserted that would cause a duplicate value in a `UNIQUE` index or `PRIMARY KEY`, an `UPDATE` of the old row is performed. For example, if column `a` is declared as `UNIQUE` and contains the value `1`, the following two statements have identical effect:

```
INSERT INTO table (a,b,c) VALUES (1,2,3)
ON DUPLICATE KEY UPDATE c=c+1;

UPDATE table SET c=c+1 WHERE a=1;
```

The rows-affected value is 1 if the row is inserted as a new record and 2 if an existing record is updated.

If column `b` is also unique, the `INSERT` is equivalent to this `UPDATE` statement instead:

```
UPDATE table SET c=c+1 WHERE a=1 OR b=2 LIMIT 1;
```

If `a=1 OR b=2` matches several rows, only *one* row is updated. In general, you should try to avoid using an `ON DUPLICATE KEY` clause on tables with multiple unique indexes.

As of MySQL 4.1.1, you can use the `VALUES(col_name)` function in the `UPDATE` clause to refer to column values from the `INSERT` portion of the `INSERT ... UPDATE` statement. In other words, `VALUES(col_name)` in the `UPDATE` clause refers to the value of `col_name` that would be inserted, had no duplicate-key conflict occurred. This function is especially useful in multiple-row inserts. The `VALUES()` function is meaningful only in `INSERT ... UPDATE` statements and returns `NULL` otherwise. Example:

```
INSERT INTO table (a,b,c) VALUES (1,2,3),(4,5,6)
ON DUPLICATE KEY UPDATE c=VALUES(a)+VALUES(b);
```

That statement is identical to the following two statements:

```
INSERT INTO table (a,b,c) VALUES (1,2,3)
ON DUPLICATE KEY UPDATE c=3;
INSERT INTO table (a,b,c) VALUES (4,5,6)
ON DUPLICATE KEY UPDATE c=9;
```

If a table contains an `AUTO_INCREMENT` column and `INSERT ... UPDATE` inserts a row, the `LAST_INSERT_ID()` function returns the `AUTO_INCREMENT` value. If the statement updates a row instead, `LAST_INSERT_ID()` is not meaningful. However, you can work around this by using `LAST_INSERT_ID(expr)`. Suppose that `id` is the `AUTO_INCREMENT` column. To make `LAST_INSERT_ID()` meaningful for updates, insert rows as follows:

```
INSERT INTO table (a,b,c) VALUES (1,2,3)
ON DUPLICATE KEY UPDATE id=LAST_INSERT_ID(id), c=3;
```

The `DELAYED` option is ignored when you use `ON DUPLICATE KEY UPDATE`.

## 13.2.5. LOAD DATA INFILE Syntax

```
LOAD DATA [LOW_PRIORITY | CONCURRENT] [LOCAL] INFILE 'file_name'
[REPLACE | IGNORE]
INTO TABLE tbl_name
[FIELDS
  [TERMINATED BY 'string']
  [[OPTIONALLY] ENCLOSED BY 'char']
  [ESCAPED BY 'char']
]
[LINES
  [STARTING BY 'string']
  [TERMINATED BY 'string']
]
[IGNORE number LINES]
```

```
[ (col_name, ... ) ]
```

The `LOAD DATA INFILE` statement reads rows from a text file into a table at a very high speed. The filename must be given as a literal string.

`LOAD DATA INFILE` is the complement of `SELECT ... INTO OUTFILE`. (See [Section 13.2.7, “SELECT Syntax”](#).) To write data from a table to a file, use `SELECT ... INTO OUTFILE`. To read the file back into a table, use `LOAD DATA INFILE`. The syntax of the `FIELDS` and `LINES` clauses is the same for both statements. Both clauses are optional, but `FIELDS` must precede `LINES` if both are specified.

For more information about the efficiency of `INSERT` versus `LOAD DATA INFILE` and speeding up `LOAD DATA INFILE`, see [Section 7.2.13, “Speed of INSERT Statements”](#).

As of MySQL 4.1, the character set indicated by the `character_set_database` system variable is used to interpret the information in the file. `SET NAMES` and the setting of `character_set_client` do not affect interpretation of input.

Note that it is currently not possible to load data files that use the `ucs2` character set.

You can also load data files by using the `mysqlimport` utility; it operates by sending a `LOAD DATA INFILE` statement to the server. The `--local` option causes `mysqlimport` to read data files from the client host. You can specify the `--compress` option to get better performance over slow networks if the client and server support the compressed protocol. See [Section 8.14, “mysqlimport — A Data Import Program”](#).

If you use `LOW_PRIORITY`, execution of the `LOAD DATA` statement is delayed until no other clients are reading from the table.

If you specify `CONCURRENT` with a `MyISAM` table that satisfies the condition for concurrent inserts (that is, it contains no free blocks in the middle), other threads can retrieve data from the table while `LOAD DATA` is executing. Using this option affects the performance of `LOAD DATA` a bit, even if no other thread is using the table at the same time.

The `LOCAL` keyword, if specified, is interpreted with respect to the client end of the connection:

- If `LOCAL` is specified, the file is read by the client program on the client host and sent to the server. The file can be given as a full pathname to specify its exact location. If given as a relative pathname, the name is interpreted relative to the directory in which the client program was started.

`LOCAL` is available in MySQL 3.22.6 or later.

- If `LOCAL` is not specified, the file must be located on the server host and is read directly by the server. The server uses the following rules to locate the file:
  - If the filename is an absolute pathname, the server uses it as given.
  - If the filename is a relative pathname with one or more leading components, the server searches for the file relative to the server's data directory.
  - If a filename with no leading components is given, the server looks for the file in the database directory of the default database.

Note that, in the non-`LOCAL` case, these rules mean that a file named as `./myfile.txt` is read from the server's data directory, whereas the file named as `myfile.txt` is read from the database directory of the default database. For example, if `db1` is the default database, the following `LOAD DATA` statement reads the file `data.txt` from the database directory for `db1`, even though the statement expli-

city loads the file into a table in the `db2` database:

```
LOAD DATA INFILE 'data.txt' INTO TABLE db2.my_table;
```

Windows pathnames are specified using forward slashes rather than backslashes. If you do use backslashes, you must double them.

For security reasons, when reading text files located on the server, the files must either reside in the database directory or be readable by all. Also, to use `LOAD DATA INFILE` on server files, you must have the `FILE` privilege. See [Section 5.7.3, “Privileges Provided by MySQL”](#).

Using `LOCAL` is a bit slower than letting the server access the files directly, because the contents of the file must be sent over the connection by the client to the server. On the other hand, you do not need the `FILE` privilege to load local files.

As of MySQL 3.23.49 and MySQL 4.0.2 (4.0.13 on Windows), `LOCAL` works only if your server and your client both have been enabled to allow it. For example, if `mysqld` was started with `-local-infile=0`, `LOCAL` does not work. See [Section 5.6.4, “Security Issues with LOAD DATA LOCAL”](#).

On Unix, if you need `LOAD DATA` to read from a pipe, you can use the following technique (here we load the listing of the `/` directory into a table):

```
mkfifo /mysql/db/x/x
chmod 666 /mysql/db/x/x
find / -ls > /mysql/db/x/x
mysql -e "LOAD DATA INFILE 'x' INTO TABLE x" x
```

If you are using a version of MySQL older than 3.23.25, you can use this technique only with `LOAD DATA LOCAL INFILE`.

If you are using MySQL before version 3.23.24, you cannot read from a FIFO with `LOAD DATA INFILE`. If you need to read from a FIFO (for example, the output from `gunzip`), use `LOAD DATA LOCAL INFILE` instead.

The `REPLACE` and `IGNORE` keywords control handling of input rows that duplicate existing rows on unique key values:

- If you specify `REPLACE`, input rows replace existing rows. In other words, rows that have the same value for a primary key or unique index as an existing row. See [Section 13.2.6, “REPLACE Syntax”](#).
- If you specify `IGNORE`, input rows that duplicate an existing row on a unique key value are skipped. If you do not specify either option, the behavior depends on whether the `LOCAL` keyword is specified. Without `LOCAL`, an error occurs when a duplicate key value is found, and the rest of the text file is ignored. With `LOCAL`, the default behavior is the same as if `IGNORE` is specified; this is because the server has no way to stop transmission of the file in the middle of the operation.

If you want to ignore foreign key constraints during the load operation, you can issue a `SET FOREIGN_KEY_CHECKS=0` statement before executing `LOAD DATA`.

If you use `LOAD DATA INFILE` on an empty `MyISAM` table, all non-unique indexes are created in a separate batch (as for `REPAIR TABLE`). Normally, this makes `LOAD DATA INFILE` much faster when you have many indexes. In some extreme cases, you can create the indexes even faster by turning them off with `ALTER TABLE ... DISABLE KEYS` before loading the file into the table and using `ALTER TABLE ... ENABLE KEYS` to re-create the indexes after loading the file. See [Section 7.2.13, “Speed of INSERT Statements”](#).

For both the `LOAD DATA INFILE` and `SELECT ... INTO OUTFILE` statements, the syntax of

the `FIELDS` and `LINES` clauses is the same. Both clauses are optional, but `FIELDS` must precede `LINES` if both are specified.

If you specify a `FIELDS` clause, each of its subclauses (`TERMINATED BY`, `[OPTIONALLY] ENCLOSED BY`, and `ESCAPED BY`) is also optional, except that you must specify at least one of them.

If you specify no `FIELDS` clause, the defaults are the same as if you had written this:

```
FIELDS TERMINATED BY '\t' ENCLOSED BY '' ESCAPED BY '\\'
```

If you specify no `LINES` clause, the defaults are the same as if you had written this:

```
LINES TERMINATED BY '\n' STARTING BY ''
```

In other words, the defaults cause `LOAD DATA INFILE` to act as follows when reading input:

- Look for line boundaries at newlines.
- Do not skip over any line prefix.
- Break lines into fields at tabs.
- Do not expect fields to be enclosed within any quoting characters.
- Interpret occurrences of tab, newline, or `'\'` preceded by `'\'` as literal characters that are part of field values.

Conversely, the defaults cause `SELECT ... INTO OUTFILE` to act as follows when writing output:

- Write tabs between fields.
- Do not enclose fields within any quoting characters.
- Use `'\'` to escape instances of tab, newline, or `'\'` that occur within field values.
- Write newlines at the ends of lines.

Backslash is the MySQL escape character within strings, so to write `FIELDS ESCAPED BY '\\'`, you must specify two backslashes for the value to be interpreted as a single backslash.

**Note:** If you have generated the text file on a Windows system, you might have to use `LINES TERMINATED BY '\r\n'` to read the file properly, because Windows programs typically use two characters as a line terminator. Some programs, such as `WordPad`, might use `\r` as a line terminator when writing files. To read such files, use `LINES TERMINATED BY '\r'`.

If all the lines you want to read in have a common prefix that you want to ignore, you can use `LINES STARTING BY 'prefix_string'` to skip over the prefix, *and anything before it*. If a line does not include the prefix, the entire line is skipped. Suppose that you issue the following statement:

```
LOAD DATA INFILE '/tmp/test.txt' INTO TABLE test
FIELDS TERMINATED BY ',' LINES STARTING BY 'xxx';
```

If the data file looks like this:

```
xxx"abc",1
something xxx"def",2
"ghi",3
```

The resulting rows will be ("abc",1) and ("def",2). The third row in the file is skipped because it does not contain the prefix.

The `IGNORE number LINES` option can be used to ignore lines at the start of the file. For example, you can use `IGNORE 1 LINES` to skip over an initial header line containing column names:

```
LOAD DATA INFILE '/tmp/test.txt' INTO TABLE test IGNORE 1 LINES;
```

When you use `SELECT ... INTO OUTFILE` in tandem with `LOAD DATA INFILE` to write data from a database into a file and then read the file back into the database later, the field- and line-handling options for both statements must match. Otherwise, `LOAD DATA INFILE` will not interpret the contents of the file properly. Suppose that you use `SELECT ... INTO OUTFILE` to write a file with fields delimited by commas:

```
SELECT * INTO OUTFILE 'data.txt'
  FIELDS TERMINATED BY ','
FROM table2;
```

To read the comma-delimited file back in, the correct statement would be:

```
LOAD DATA INFILE 'data.txt' INTO TABLE table2
  FIELDS TERMINATED BY ',';
```

If instead you tried to read in the file with the statement shown following, it wouldn't work because it instructs `LOAD DATA INFILE` to look for tabs between fields:

```
LOAD DATA INFILE 'data.txt' INTO TABLE table2
  FIELDS TERMINATED BY '\t';
```

The likely result is that each input line would be interpreted as a single field.

`LOAD DATA INFILE` can be used to read files obtained from external sources. For example, many programs can export data in comma-separated values (CSV) format, such that lines have fields separated by commas and enclosed within double quotes. If lines in such a file are terminated by newlines, the statement shown here illustrates the field- and line-handling options you would use to load the file:

```
LOAD DATA INFILE 'data.txt' INTO TABLE tbl_name
  FIELDS TERMINATED BY ',' ENCLOSED BY '"'
  LINES TERMINATED BY '\n';
```

Any of the field- or line-handling options can specify an empty string (''). If not empty, the `FIELDS [OPTIONALLY] ENCLOSED BY` and `FIELDS ESCAPED BY` values must be a single character. The `FIELDS TERMINATED BY`, `LINES STARTING BY`, and `LINES TERMINATED BY` values can be more than one character. For example, to write lines that are terminated by carriage return/line-feed pairs, or to read a file containing such lines, specify a `LINES TERMINATED BY '\r\n'` clause.

To read a file containing jokes that are separated by lines consisting of `%%`, you can do this

```
CREATE TABLE jokes
  (a INT NOT NULL AUTO_INCREMENT PRIMARY KEY,
  joke TEXT NOT NULL);
LOAD DATA INFILE '/tmp/jokes.txt' INTO TABLE jokes
  FIELDS TERMINATED BY ''
  LINES TERMINATED BY '\n%%\n' (joke);
```

`FIELDS [OPTIONALLY] ENCLOSED BY` controls quoting of fields. For output (`SELECT ... INTO OUTFILE`), if you omit the word `OPTIONALLY`, all fields are enclosed by the `ENCLOSED BY`

character. An example of such output (using a comma as the field delimiter) is shown here:

```
"1","a string","100.20"
"2","a string containing a , comma","102.20"
"3","a string containing a \" quote","102.20"
"4","a string containing a \", quote and comma","102.20"
```

If you specify **OPTIONALLY**, the **ENCLOSED BY** character is used only to enclose values from columns that have a string data type (such as **CHAR**, **BINARY**, **TEXT**, or **ENUM**):

```
1,"a string",100.20
2,"a string containing a , comma",102.20
3,"a string containing a \" quote",102.20
4,"a string containing a \", quote and comma",102.20
```

Note that occurrences of the **ENCLOSED BY** character within a field value are escaped by prefixing them with the **ESCAPED BY** character. Also note that if you specify an empty **ESCAPED BY** value, it is possible to inadvertently generate output that cannot be read properly by **LOAD DATA INFILE**. For example, the preceding output just shown would appear as follows if the escape character is empty. Observe that the second field in the fourth line contains a comma following the quote, which (erroneously) appears to terminate the field:

```
1,"a string",100.20
2,"a string containing a , comma",102.20
3,"a string containing a " quote",102.20
4,"a string containing a ", quote and comma",102.20
```

For input, the **ENCLOSED BY** character, if present, is stripped from the ends of field values. (This is true regardless of whether **OPTIONALLY** is specified; **OPTIONALLY** has no effect on input interpretation.) Occurrences of the **ENCLOSED BY** character preceded by the **ESCAPED BY** character are interpreted as part of the current field value.

If the field begins with the **ENCLOSED BY** character, instances of that character are recognized as terminating a field value only if followed by the field or line **TERMINATED BY** sequence. To avoid ambiguity, occurrences of the **ENCLOSED BY** character within a field value can be doubled and are interpreted as a single instance of the character. For example, if **ENCLOSED BY ' ''** is specified, quotes are handled as shown here:

```
"The "BIG" boss" -> The "BIG" boss
The "BIG" boss  -> The "BIG" boss
The "BIG" boss  -> The "BIG" boss
```

**FIELDS ESCAPED BY** controls how to write or read special characters. If the **FIELDS ESCAPED BY** character is not empty, it is used to prefix the following characters on output:

- The **FIELDS ESCAPED BY** character
- The **FIELDS [OPTIONALLY] ENCLOSED BY** character
- The first character of the **FIELDS TERMINATED BY** and **LINES TERMINATED BY** values
- ASCII 0 (what is actually written following the escape character is ASCII '0', not a zero-valued byte)

If the **FIELDS ESCAPED BY** character is empty, no characters are escaped and **NULL** is output as **NULL**, not **\N**. It is probably not a good idea to specify an empty escape character, particularly if field values in your data contain any of the characters in the list just given.

For input, if the **FIELDS ESCAPED BY** character is not empty, occurrences of that character are

stripped and the following character is taken literally as part of a field value. The exceptions are an escaped '0' or 'N' (for example, \0 or \N if the escape character is '\'). These sequences are interpreted as ASCII NUL (a zero-valued byte) and NULL. The rules for NULL handling are described later in this section.

For more information about '\'-escape syntax, see [Section 9.1, “Literal Values”](#).

In certain cases, field- and line-handling options interact:

- If `LINES TERMINATED BY` is an empty string and `FIELDS TERMINATED BY` is non-empty, lines are also terminated with `FIELDS TERMINATED BY`.
- If the `FIELDS TERMINATED BY` and `FIELDS ENCLOSED BY` values are both empty (' '), a fixed-row (non-delimited) format is used. With fixed-row format, no delimiters are used between fields (but you can still have a line terminator). Instead, column values are read and written using a field width wide enough to hold all values in the field. For `TINYINT`, `SMALLINT`, `MEDIUMINT`, `INT`, and `BIGINT`, the field widths are 4, 6, 8, 11, and 20, respectively, no matter what the declared display width is.

`LINES TERMINATED BY` is still used to separate lines. If a line does not contain all fields, the rest of the columns are set to their default values. If you do not have a line terminator, you should set this to ' '. In this case, the text file must contain all fields for each row.

Fixed-row format also affects handling of NULL values, as described later. Note that fixed-size format does not work if you are using a multi-byte character set.

**Note:** Before MySQL 4.1.12, fixed-row format used the display width of the column. For example, `INT(4)` was read or written using a field with a width of 4. However, if the column contained wider values, they were dumped to their full width, leading to the possibility of a “ragged” field holding values of different widths. Using a field wide enough to hold all values in the field prevents this problem. However, data files written before this change was made might not be reloaded correctly with `LOAD DATA INFILE` for MySQL 4.1.12 and up. This change also affects data files read by `mysqlimport` and written by `mysqldump --tab`, which use `LOAD DATA INFILE` and `SELECT ... INTO outfile`.

Handling of NULL values varies according to the `FIELDS` and `LINES` options in use:

- For the default `FIELDS` and `LINES` values, NULL is written as a field value of \N for output, and a field value of \N is read as NULL for input (assuming that the `ESCAPED BY` character is '\').
- If `FIELDS ENCLOSED BY` is not empty, a field containing the literal word NULL as its value is read as a NULL value. This differs from the word NULL enclosed within `FIELDS ENCLOSED BY` characters, which is read as the string 'NULL'.
- If `FIELDS ESCAPED BY` is empty, NULL is written as the word NULL.
- With fixed-row format (which is used when `FIELDS TERMINATED BY` and `FIELDS ENCLOSED BY` are both empty), NULL is written as an empty string. Note that this causes both NULL values and empty strings in the table to be indistinguishable when written to the file because both are written as empty strings. If you need to be able to tell the two apart when reading the file back in, you should not use fixed-row format.

An attempt to load NULL into a `NOT NULL` column causes assignment of the implicit default value for the column's data type and a warning. Implicit default values are discussed in [Section 11.1.4, “Data Type Default Values”](#).



Some cases are not supported by `LOAD DATA INFILE`:

- Fixed-size rows (`FIELDS TERMINATED BY` and `FIELDS ENCLOSED BY` both empty) and `BLOB` or `TEXT` columns.
- If you specify one separator that is the same as or a prefix of another, `LOAD DATA INFILE` cannot interpret the input properly. For example, the following `FIELDS` clause would cause problems:

```
FIELDS TERMINATED BY '' ENCLOSED BY ''
```

- If `FIELDS ESCAPED BY` is empty, a field value that contains an occurrence of `FIELDS ENCLOSED BY` or `LINES TERMINATED BY` followed by the `FIELDS TERMINATED BY` value causes `LOAD DATA INFILE` to stop reading a field or line too early. This happens because `LOAD DATA INFILE` cannot properly determine where the field or line value ends.

The following example loads all columns of the `persondata` table:

```
LOAD DATA INFILE 'persondata.txt' INTO TABLE persondata;
```

By default, when no column list is provided at the end of the `LOAD DATA INFILE` statement, input lines are expected to contain a field for each table column. If you want to load only some of a table's columns, specify a column list:

```
LOAD DATA INFILE 'persondata.txt' INTO TABLE persondata (col1,col2,...);
```

You must also specify a column list if the order of the fields in the input file differs from the order of the columns in the table. Otherwise, MySQL cannot tell how to match input fields with table columns.

If an input line has too many fields, the extra fields are ignored and the number of warnings is incremented.

If an input line has too few fields, the table columns for which input fields are missing are set to their default values. Default value assignment is described in [Section 11.1.4, “Data Type Default Values”](#).

An empty field value is interpreted differently than if the field value is missing:

- For string types, the column is set to the empty string.
- For numeric types, the column is set to 0.
- For date and time types, the column is set to the appropriate “zero” value for the type. See [Section 11.3, “Date and Time Types”](#).

These are the same values that result if you assign an empty string explicitly to a string, numeric, or date or time type explicitly in an `INSERT` or `UPDATE` statement.

`TIMESTAMP` columns are set to the current date and time only if there is a `NULL` value for the column (that is, `\N`), or if the `TIMESTAMP` column's default value is the current timestamp and it is omitted from the field list when a field list is specified.

`LOAD DATA INFILE` regards all input as strings, so you cannot use numeric values for `ENUM` or `SET` columns the way you can with `INSERT` statements. All `ENUM` and `SET` values must be specified as strings.

When the `LOAD DATA INFILE` statement finishes, it returns an information string in the following format:

```
Records: 1 Deleted: 0 Skipped: 0 Warnings: 0
```

If you are using the C API, you can get information about the statement by calling the `mysql_info()` function. See [Section 17.2.3.33](#), “`mysql_info()`”.

Warnings occur under the same circumstances as when values are inserted via the `INSERT` statement (see [Section 13.2.4](#), “`INSERT Syntax`”), except that `LOAD DATA INFILE` also generates warnings when there are too few or too many fields in the input row. The warnings are not stored anywhere; the number of warnings can be used only as an indication of whether everything went well.

From MySQL 4.1.1 on, you can use `SHOW WARNINGS` to get a list of the first `max_error_count` warnings as information about what went wrong. See [Section 13.5.4.21](#), “`SHOW WARNINGS Syntax`”.

Before MySQL 4.1.1, only a warning count is available to indicate that something went wrong. If you get warnings and want to know exactly why you got them, one way to do this is to dump the table into another file using `SELECT ... INTO OUTFILE` and compare the file to your original input file.

## 13.2.6. REPLACE Syntax

```
REPLACE [LOW_PRIORITY | DELAYED]
  [INTO] tbl_name [(col_name,...)]
  VALUES ({expr | DEFAULT},...),(...),...
```

Or:

```
REPLACE [LOW_PRIORITY | DELAYED]
  [INTO] tbl_name
  SET col_name={expr | DEFAULT}, ...
```

Or:

```
REPLACE [LOW_PRIORITY | DELAYED]
  [INTO] tbl_name [(col_name,...)]
  SELECT ...
```

`REPLACE` works exactly like `INSERT`, except that if an old row in the table has the same value as a new row for a `PRIMARY KEY` or a `UNIQUE` index, the old row is deleted before the new row is inserted. See [Section 13.2.4](#), “`INSERT Syntax`”.

`REPLACE` is a MySQL extension to the SQL standard. It either inserts, or *deletes* and inserts. For another MySQL extension to standard SQL — that either inserts or *updates* — see [Section 13.2.4.3](#), “`INSERT ... ON DUPLICATE KEY UPDATE Syntax`”. `INSERT ... ON DUPLICATE KEY UPDATE` is available as of MySQL 4.1.0.

Note that unless the table has a `PRIMARY KEY` or `UNIQUE` index, using a `REPLACE` statement makes no sense. It becomes equivalent to `INSERT`, because there is no index to be used to determine whether a new row duplicates another.

Values for all columns are taken from the values specified in the `REPLACE` statement. Any missing columns are set to their default values, just as happens for `INSERT`. You cannot refer to values from the current row and use them in the new row. If you use an assignment such as `SET col_name = col_name + 1`, the reference to the column name on the right hand side is treated as `DEFAULT(col_name)`, so the assignment is equivalent to `SET col_name = DEFAULT(col_name) + 1`.

To use `REPLACE`, you must have both the `INSERT` and `DELETE` privileges for the table.

The `REPLACE` statement returns a count to indicate the number of rows affected. This is the sum of the rows deleted and inserted. If the count is 1 for a single-row `REPLACE`, a row was inserted and no rows were deleted. If the count is greater than 1, one or more old rows were deleted before the new row was inserted. It is possible for a single row to replace more than one old row if the table contains multiple unique indexes and the new row duplicates values for different old rows in different unique indexes.

The affected-rows count makes it easy to determine whether `REPLACE` only added a row or whether it also replaced any rows: Check whether the count is 1 (added) or greater (replaced).

If you are using the C API, the affected-rows count can be obtained using the `mysql_affected_rows()` function.

Currently, you cannot replace into a table and select from the same table in a subquery.

MySQL uses the following algorithm for `REPLACE` (and `LOAD DATA . . . REPLACE`):

1. Try to insert the new row into the table
2. While the insertion fails because a duplicate-key error occurs for a primary key or unique index:
  - a. Delete from the table the conflicting row that has the duplicate key value
  - b. Try again to insert the new row into the table

## 13.2.7. `SELECT` Syntax

```
SELECT
  [ALL | DISTINCT | DISTINCTROW ]
  [HIGH_PRIORITY]
  [STRAIGHT_JOIN]
  [SQL_SMALL_RESULT] [SQL_BIG_RESULT] [SQL_BUFFER_RESULT]
  [SQL_CACHE | SQL_NO_CACHE] [SQL_CALC_FOUND_ROWS]
  select_expr, ...
  [FROM table_references]
  [WHERE where_condition]
  [GROUP BY {col_name | expr | position}
  [ASC | DESC], ... [WITH ROLLUP]]
  [HAVING where_condition]
  [ORDER BY {col_name | expr | position}
  [ASC | DESC], ...]
  [LIMIT {[offset,] row_count | row_count OFFSET offset}]
  [PROCEDURE procedure_name(argument_list)]
  [INTO OUTFILE 'file_name' export_options
  | INTO DUMPFILE 'file_name'
  | INTO @var_name [, @var_name]]
  [FOR UPDATE | LOCK IN SHARE MODE]]
```

`SELECT` is used to retrieve rows selected from one or more tables. Support for `UNION` statements and subqueries is available as of MySQL 4.0 and 4.1, respectively. See [Section 13.2.7.2, “UNION Syntax”](#), and [Section 13.2.8, “Subquery Syntax”](#).

The most commonly used clauses of `SELECT` statements are these:

- Each *select\_expr* indicates a column that you want to retrieve. There must be at least one *select\_expr*.
- *table\_references* indicates the table or tables from which to retrieve rows. Its syntax is described in [Section 13.2.7.1, “JOIN Syntax”](#).

- The `WHERE` clause, if given, indicates the condition or conditions that rows must satisfy to be selected. `where_condition` is an expression that evaluates to true for each row to be selected. The statement selects all rows if there is no `WHERE` clause.

In the `WHERE` clause, you can use any of the functions and operators that MySQL supports, except for aggregate (summary) functions. See [Chapter 12, Functions and Operators](#).

`SELECT` can also be used to retrieve rows computed without reference to any table.

For example:

```
mysql> SELECT 1 + 1;
-> 2
```

From MySQL 4.1.0 on, you are allowed to specify `DUAL` as a dummy table name in situations where no tables are referenced:

```
mysql> SELECT 1 + 1 FROM DUAL;
-> 2
```

`DUAL` is purely for the convenience of people who require that all `SELECT` statements should have `FROM` and possibly other clauses. MySQL may ignore the clauses. MySQL does not require `FROM DUAL` if no tables are referenced.

In general, clauses used must be given in exactly the order shown in the syntax description. For example, a `HAVING` clause must come after any `GROUP BY` clause and before any `ORDER BY` clause. The exception is that the `INTO` clause can appear either as shown in the syntax description or immediately preceding the `FROM` clause.

- A `select_expr` can be given an alias using `AS alias_name`. The alias is used as the expression's column name and can be used in `GROUP BY`, `ORDER BY`, or `HAVING` clauses. For example:

```
SELECT CONCAT(last_name,', ',first_name) AS full_name
FROM mytable ORDER BY full_name;
```

The `AS` keyword is optional when aliasing a `select_expr`. The preceding example could have been written like this:

```
SELECT CONCAT(last_name,', ',first_name) full_name
FROM mytable ORDER BY full_name;
```

However, because the `AS` is optional, a subtle problem can occur if you forget the comma between two `select_expr` expressions: MySQL interprets the second as an alias name. For example, in the following statement, `columnb` is treated as an alias name:

```
SELECT columna columnb FROM mytable;
```

For this reason, it is good practice to be in the habit of using `AS` explicitly when specifying column aliases.

- It is not allowable to use a column alias in a `WHERE` clause, because the column value might not yet be determined when the `WHERE` clause is executed. See [Section A.5.4, “Problems with Column Aliases”](#).
- The `FROM table_references` clause indicates the table or tables from which to retrieve rows. If you name more than one table, you are performing a join. For information on join syntax, see [Sec-](#)

tion 13.2.7.1, “JOIN Syntax”. For each table specified, you can optionally specify an alias.

```
tbl_name [[AS] alias]  
[ {USE|IGNORE|FORCE} INDEX (key_list) ]
```

The use of `USE INDEX`, `IGNORE INDEX`, `FORCE INDEX` to give the optimizer hints about how to choose indexes is described in Section 13.2.7.1, “JOIN Syntax”.

In MySQL 4.0.14, you can use `SET max_seeks_for_key=value` as an alternative way to force MySQL to prefer key scans instead of table scans. See Section 5.2.2, “System Variables”.

- You can refer to a table within the default database as *tbl\_name*, or as *db\_name.tbl\_name* to specify a database explicitly. You can refer to a column as *col\_name*, *tbl\_name.col\_name*, or *db\_name.tbl\_name.col\_name*. You need not specify a *tbl\_name* or *db\_name.tbl\_name* prefix for a column reference unless the reference would be ambiguous. See Section 9.2.1, “Identifier Qualifiers”, for examples of ambiguity that require the more explicit column reference forms.
- A table reference can be aliased using *tbl\_name AS alias\_name* or *tbl\_name alias\_name*:

```
SELECT t1.name, t2.salary FROM employee AS t1, info AS t2  
WHERE t1.name = t2.name;  
  
SELECT t1.name, t2.salary FROM employee t1, info t2  
WHERE t1.name = t2.name;
```

- Columns selected for output can be referred to in `ORDER BY` and `GROUP BY` clauses using column names, column aliases, or column positions. Column positions are integers and begin with 1:

```
SELECT college, region, seed FROM tournament  
ORDER BY region, seed;  
  
SELECT college, region AS r, seed AS s FROM tournament  
ORDER BY r, s;  
  
SELECT college, region, seed FROM tournament  
ORDER BY 2, 3;
```

To sort in reverse order, add the `DESC` (descending) keyword to the name of the column in the `ORDER BY` clause that you are sorting by. The default is ascending order; this can be specified explicitly using the `ASC` keyword.

Use of column positions is deprecated because the syntax has been removed from the SQL standard.

- If you use `GROUP BY`, output rows are sorted according to the `GROUP BY` columns as if you had an `ORDER BY` for the same columns. To avoid the overhead of sorting that `GROUP BY` produces, add `ORDER BY NULL`:

```
SELECT a, COUNT(b) FROM test_table GROUP BY a ORDER BY NULL;
```

- MySQL extends the `GROUP BY` clause as of version 3.23.34 so that you can also specify `ASC` and `DESC` after columns named in the clause:

```
SELECT a, COUNT(b) FROM test_table GROUP BY a DESC
```

- MySQL extends the use of `GROUP BY` to allow selecting fields that are not mentioned in the `GROUP BY` clause. If you are not getting the results that you expect from your query, please read the description of `GROUP BY` found in Section 12.10, “Functions and Modifiers for Use with `GROUP BY` Clauses”.

- As of MySQL 4.1.1, `GROUP BY` allows a `WITH ROLLUP` modifier. See [Section 12.10.2, “GROUP BY Modifiers”](#).
- The `HAVING` clause is applied nearly last, just before items are sent to the client, with no optimization. (`LIMIT` is applied after `HAVING`.)

A `HAVING` clause can refer to any column or alias named in a `select_expr` in the `SELECT` list or in outer subqueries, and to aggregate functions. (Standard SQL requires that `HAVING` must reference only columns in the `GROUP BY` clause or columns used in aggregate functions.)

- Do not use `HAVING` for items that should be in the `WHERE` clause. For example, do not write the following:

```
SELECT col_name FROM tbl_name HAVING col_name > 0;
```

Write this instead:

```
SELECT col_name FROM tbl_name WHERE col_name > 0;
```

- The `HAVING` clause can refer to aggregate functions, which the `WHERE` clause cannot:

```
SELECT user, MAX(salary) FROM users
GROUP BY user HAVING MAX(salary) > 10;
```

However, that does not work in older MySQL servers (before version 3.22.5). In those versions, you can use a column alias in the select list and refer to the alias in the `HAVING` clause:

```
SELECT user, MAX(salary) AS max_salary FROM users
GROUP BY user HAVING max_salary>10;
```

- MySQL allows duplicate column names. That is, there can be more than one `select_expr` with the same name. This is an extension to standard SQL. Because MySQL also allows `GROUP BY` and `HAVING` to refer to `select_expr` values, this can result in an ambiguity:

```
SELECT 12 AS a, a FROM t GROUP BY a;
```

In that statement, both columns have the name `a`. To ensure that the correct column is used for grouping, use different names for each `select_expr`.

- When MySQL resolves an unqualified column or alias reference in an `ORDER BY`, `GROUP BY`, or `HAVING` clause, it first searches for the name in the `select_expr` values. If the name is not found, it looks in the columns of the tables named in the `FROM` clause.
- The `LIMIT` clause can be used to constrain the number of rows returned by the `SELECT` statement. `LIMIT` takes one or two numeric arguments, which must both be non-negative integer constants (except when using prepared statements).

With two arguments, the first argument specifies the offset of the first row to return, and the second specifies the maximum number of rows to return. The offset of the initial row is 0 (not 1):

```
SELECT * FROM tbl LIMIT 5,10; # Retrieve rows 6-15
```

To retrieve all rows from a certain offset up to the end of the result set, you can use some large number for the second parameter. This statement retrieves all rows from the 96th row to the last:

```
SELECT * FROM tbl LIMIT 95,18446744073709551615;
```

With one argument, the value specifies the number of rows to return from the beginning of the result set:

```
SELECT * FROM tbl LIMIT 5;      # Retrieve first 5 rows
```

In other words, `LIMIT row_count` is equivalent to `LIMIT 0, row_count`.

For prepared statements, you can use placeholders (supported as of MySQL version 5.0.7). The following statements will return one row from the `tbl` table:

```
SET @a=1;
PREPARE STMT FROM 'SELECT * FROM tbl LIMIT ?';
EXECUTE STMT USING @a;
```

The following statements will return the second to sixth row from the `tbl` table:

```
SET @skip=1; SET @numrows=5;
PREPARE STMT FROM 'SELECT * FROM tbl LIMIT ?, ?';
EXECUTE STMT USING @skip, @numrows;
```

For compatibility with PostgreSQL, MySQL also supports the `LIMIT row_count OFFSET offset` syntax.

- The `SELECT ... INTO OUTFILE 'file_name'` form of `SELECT` writes the selected rows to a file. The file is created on the server host, so you must have the `FILE` privilege to use this syntax. `file_name` cannot be an existing file, which among other things prevents files such as `/etc/passwd` and database tables from being destroyed.

The `SELECT ... INTO OUTFILE` statement is intended primarily to let you very quickly dump a table to a text file on the server machine. If you want to create the resulting file on some client host other than the server host, you cannot use `SELECT ... INTO OUTFILE`. In that case, you should instead use a command such as `mysql -e "SELECT ..." > file_name` to generate the file on the client host.

`SELECT ... INTO OUTFILE` is the complement of `LOAD DATA INFILE`; the syntax for the `export_options` part of the statement consists of the same `FIELDS` and `LINES` clauses that are used with the `LOAD DATA INFILE` statement. See [Section 13.2.5, “LOAD DATA INFILE Syntax”](#).

`FIELDS ESCAPED BY` controls how to write special characters. If the `FIELDS ESCAPED BY` character is not empty, it is used as a prefix that precedes following characters on output:

- The `FIELDS ESCAPED BY` character
- The `FIELDS [OPTIONALLY] ENCLOSED BY` character
- The first character of the `FIELDS TERMINATED BY` and `LINES TERMINATED BY` values
- ASCII `NUL` (the zero-valued byte; what is actually written following the escape character is ASCII `'0'`, not a zero-valued byte)

The `FIELDS TERMINATED BY`, `ENCLOSED BY`, `ESCAPED BY`, or `LINES TERMINATED BY` characters *must* be escaped so that you can read the file back in reliably. ASCII `NUL` is escaped to make it easier to view with some paggers.

The resulting file does not have to conform to SQL syntax, so nothing else need be escaped.

If the `FIELDS ESCAPED BY` character is empty, no characters are escaped and `NULL` is output as

`NULL`, not `\N`. It is probably not a good idea to specify an empty escape character, particularly if field values in your data contain any of the characters in the list just given.

Here is an example that produces a file in the comma-separated values (CSV) format used by many programs:

```
SELECT a,b,a+b INTO OUTFILE '/tmp/result.txt'
  FIELDS TERMINATED BY ',' OPTIONALLY ENCLOSED BY '"'
  LINES TERMINATED BY '\n'
FROM test_table;
```

- If you use `INTO DUMPFILE` instead of `INTO OUTFILE`, MySQL writes only one row into the file, without any column or line termination and without performing any escape processing. This is useful if you want to store a `BLOB` value in a file.
- The `INTO` clause can name a list of one or more user-defined variables. The selected values are assigned to the variables. The number of variables must match the number of columns.
- **Note:** Any file created by `INTO OUTFILE` or `INTO DUMPFILE` is writable by all users on the server host. The reason for this is that the MySQL server cannot create a file that is owned by anyone other than the user under whose account it is running. (You should *never* run `mysqld` as `root` for this and other reasons.) The file thus must be world-writable so that you can manipulate its contents.
- The `SELECT` syntax description at the beginning of this section shows the `INTO` clause near the end of the statement. It is also possible to use `INTO OUTFILE` or `INTO DUMPFILE` immediately preceding the `FROM` clause.
- A `PROCEDURE` clause names a procedure that should process the data in the result set. For an example, see [Section 19.3.1, “Procedure Analyse”](#).
- If you use `FOR UPDATE` with a storage engine that uses page or row locks, rows examined by the query are write-locked until the end of the current transaction. Using `LOCK IN SHARE MODE` sets a shared lock that allows other transactions to read the examined rows but not to update or delete them. See [Section 14.2.11.5, “SELECT ... FOR UPDATE and SELECT ... LOCK IN SHARE MODE Locking Reads”](#).

Following the `SELECT` keyword, you can use a number of options that affect the operation of the statement.

The `ALL`, `DISTINCT`, and `DISTINCTROW` options specify whether duplicate rows should be returned. If none of these options are given, the default is `ALL` (all matching rows are returned). `DISTINCT` and `DISTINCTROW` are synonyms and specify removal of duplicate rows from the result set.

`HIGH_PRIORITY`, `STRAIGHT_JOIN`, and options beginning with `SQL_` are MySQL extensions to standard SQL.

- `HIGH_PRIORITY` gives the `SELECT` higher priority than a statement that updates a table. You should use this only for queries that are very fast and must be done at once. A `SELECT HIGH_PRIORITY` query that is issued while the table is locked for reading runs even if there is an update statement waiting for the table to be free.

`HIGH_PRIORITY` cannot be used with `SELECT` statements that are part of a `UNION`.

- `STRAIGHT_JOIN` forces the optimizer to join the tables in the order in which they are listed in the `FROM` clause. You can use this to speed up a query if the optimizer joins the tables in non-optimal order. See [Section 7.2.1, “Optimizing Queries with EXPLAIN”](#). `STRAIGHT_JOIN` also can be used



in the *table\_references* list. See [Section 13.2.7.1, “JOIN Syntax”](#).

- `SQL_BIG_RESULT` can be used with `GROUP BY` or `DISTINCT` to tell the optimizer that the result set has many rows. In this case, MySQL directly uses disk-based temporary tables if needed, and prefers sorting to using a temporary table with a key on the `GROUP BY` elements.
- `SQL_BUFFER_RESULT` forces the result to be put into a temporary table. This helps MySQL free the table locks early and helps in cases where it takes a long time to send the result set to the client.
- `SQL_SMALL_RESULT` can be used with `GROUP BY` or `DISTINCT` to tell the optimizer that the result set is small. In this case, MySQL uses fast temporary tables to store the resulting table instead of using sorting. In MySQL 3.23 and up, this should not normally be needed.
- `SQL_CALC_FOUND_ROWS` (available in MySQL 4.0.0 and up) tells MySQL to calculate how many rows there would be in the result set, disregarding any `LIMIT` clause. The number of rows can then be retrieved with `SELECT FOUND_ROWS()`. See [Section 12.9.3, “Information Functions”](#).

Before MySQL 4.1.0, this option does not work with `LIMIT 0`, which is optimized to return instantly (resulting in a row count of 0). See [Section 7.2.11, “LIMIT Optimization”](#).

- `SQL_CACHE` tells MySQL to store the query result in the query cache if you are using a `query_cache_type` value of 2 or `DEMAND`. For a query that uses `UNION` or subqueries, this option effects any `SELECT` in the query. See [Section 5.13, “The MySQL Query Cache”](#).
- `SQL_NO_CACHE` tells MySQL not to store the query result in the query cache. See [Section 5.13, “The MySQL Query Cache”](#). For a query that uses `UNION` or subqueries, this option effects any `SELECT` in the query.

### 13.2.7.1. JOIN Syntax

MySQL supports the following `JOIN` syntaxes for the *table\_references* part of `SELECT` statements and multiple-table `DELETE` and `UPDATE` statements:

```
table_reference, table_reference
table_reference [INNER | CROSS] JOIN table_reference [join_condition]
table_reference STRAIGHT_JOIN table_reference
table_reference LEFT [OUTER] JOIN table_reference join_condition
table_reference NATURAL [LEFT [OUTER]] JOIN table_reference
{ ON table_reference LEFT OUTER JOIN table_reference
  ON conditional_expr }
table_reference RIGHT [OUTER] JOIN table_reference join_condition
table_reference NATURAL [RIGHT [OUTER]] JOIN table_reference
```

*table\_reference* is defined as:

```
tbl_name [[AS] alias]
[ {USE | IGNORE | FORCE} INDEX (key_list) ]
```

*join\_condition* is defined as:

```
ON conditional_expr | USING (column_list)
```

Note that several changes in join processing were made in MySQL 5.0.12 to make MySQL more compliant with standard SQL. These changes include the ability to handle nested joins (including outer joins) according to the standard. If a nested join returns results that are not what you expect, please consider upgrading to MySQL 5.0. Further details about the changes in join processing can be found at <http://dev.mysql.com/doc/refman/5.0/en/join.html>.

You should generally not have any conditions in the `ON` part that are used to restrict which rows you

want in the result set, but rather specify these conditions in the `WHERE` clause. There are exceptions to this rule.

Note that `INNER JOIN` syntax allows a `join_condition` only from MySQL 3.23.17 on. The same is true for `JOIN` and `CROSS JOIN` only as of MySQL 4.0.11.

The `{ OJ ... LEFT OUTER JOIN ... }` syntax shown in the preceding list exists only for compatibility with ODBC. The curly braces in the syntax should be written literally; they are not metasyntax as used elsewhere in syntax descriptions.

- A table reference can be aliased using `tbl_name AS alias_name` or `tbl_name alias_name`:

```
SELECT t1.name, t2.salary FROM employee AS t1, info AS t2
WHERE t1.name = t2.name;

SELECT t1.name, t2.salary FROM employee t1, info t2
WHERE t1.name = t2.name;
```

- The `ON` conditional is any conditional expression of the form that can be used in a `WHERE` clause.
- If there is no matching row for the right table in the `ON` or `USING` part in a `LEFT JOIN`, a row with all columns set to `NULL` is used for the right table. You can use this fact to find rows in a table that have no counterpart in another table:

```
SELECT table1.* FROM table1
LEFT JOIN table2 ON table1.id=table2.id
WHERE table2.id IS NULL;
```

This example finds all rows in `table1` with an `id` value that is not present in `table2` (that is, all rows in `table1` with no corresponding row in `table2`). This assumes that `table2.id` is declared `NOT NULL`. See [Section 7.2.8, “LEFT JOIN and RIGHT JOIN Optimization”](#).

- The `USING(column_list)` clause names a list of columns that must exist in both tables. The following two clauses are semantically identical:

```
a LEFT JOIN b USING (c1,c2,c3)
a LEFT JOIN b ON a.c1=b.c1 AND a.c2=b.c2 AND a.c3=b.c3
```

- The `NATURAL [LEFT] JOIN` of two tables is defined to be semantically equivalent to an `INNER JOIN` or a `LEFT JOIN` with a `USING` clause that names all columns that exist in both tables.
- `INNER JOIN` and `,` (comma) are semantically equivalent in the absence of a join condition: both produce a Cartesian product between the specified tables (that is, each and every row in the first table is joined to each and every row in the second table).
- `RIGHT JOIN` works analogously to `LEFT JOIN`. To keep code portable across databases, it is recommended that you use `LEFT JOIN` instead of `RIGHT JOIN`.
- `STRAIGHT_JOIN` is identical to `JOIN`, except that the left table is always read before the right table. This can be used for those (few) cases for which the join optimizer puts the tables in the wrong order.

As of MySQL 3.23.12, you can give hints about which index MySQL should use when retrieving information from a table. By specifying `USE INDEX (key_list)`, you can tell MySQL to use only one of the possible indexes to find rows in the table. The alternative syntax `IGNORE INDEX (key_list)` can be used to tell MySQL to not use some particular index. These hints are useful if `EXPLAIN` shows that MySQL is using the wrong index from the list of possible indexes.

From MySQL 4.0.9 on, you can also use `FORCE INDEX`, which acts like `USE INDEX (key_list)` but with the addition that a table scan is assumed to be *very* expensive. In other words, a table scan is used only if there is no way to use one of the given indexes to find rows in the table.

`USE INDEX`, `IGNORE INDEX`, and `FORCE INDEX` affect only which indexes are used when MySQL decides how to find rows in the table and how to do the join. They do not affect whether an index is used when resolving an `ORDER BY` or `GROUP BY`.

`USE KEY`, `IGNORE KEY`, and `FORCE KEY` are synonyms for `USE INDEX`, `IGNORE INDEX`, and `FORCE INDEX`.

Some join examples:

```
SELECT * FROM table1,table2 WHERE table1.id=table2.id;
SELECT * FROM table1 LEFT JOIN table2 ON table1.id=table2.id;
SELECT * FROM table1 LEFT JOIN table2 USING (id);
SELECT * FROM table1 LEFT JOIN table2 ON table1.id=table2.id
LEFT JOIN table3 ON table2.id=table3.id;
SELECT * FROM table1 USE INDEX (key1,key2)
WHERE key1=1 AND key2=2 AND key3=3;
SELECT * FROM table1 IGNORE INDEX (key3)
WHERE key1=1 AND key2=2 AND key3=3;
```

### 13.2.7.2. UNION Syntax

```
SELECT ...
UNION [ALL | DISTINCT] SELECT ...
[UNION [ALL | DISTINCT] SELECT ...]
```

`UNION` is used to combine the result from multiple `SELECT` statements into a single result set. `UNION` is available from MySQL 4.0.0 on.

The column names from the first `SELECT` statement are used as the column names for the results returned. Selected columns listed in corresponding positions of each `SELECT` statement should have the same data type. (For example, the first column selected by the first statement should have the same type as the first column selected by the other statements.)

As of MySQL 4.1.1, if the data types of corresponding `SELECT` columns do not match, the types and lengths of the columns in the `UNION` result take into account the values retrieved by all of the `SELECT` statements. For example, consider the following:

```
mysql> SELECT REPEAT('a',1) UNION SELECT REPEAT('b',10);
+-----+
| REPEAT('a',1) |
+-----+
| a              |
| bbbbbbbbbbb   |
+-----+
```

Before MySQL 4.1.1, only the type and length from the first `SELECT` would have been used and the second row would have been truncated to a length of 1:

```
mysql> SELECT REPEAT('a',1) UNION SELECT REPEAT('b',10);
+-----+
| REPEAT('a',1) |
+-----+
| a              |
| b              |
+-----+
```

The `SELECT` statements are normal select statements, but with the following restrictions:

- Only the last `SELECT` statement can use `INTO OUTFILE`.
- `HIGH_PRIORITY` cannot be used with `SELECT` statements that are part of a `UNION`. If you specify it for the first `SELECT`, it has no effect. If you specify it for any subsequent `SELECT` statements, a syntax error results.

The default behavior for `UNION` is that duplicate rows are removed from the result. The optional `DISTINCT` keyword (introduced in MySQL 4.0.17) has no effect other than the default because it also specifies duplicate-row removal. With the optional `ALL` keyword, duplicate-row removal does not occur and the result includes all matching rows from all the `SELECT` statements.

Before MySQL 4.1.2, you cannot mix `UNION ALL` and `UNION DISTINCT` in the same query. If you use `ALL` for one `UNION`, it is used for all of them. As of MySQL 4.1.2, mixed `UNION` types are treated such that a `DISTINCT` union overrides any `ALL` union to its left. A `DISTINCT` union can be produced explicitly by using `UNION DISTINCT` or implicitly by using `UNION` with no following `DISTINCT` or `ALL` keyword.

To use an `ORDER BY` or `LIMIT` clause to sort or limit the entire `UNION` result, parenthesize the individual `SELECT` statements and place the `ORDER BY` or `LIMIT` after the last one. The following example uses both clauses:

```
(SELECT a FROM t1 WHERE a=10 AND B=1)
UNION
(SELECT a FROM t2 WHERE a=11 AND B=2)
ORDER BY a LIMIT 10;
```

This kind of `ORDER BY` cannot use column references that include a table name (that is, names in `tbl_name.col_name` format). Instead, provide a column alias in the first `SELECT` statement and refer to the alias in the `ORDER BY`. (Alternatively, refer to the column in the `ORDER BY` using its column position. However, use of column positions is deprecated.)

Also, if a column to be sorted is aliased, the `ORDER BY` clause *must* refer to the alias, not the column name. The first of the following statements will work, but the second will fail with an `Unknown column 'a' in 'order clause'` error:

```
(SELECT a AS b FROM t) UNION (SELECT ...) ORDER BY b;
(SELECT a AS b FROM t) UNION (SELECT ...) ORDER BY a;
```

To apply `ORDER BY` or `LIMIT` to an individual `SELECT`, place the clause inside the parentheses that enclose the `SELECT`:

```
(SELECT a FROM t1 WHERE a=10 AND B=1 ORDER BY a LIMIT 10)
UNION
(SELECT a FROM t2 WHERE a=11 AND B=2 ORDER BY a LIMIT 10);
```

Use of `ORDER BY` for individual `SELECT` statements implies nothing about the order in which the rows appear in the final result because `UNION` by default produces an unordered set of rows. If `ORDER BY` appears with `LIMIT`, it is used to determine the subset of the selected rows to retrieve for the `SELECT`, but does not necessarily affect the order of those rows in the final `UNION` result. If `ORDER BY` appears without `LIMIT` in a `SELECT`, it is optimized away because it will have no effect anyway.

To cause rows in a `UNION` result to consist of the sets of rows retrieved by each `SELECT` one after the other, select an additional column in each `SELECT` to use as a sort column and add an `ORDER BY` following the last `SELECT`:

```
(SELECT 1 AS sort_col, colla, collb, ... FROM t1)
```

```
UNION
(SELECT 2, col2a, col2b, ... FROM t2) ORDER BY sort_col;
```

To additionally maintain sort order within individual `SELECT` results, add a secondary column to the `ORDER BY` clause:

```
(SELECT 1 AS sort_col, col1a, col1b, ... FROM t1)
UNION
(SELECT 2, col2a, col2b, ... FROM t2) ORDER BY sort_col, col1a;
```

## 13.2.8. Subquery Syntax

A subquery is a `SELECT` statement within another statement.

Starting with MySQL 4.1, all subquery forms and operations that the SQL standard requires are supported, as well as a few features that are MySQL-specific.

With MySQL versions prior to 4.1, it was necessary to work around or avoid the use of subqueries. In many cases, subqueries can successfully be rewritten using joins and other methods. See [Section 13.2.8.11, “Rewriting Subqueries as Joins for Earlier MySQL Versions”](#).

Here is an example of a subquery:

```
SELECT * FROM t1 WHERE column1 = (SELECT column1 FROM t2);
```

In this example, `SELECT * FROM t1 ...` is the *outer query* (or *outer statement*), and `(SELECT column1 FROM t2)` is the *subquery*. We say that the subquery is *nested* within the outer query, and in fact it is possible to nest subqueries within other subqueries, to a considerable depth. A subquery must always appear within parentheses.

The main advantages of subqueries are:

- They allow queries that are *structured* so that it is possible to isolate each part of a statement.
- They provide alternative ways to perform operations that would otherwise require complex joins and unions.
- They are, in many people's opinion, readable. Indeed, it was the innovation of subqueries that gave people the original idea of calling the early SQL “Structured Query Language.”

Here is an example statement that shows the major points about subquery syntax as specified by the SQL standard and supported in MySQL:

```
DELETE FROM t1
WHERE s11 > ANY
  (SELECT COUNT(*) /* no hint */ FROM t2
  WHERE NOT EXISTS
   (SELECT * FROM t3
    WHERE ROW(5*t2.s1,77)=
     (SELECT 50,11*s1 FROM t4 UNION SELECT 50,77 FROM
      (SELECT * FROM t5) AS t5)));
```

A subquery can return a scalar (a single value), a single row, a single column, or a table (one or more rows of one or more columns). These are called scalar, column, row, and table subqueries. Subqueries that return a particular kind of result often can be used only in certain contexts, as described in the following sections.

There are few restrictions on the type of statements in which subqueries can be used. A subquery can

contain any of the keywords or clauses that an ordinary `SELECT` can contain: `DISTINCT`, `GROUP BY`, `ORDER BY`, `LIMIT`, joins, index hints, `UNION` constructs, comments, functions, and so on.

One restriction is that a subquery's outer statement must be one of: `SELECT`, `INSERT`, `UPDATE`, `DELETE`, `SET`, or `DO`. Another restriction is that currently you cannot modify a table and select from the same table in a subquery. This applies to statements such as `DELETE`, `INSERT`, `REPLACE`, and `UPDATE`.

A more comprehensive discussion of restrictions on subquery use, including performance issues for certain forms of subquery syntax, is given in [Section I.1, “Restrictions on Subqueries”](#).

### 13.2.8.1. The Subquery as Scalar Operand

In its simplest form, a subquery is a scalar subquery that returns a single value. A scalar subquery is a simple operand, and you can use it almost anywhere a single column value or literal is legal, and you can expect it to have those characteristics that all operands have: a data type, a length, an indication whether it can be `NULL`, and so on. For example:

```
CREATE TABLE t1 (s1 INT, s2 CHAR(5) NOT NULL);
INSERT INTO t1 VALUES(100, 'abcde');
SELECT (SELECT s2 FROM t1);
```

The subquery in this `SELECT` returns a single value ('abcde') that has a data type of `CHAR`, a length of 5, a character set and collation equal to the defaults in effect at `CREATE TABLE` time, and an indication that the value in the column can be `NULL`. In fact, almost all subqueries can be `NULL`. If the table used in the example were empty, the value of the subquery would be `NULL`.

There are a few contexts in which a scalar subquery cannot be used. If a statement allows only a literal value, you cannot use a subquery. For example, `LIMIT` requires literal integer arguments, and `LOAD DATA INFILE` requires a literal string filename. You cannot use subqueries to supply these values.

When you see examples in the following sections that contain the rather spartan construct (`SELECT column1 FROM t1`), imagine that your own code contains much more diverse and complex constructions.

Suppose that we make two tables:

```
CREATE TABLE t1 (s1 INT);
INSERT INTO t1 VALUES (1);
CREATE TABLE t2 (s1 INT);
INSERT INTO t2 VALUES (2);
```

Then perform a `SELECT`:

```
SELECT (SELECT s1 FROM t2) FROM t1;
```

The result is 2 because there is a row in `t2` containing a column `s1` that has a value of 2.

A scalar subquery can be part of an expression, but remember the parentheses, even if the subquery is an operand that provides an argument for a function. For example:

```
SELECT UPPER((SELECT s1 FROM t1)) FROM t2;
```

### 13.2.8.2. Comparisons Using Subqueries

The most common use of a subquery is in the form:

```
non_subquery_operand comparison_operator (subquery)
```

Where *comparison\_operator* is one of these operators:

```
= > < >= <= <>
```

For example:

```
... 'a' = (SELECT column1 FROM t1)
```

At one time the only legal place for a subquery was on the right side of a comparison, and you might still find some old DBMSs that insist on this.

Here is an example of a common-form subquery comparison that you cannot do with a join. It finds all the values in table *t1* that are equal to a maximum value in table *t2*:

```
SELECT column1 FROM t1
WHERE column1 = (SELECT MAX(column2) FROM t2);
```

Here is another example, which again is impossible with a join because it involves aggregating for one of the tables. It finds all rows in table *t1* containing a value that occurs twice in a given column:

```
SELECT * FROM t1 AS t
WHERE 2 = (SELECT COUNT(*) FROM t1 WHERE t1.id = t.id);
```

For a comparison performed with one of these operators, the subquery must return a scalar, with the exception that = can be used with row subqueries. See [Section 13.2.8.5, “Row Subqueries”](#).

### 13.2.8.3. Subqueries with **ANY**, **IN**, and **SOME**

Syntax:

```
operand comparison_operator ANY (subquery)
operand IN (subquery)
operand comparison_operator SOME (subquery)
```

The **ANY** keyword, which must follow a comparison operator, means “return **TRUE** if the comparison is **TRUE** for **ANY** of the values in the column that the subquery returns.” For example:

```
SELECT s1 FROM t1 WHERE s1 > ANY (SELECT s1 FROM t2);
```

Suppose that there is a row in table *t1* containing (10). The expression is **TRUE** if table *t2* contains (21, 14, 7) because there is a value 7 in *t2* that is less than 10. The expression is **FALSE** if table *t2* contains (20, 10), or if table *t2* is empty. The expression is **UNKNOWN** if table *t2* contains (NULL, NULL, NULL).

The word **IN** is an alias for = **ANY**. Thus, these two statements are the same:

```
SELECT s1 FROM t1 WHERE s1 = ANY (SELECT s1 FROM t2);
SELECT s1 FROM t1 WHERE s1 IN (SELECT s1 FROM t2);
```

However, **NOT IN** is not an alias for <> **ANY**, but for <> **ALL**. See [Section 13.2.8.4, “Subqueries with ALL”](#).

The word **SOME** is an alias for **ANY**. Thus, these two statements are the same:

```
SELECT s1 FROM t1 WHERE s1 <> ANY (SELECT s1 FROM t2);
SELECT s1 FROM t1 WHERE s1 <> SOME (SELECT s1 FROM t2);
```

Use of the word **SOME** is rare, but this example shows why it might be useful. To most people's ears, the English phrase “a is not equal to any b” means “there is no b which is equal to a,” but that is not what is meant by the SQL syntax. The syntax means “there is some b to which a is not equal.” Using **<> SOME** instead helps ensure that everyone understands the true meaning of the query.

### 13.2.8.4. Subqueries with **ALL**

Syntax:

```
operand comparison_operator ALL (subquery)
```

The word **ALL**, which must follow a comparison operator, means “return **TRUE** if the comparison is **TRUE** for **ALL** of the values in the column that the subquery returns.” For example:

```
SELECT s1 FROM t1 WHERE s1 > ALL (SELECT s1 FROM t2);
```

Suppose that there is a row in table **t1** containing **(10)**. The expression is **TRUE** if table **t2** contains **(-5, 0, +5)** because **10** is greater than all three values in **t2**. The expression is **FALSE** if table **t2** contains **(12, 6, NULL, -100)** because there is a single value **12** in table **t2** that is greater than **10**. The expression is *unknown* (that is, **NULL**) if table **t2** contains **(0, NULL, 1)**.

Finally, if table **t2** is empty, the result is **TRUE**. So, the following statement is **TRUE** when table **t2** is empty:

```
SELECT * FROM t1 WHERE 1 > ALL (SELECT s1 FROM t2);
```

But this statement is **NULL** when table **t2** is empty:

```
SELECT * FROM t1 WHERE 1 > (SELECT s1 FROM t2);
```

In addition, the following statement is **NULL** when table **t2** is empty:

```
SELECT * FROM t1 WHERE 1 > ALL (SELECT MAX(s1) FROM t2);
```

In general, *tables containing NULL values* and *empty tables* are “edge cases.” When writing subquery code, always consider whether you have taken those two possibilities into account.

**NOT IN** is an alias for **<> ALL**. Thus, these two statements are the same:

```
SELECT s1 FROM t1 WHERE s1 <> ALL (SELECT s1 FROM t2);
SELECT s1 FROM t1 WHERE s1 NOT IN (SELECT s1 FROM t2);
```

### 13.2.8.5. Row Subqueries

The discussion to this point has been of scalar or column subqueries; that is, subqueries that return a single value or a column of values. A *row subquery* is a subquery variant that returns a single row and can thus return more than one column value. Here are two examples:

```
SELECT * FROM t1 WHERE (1,2) = (SELECT column1, column2 FROM t2);
SELECT * FROM t1 WHERE ROW(1,2) = (SELECT column1, column2 FROM t2);
```

The queries here are both **TRUE** if table **t2** has a row where **column1 = 1** and **column2 = 2**.

The expressions **(1, 2)** and **ROW(1, 2)** are sometimes called *row constructors*. The two are equivalent. They are legal in other contexts as well. For example, the following two statements are semantically equivalent (although in MySQL 4.1 only the second one can be optimized):



```
SELECT * FROM t1 WHERE (column1,column2) = (1,1);
SELECT * FROM t1 WHERE column1 = 1 AND column2 = 1;
```

The normal use of row constructors is for comparisons with subqueries that return two or more columns. For example, the following query answers the request, “find all rows in table `t1` that also exist in table `t2`”:

```
SELECT column1,column2,column3
FROM t1
WHERE (column1,column2,column3) IN
      (SELECT column1,column2,column3 FROM t2);
```

### 13.2.8.6. EXISTS and NOT EXISTS

If a subquery returns any rows at all, `EXISTS subquery` is `TRUE`, and `NOT EXISTS subquery` is `FALSE`. For example:

```
SELECT column1 FROM t1 WHERE EXISTS (SELECT * FROM t2);
```

Traditionally, an `EXISTS` subquery starts with `SELECT *`, but it could begin with `SELECT 5` or `SELECT column1` or anything at all. MySQL ignores the `SELECT` list in such a subquery, so it makes no difference.

For the preceding example, if `t2` contains any rows, even rows with nothing but `NULL` values, the `EXISTS` condition is `TRUE`. This is actually an unlikely example because a `[NOT] EXISTS` subquery almost always contains correlations. Here are some more realistic examples:

- What kind of store is present in one or more cities?

```
SELECT DISTINCT store_type FROM stores
WHERE EXISTS (SELECT * FROM cities_stores
              WHERE cities_stores.store_type = stores.store_type);
```

- What kind of store is present in no cities?

```
SELECT DISTINCT store_type FROM stores
WHERE NOT EXISTS (SELECT * FROM cities_stores
                  WHERE cities_stores.store_type = stores.store_type);
```

- What kind of store is present in all cities?

```
SELECT DISTINCT store_type FROM stores s1
WHERE NOT EXISTS (
  SELECT * FROM cities WHERE NOT EXISTS (
    SELECT * FROM cities_stores
    WHERE cities_stores.city = cities.city
    AND cities_stores.store_type = stores.store_type));
```

The last example is a double-nested `NOT EXISTS` query. That is, it has a `NOT EXISTS` clause within a `NOT EXISTS` clause. Formally, it answers the question “does a city exist with a store that is not in `Stores`”? But it is easier to say that a nested `NOT EXISTS` answers the question “is `x TRUE` for all `y`?”

### 13.2.8.7. Correlated Subqueries

A *correlated subquery* is a subquery that contains a reference to a table that also appears in the outer query. For example:

```
SELECT * FROM t1 WHERE column1 = ANY
      (SELECT column1 FROM t2 WHERE t2.column2 = t1.column2);
```

Notice that the subquery contains a reference to a column of `t1`, even though the subquery's `FROM` clause does not mention a table `t1`. So, MySQL looks outside the subquery, and finds `t1` in the outer query.

Suppose that table `t1` contains a row where `column1 = 5` and `column2 = 6`; meanwhile, table `t2` contains a row where `column1 = 5` and `column2 = 7`. The simple expression `... WHERE column1 = ANY (SELECT column1 FROM t2)` would be `TRUE`, but in this example, the `WHERE` clause within the subquery is `FALSE` (because `(5, 6)` is not equal to `(5, 7)`), so the subquery as a whole is `FALSE`.

**Scoping rule:** MySQL evaluates from inside to outside. For example:

```
SELECT column1 FROM t1 AS x
      WHERE x.column1 = (SELECT column1 FROM t2 AS x
                        WHERE x.column1 = (SELECT column1 FROM t3
                                          WHERE x.column2 = t3.column1));
```

In this statement, `x.column2` must be a column in table `t2` because `SELECT column1 FROM t2 AS x ...` renames `t2`. It is not a column in table `t1` because `SELECT column1 FROM t1 ...` is an outer query that is *farther out*.

For subqueries in `HAVING` or `ORDER BY` clauses, MySQL also looks for column names in the outer select list.

For certain cases, a correlated subquery is optimized. For example:

```
val IN (SELECT key_val FROM tbl_name WHERE correlated_condition)
```

Otherwise, they are inefficient and likely to be slow. Rewriting the query as a join might improve performance.

Correlated subqueries cannot refer to the results of aggregate functions from the outer query.

### 13.2.8.8. Subqueries in the `FROM` clause

Subqueries are legal in a `SELECT` statement's `FROM` clause. The actual syntax is:

```
SELECT ... FROM (subquery) [AS] name ...
```

The `[AS] name` clause is mandatory, because every table in a `FROM` clause must have a name. Any columns in the `subquery` select list must have unique names. You can find this syntax described elsewhere in this manual, where the term used is “derived tables.”

For the sake of illustration, assume that you have this table:

```
CREATE TABLE t1 (s1 INT, s2 CHAR(5), s3 FLOAT);
```

Here is how to use a subquery in the `FROM` clause, using the example table:

```
INSERT INTO t1 VALUES (1, '1', 1.0);
INSERT INTO t1 VALUES (2, '2', 2.0);
SELECT sb1, sb2, sb3
      FROM (SELECT s1 AS sb1, s2 AS sb2, s3*2 AS sb3 FROM t1) AS sb
      WHERE sb1 > 1;
```

Result: 2, '2', 4.0.

Here is another example: Suppose that you want to know the average of a set of sums for a grouped table. This does not work:

```
SELECT AVG(SUM(column1)) FROM t1 GROUP BY column1;
```

However, this query provides the desired information:

```
SELECT AVG(sum_column1)
       FROM (SELECT SUM(column1) AS sum_column1
             FROM t1 GROUP BY column1) AS t1;
```

Notice that the column name used within the subquery (`sum_column1`) is recognized in the outer query.

Subqueries in the `FROM` clause can return a scalar, column, row, or table. Subqueries in the `FROM` clause cannot be correlated subqueries.

Subqueries in the `FROM` clause are executed even for the `EXPLAIN` statement (that is, derived temporary tables are built). This occurs because upper level queries need information about all tables during optimization phase.

### 13.2.8.9. Subquery Errors

There are some errors that apply only to subqueries. This section describes them.

- Unsupported subquery syntax:

```
ERROR 1235 (ER_NOT_SUPPORTED_YET)
SQLSTATE = 42000
Message = "This version of MySQL does not yet support
'LIMIT & IN/ALL/ANY/SOME subquery'"
```

This means that statements of the following form do not work yet:

```
SELECT * FROM t1 WHERE s1 IN (SELECT s2 FROM t2 ORDER BY s1 LIMIT 1)
```

- Incorrect number of columns from subquery:

```
ERROR 1241 (ER_OPERAND_COL)
SQLSTATE = 21000
Message = "Operand should contain 1 column(s)"
```

This error occurs in cases like this:

```
SELECT (SELECT column1, column2 FROM t2) FROM t1;
```

You may use a subquery that returns multiple columns, if the purpose is comparison. See [Section 13.2.8.5, “Row Subqueries”](#). However, in other contexts, the subquery must be a scalar operand.

- Incorrect number of rows from subquery:

```
ERROR 1242 (ER_SUBSELECT_NO_1_ROW)
SQLSTATE = 21000
Message = "Subquery returns more than 1 row"
```

This error occurs for statements where the subquery returns more than one row. Consider the following example:

```
SELECT * FROM t1 WHERE column1 = (SELECT column1 FROM t2);
```

If `SELECT column1 FROM t2` returns just one row, the previous query will work. If the subquery returns more than one row, error 1242 will occur. In that case, the query should be rewritten as:

```
SELECT * FROM t1 WHERE column1 = ANY (SELECT column1 FROM t2);
```

- Incorrectly used table in subquery:

```
Error 1093 (ER_UPDATE_TABLE_USED)
SQLSTATE = HY000
Message = "You can't specify target table 'x'
for update in FROM clause"
```

This error occurs in cases such as the following:

```
UPDATE t1 SET column2 = (SELECT MAX(column1) FROM t1);
```

You can use a subquery for assignment within an `UPDATE` statement because subqueries are legal in `UPDATE` and `DELETE` statements as well as in `SELECT` statements. However, you cannot use the same table (in this case, table `t1`) for both the subquery's `FROM` clause and the update target.

For transactional storage engines, the failure of a subquery causes the entire statement to fail. For non-transactional storage engines, data modifications made before the error was encountered are preserved.

### 13.2.8.10. Optimizing Subqueries

Development is ongoing, so no optimization tip is reliable for the long term. The following list provides some interesting tricks that you might want to play with:

- Use subquery clauses that affect the number or order of the rows in the subquery. For example:

```
SELECT * FROM t1 WHERE t1.column1 IN
  (SELECT column1 FROM t2 ORDER BY column1);
SELECT * FROM t1 WHERE t1.column1 IN
  (SELECT DISTINCT column1 FROM t2);
SELECT * FROM t1 WHERE EXISTS
  (SELECT * FROM t2 LIMIT 1);
```

- Replace a join with a subquery. For example, try this:

```
SELECT DISTINCT column1 FROM t1 WHERE t1.column1 IN (
  SELECT column1 FROM t2);
```

Instead of this:

```
SELECT DISTINCT t1.column1 FROM t1, t2
WHERE t1.column1 = t2.column1;
```

- Some subqueries can be transformed to joins for compatibility with older versions of MySQL before 4.1 that do not support subqueries. However, in some cases, converting a subquery to a join may also improve performance. See [Section 13.2.8.11, “Rewriting Subqueries as Joins for Earlier MySQL](#)

**Versions”.**

- Move clauses from outside to inside the subquery. For example, use this query:

```
SELECT * FROM t1
WHERE s1 IN (SELECT s1 FROM t1 UNION ALL SELECT s1 FROM t2);
```

Instead of this query:

```
SELECT * FROM t1
WHERE s1 IN (SELECT s1 FROM t1) OR s1 IN (SELECT s1 FROM t2);
```

For another example, use this query:

```
SELECT (SELECT column1 + 5 FROM t1) FROM t2;
```

Instead of this query:

```
SELECT (SELECT column1 FROM t1) + 5 FROM t2;
```

- Use a row subquery instead of a correlated subquery. For example, use this query:

```
SELECT * FROM t1
WHERE (column1,column2) IN (SELECT column1,column2 FROM t2);
```

Instead of this query:

```
SELECT * FROM t1
WHERE EXISTS (SELECT * FROM t2 WHERE t2.column1=t1.column1
AND t2.column2=t1.column2);
```

- Use `NOT (a = ANY (...))` rather than `a <> ALL (...)`.
- Use `x = ANY (table containing (1,2))` rather than `x=1 OR x=2`.
- Use `= ANY` rather than `EXISTS`.
- For uncorrelated subqueries that always return one row, `IN` is always slower than `=`. For example, use this query:

```
SELECT * FROM t1 WHERE t1.col_name
= (SELECT a FROM t2 WHERE b = some_const);
```

Instead of this query:

```
SELECT * FROM t1 WHERE t1.col_name
IN (SELECT a FROM t2 WHERE b = some_const);
```

These tricks might cause programs to go faster or slower. Using MySQL facilities like the `BENCHMARK()` function, you can get an idea about what helps in your own situation. See [Section 12.9.3, “Information Functions”](#).

Some optimizations that MySQL itself makes are:

- MySQL executes non-correlated subqueries only once. Use `EXPLAIN` to make sure that a given subquery really is non-correlated.

- MySQL rewrites `IN`, `ALL`, `ANY`, and `SOME` subqueries in an attempt to take advantage of the possibility that the select-list columns in the subquery are indexed.
- MySQL replaces subqueries of the following form with an index-lookup function, which `EXPLAIN` describes as a special join type (`unique_subquery` or `index_subquery`):

```
... IN (SELECT indexed_column FROM single_table ...)
```

- MySQL enhances expressions of the following form with an expression involving `MIN()` or `MAX()`, unless `NULL` values or empty sets are involved:

```
value {ALL|ANY|SOME} {> | < | >= | <=} (non-correlated subquery)
```

For example, this `WHERE` clause:

```
WHERE 5 > ALL (SELECT x FROM t)
```

might be treated by the optimizer like this:

```
WHERE 5 > (SELECT MAX(x) FROM t)
```

There is a chapter titled “How MySQL Transforms Subqueries” in the MySQL Internals Manual, available at <http://dev.mysql.com/doc/>.

### 13.2.8.11. Rewriting Subqueries as Joins for Earlier MySQL Versions

Before MySQL 4.1, only nested queries of the form `INSERT ... SELECT ...` and `REPLACE ... SELECT ...` are supported. The `IN()` construct can be used in other contexts to test membership in a set of values.

It is often possible to rewrite a query without a subquery:

```
SELECT * FROM t1 WHERE id IN (SELECT id FROM t2);
```

This can be rewritten as:

```
SELECT DISTINCT t1.* FROM t1, t2 WHERE t1.id=t2.id;
```

The queries:

```
SELECT * FROM t1 WHERE id NOT IN (SELECT id FROM t2);
SELECT * FROM t1 WHERE NOT EXISTS (SELECT id FROM t2 WHERE t1.id=t2.id);
```

Can be be rewritten using `IN()`:

```
SELECT table1.* FROM table1
LEFT JOIN table2 ON table1.id=table2.id
WHERE table2.id IS NULL;
```

A `LEFT [OUTER] JOIN` can be faster than an equivalent subquery because the server might be able to optimize it better — a fact that is not specific to MySQL Server alone. Prior to SQL-92, outer joins did not exist, so subqueries were the only way to do certain things. Today, MySQL Server and many other modern database systems offer a wide range of outer join types.

For more complicated subqueries, you can often create temporary tables to hold the subquery. In some

cases, however, this option does not work. The most frequently encountered of these cases arises with `DELETE` statements, for which standard SQL does not support joins (except in subqueries). For this situation, there are three options available:

- The first option is to upgrade to MySQL 4.1, which does support subqueries in `DELETE` statements.
- The second option is to use a procedural programming language (such as Perl or PHP) to submit a `SELECT` query which obtains the primary keys for the rows to be deleted, and then use these values to construct the appropriate `DELETE` statement (`DELETE FROM ... WHERE key_col IN (key1, key2, ...)`).
- The third option is to use interactive SQL to construct a set of `DELETE` statements automatically, using the MySQL extension `CONCAT()` (in lieu of the standard `||` operator). For example:

```
SELECT
  CONCAT('DELETE FROM tabl WHERE pkid = ', "", tabl.pkid, "", ';')
  FROM tabl, tab2
  WHERE tabl.col1 = tab2.col2;
```

You can place this query in a script file, use the file as input to one instance of the `mysql` program, and use the program output as input to a second instance of `mysql`:

```
shell> mysql --skip-column-names mydb < myscript.sql | mysql mydb
```

MySQL Server 4.0 supports multiple-table `DELETE` statements that can be used to efficiently delete rows based on information from one table or even from many tables at the same time. Multiple-table `UPDATE` statements are also supported as of MySQL 4.0.

## 13.2.9. TRUNCATE Syntax

```
TRUNCATE [TABLE] tbl_name
```

`TRUNCATE TABLE` empties a table completely. Logically, this is equivalent to a `DELETE` statement that deletes all rows, but there are practical differences under some circumstances.

For `InnoDB`, `TRUNCATE TABLE` is mapped to `DELETE`, so there is no difference.

For other storage engines, `TRUNCATE TABLE` differs from `DELETE` in the following ways from MySQL 4.0 onward:

- Truncate operations drop and re-create the table, which is much faster than deleting rows one by one.
- Truncate operations are not transaction-safe; an error occurs when attempting one in the course of an active transaction or active table lock.
- The number of deleted rows is not returned.
- As long as the table format file `tbl_name.frm` is valid, the table can be re-created as an empty table with `TRUNCATE TABLE`, even if the data or index files have become corrupted.
- The table handler does not remember the last used `AUTO_INCREMENT` value, but starts counting from the beginning. This is true even for `MyISAM` and `InnoDB`, which normally do not reuse sequence values.

In MySQL 3.23, `TRUNCATE TABLE` is mapped to `COMMIT; DELETE FROM tbl_name`, so it behaves like `DELETE`. See [Section 13.2.1, “DELETE Syntax”](#).

`TRUNCATE TABLE` is an Oracle SQL extension. This statement was added in MySQL 3.23.28, although from 3.23.28 to 3.23.32, the keyword `TABLE` must be omitted.

## 13.2.10. UPDATE Syntax

Single-table syntax:

```
UPDATE [LOW_PRIORITY] [IGNORE] tbl_name
SET col_name1=expr1 [, col_name2=expr2 ...]
[WHERE where_condition]
[ORDER BY ...]
[LIMIT row_count]
```

Multiple-table syntax:

```
UPDATE [LOW_PRIORITY] [IGNORE] table_references
SET col_name1=expr1 [, col_name2=expr2 ...]
[WHERE where_condition]
```

For the single-table syntax, the `UPDATE` statement updates columns of existing rows in `tbl_name` with new values. The `SET` clause indicates which columns to modify and the values they should be given. The `WHERE` clause, if given, specifies the conditions that identify which rows to update. With no `WHERE` clause, all rows are updated. If the `ORDER BY` clause is specified, the rows are updated in the order that is specified. The `LIMIT` clause places a limit on the number of rows that can be updated.

For the multiple-table syntax, `UPDATE` updates rows in each table named in `table_references` that satisfy the conditions. In this case, `ORDER BY` and `LIMIT` cannot be used.

`where_condition` is an expression that evaluates to true for each row to be updated. It is specified as described in [Section 13.2.7, “SELECT Syntax”](#).

The `UPDATE` statement supports the following modifiers:

- If you use the `LOW_PRIORITY` keyword, execution of the `UPDATE` is delayed until no other clients are reading from the table.
- If you use the `IGNORE` keyword, the update statement does not abort even if errors occur during the update. Rows for which duplicate-key conflicts occur are not updated. Rows for which columns are updated to values that would cause data conversion errors are updated to the closest valid values instead.

If you access a column from `tbl_name` in an expression, `UPDATE` uses the current value of the column. For example, the following statement sets the `age` column to one more than its current value:

```
UPDATE persondata SET age=age+1;
```

Single-table `UPDATE` assignments are generally evaluated from left to right. For multiple-table updates, there is no guarantee that assignments are carried out in any particular order.

If you set a column to the value it currently has, MySQL notices this and does not update it.

If you update a column that has been declared `NOT NULL` by setting to `NULL`, the column is set to the default value appropriate for the data type and the warning count is incremented. The default value is 0 for numeric types, the empty string ( ' ' ) for string types, and the “zero” value for date and time types.



`UPDATE` returns the number of rows that were actually changed. In MySQL 3.22 or later, the `mysql_info()` C API function returns the number of rows that were matched and updated and the number of warnings that occurred during the `UPDATE`.

Starting from MySQL 3.23, you can use `LIMIT row_count` to restrict the scope of the `UPDATE`. A `LIMIT` clause works as follows:

- Before MySQL 4.0.13, `LIMIT` is a rows-affected restriction. The statement stops as soon as it has changed `row_count` rows that satisfy the `WHERE` clause.
- From 4.0.13 on, `LIMIT` is a rows-matched restriction. The statement stops as soon as it has found `row_count` rows that satisfy the `WHERE` clause, whether or not they actually were changed.

If an `UPDATE` statement includes an `ORDER BY` clause, the rows are updated in the order specified by the clause. `ORDER BY` can be used from MySQL 4.0.0. This can be useful in certain situations that might otherwise result in an error. Suppose that a table `t` contains a column `id` that has a unique index. The following statement could fail with a duplicate-key error, depending on the order in which rows are updated:

```
UPDATE t SET id = id + 1;
```

For example, if the table contains 1 and 2 in the `id` column and 1 is updated to 2 before 2 is updated to 3, an error occurs. To avoid this problem, add an `ORDER BY` clause to cause the rows with larger `id` values to be updated before those with smaller values:

```
UPDATE t SET id = id + 1 ORDER BY id DESC;
```

Starting with MySQL 4.0.4, you can also perform `UPDATE` operations covering multiple tables. However, you cannot use `ORDER BY` or `LIMIT` with a multiple-table `UPDATE`. The `table_references` clause lists the tables involved in the join. Its syntax is described in [Section 13.2.7.1, “JOIN Syntax”](#). Here is an example:

```
UPDATE items,month SET items.price=month.price
WHERE items.id=month.id;
```

The preceding example shows an inner join that uses the comma operator, but multiple-table `UPDATE` statements can use any type of join allowed in `SELECT` statements, such as `LEFT JOIN`.

Before MySQL 4.0.18, you need the `UPDATE` privilege for all tables used in a multiple-table `UPDATE`, even if they were not updated. As of MySQL 4.0.18, you need only the `SELECT` privilege for any columns that are read but not modified.

If you use a multiple-table `UPDATE` statement involving `InnoDB` tables for which there are foreign key constraints, the MySQL optimizer might process tables in an order that differs from that of their parent/child relationship. In this case, the statement fails and rolls back. Instead, update a single table and rely on the `ON UPDATE` capabilities that `InnoDB` provides to cause the other tables to be modified accordingly. See [Section 14.2.7.4, “FOREIGN KEY Constraints”](#).

Currently, you cannot update a table and select from the same table in a subquery.

## 13.3. MySQL Utility Statements

### 13.3.1. DESCRIBE Syntax

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

`DESCRIBE` provides information about the columns in a table. It is a shortcut for `SHOW COLUMNS FROM`. (See [Section 13.5.4.3](#), “`SHOW COLUMNS Syntax`”.)

`col_name` can be a column name, or a string containing the SQL ‘%’ and ‘\_’ wildcard characters to obtain output only for the columns with names matching the string. There is no need to enclose the string within quotes unless it contains spaces or other special characters.

```
mysql> DESCRIBE city;
+-----+-----+-----+-----+-----+-----+
| Field      | Type      | Null | Key | Default | Extra      |
+-----+-----+-----+-----+-----+-----+
| Id         | int(11)   |      | PRI | NULL    | auto_increment |
| Name      | char(35)  |      |     |         |              |
| Country   | char(3)   |      |     |         |              |
| District  | char(20)  | YES  | UNI |         |              |
| Population| int(11)   |      | MUL | 0       |              |
+-----+-----+-----+-----+-----+-----+
5 rows in set (0.00 sec)
```

`Field` indicates the column name.

The `Null` field indicates whether `NULL` values can be stored in the column, with `YES` displayed when `NULL` values are allowed.

The `Key` field indicates whether the column is indexed. A value of `PRI` indicates that the column is part of the table's primary key. `UNI` indicates that the column is part of a `UNIQUE` index. The `MUL` value indicates that multiple occurrences of a given value are allowed within the column.

One reason for `MUL` to be displayed on a `UNIQUE` index is that several columns form a composite `UNIQUE` index; although the combination of the columns is unique, each column can still hold multiple occurrences of a given value. Note that in a composite index, only the leftmost column of the index has an entry in the `Key` field.

If the column allows `NULL` values, the `Key` value can be `MUL` even when a `UNIQUE` index is used. The rationale is that multiple rows in a `UNIQUE` index can hold a `NULL` value if the column is not declared `NOT NULL`. (This behavior changes in MySQL 5.0.)

The `Default` field indicates the default value that is assigned to the column.

The `Extra` field contains any additional information that is available about a given column. In the example shown, the `Extra` field indicates that the `Id` column was created with the `AUTO_INCREMENT` keyword.

If the data types are different from what you expect them to be based on a `CREATE TABLE` statement, note that MySQL sometimes changes data types. See [Section 13.1.5.1](#), “`Silent Column Specification Changes`”.

The `DESCRIBE` statement is provided for compatibility with Oracle.

The `SHOW CREATE TABLE` and `SHOW TABLE STATUS` statements also provide information about tables. See [Section 13.5.4](#), “`SHOW Syntax`”.

## 13.3.2. `HELP Syntax`

```
HELP 'search_string'
```

The `HELP` statement returns online information from the MySQL Reference manual. Its proper operation requires that the help tables in the `mysql` database be initialized with help topic information (see [Section 5.2.7](#), “`Server-Side Help`”).

The `HELP` statement searches the help tables for the given search string and displays the result of the search. The search string is not case sensitive.

The `HELP` statement understands several types of search strings:

- At the most general level, use `contents` to retrieve a list of the top-level help categories:

```
HELP 'contents'
```

- For a list of topics in a given help category, such as `Data Types`, use the category name:

```
HELP 'data types'
```

- For help on a specific help topic, such as as the `ASCII ( )` function or the `CREATE TABLE` statement, use the associated keyword or keywords:

```
HELP 'ascii'
HELP 'create table'
```

In other words, the search string matches a category, many topics, or a single topic. You cannot necessarily tell in advance whether a given search string will return a list of items or the help information for a single help topic. However, you can tell what kind of response `HELP` returned by examining the number of rows and columns in the result set.

The following descriptions indicate the forms that the result set can take. Output for the example statements is shown using the familiar “tabular” or “vertical” format that you see when using the `mysql` client, but note that `mysql` itself reformats `HELP` result sets in a different way.

- Empty result set

No match could be found for the search string.

- Result set containing a single row with three columns

This means that the search string yielded a hit for the help topic. The result has three columns:

- `name`: The topic name.
- `description`: Descriptive help text for the topic.
- `example`: Usage example or examples. This column might be blank.

Example: `HELP 'replace'`

Yields:

```
name: REPLACE
description: Syntax:
REPLACE(str,from_str,to_str)

Returns the string str with all occurrences of the string from_str
replaced by the string to_str. REPLACE() performs a case-sensitive
match when searching for from_str.
example: mysql> SELECT REPLACE('www.mysql.com', 'w', 'Ww');
-> 'WwWwWw.mysql.com'
```

- Result set containing multiple rows with two columns

This means that the search string matched many help topics. The result set indicates the help topic names:

- `name`: The help topic name.
- `is_it_category`: `Y` if the name represents a help category, `N` if it does not. If it does not, the `name` value when specified as the argument to the `HELP` statement should yield a single-row result set containing a description for the named item.

Example: `HELP 'status'`

Yields:

name	is_it_category
SHOW	N
SHOW ENGINE	N
SHOW INNODB STATUS	N
SHOW MASTER STATUS	N
SHOW SLAVE STATUS	N
SHOW STATUS	N
SHOW TABLE STATUS	N

- Result set containing multiple rows with three columns

This means the search string matches a category. The result set contains category entries:

- `source_category_name`: The help category name.
- `name`: The category or topic name
- `is_it_category`: `Y` if the name represents a help category, `N` if it does not. If it does not, the `name` value when specified as the argument to the `HELP` statement should yield a single-row result set containing a description for the named item.

Example: `HELP 'functions'`

Yields:

source_category_name	name	is_it_category
Functions	CREATE FUNCTION	N
Functions	DROP FUNCTION	N
Functions	Bit Functions	Y
Functions	Comparison operators	Y
Functions	Control flow functions	Y
Functions	Date and Time Functions	Y
Functions	Encryption Functions	Y
Functions	Information Functions	Y
Functions	Logical operators	Y
Functions	Miscellaneous Functions	Y
Functions	Numeric Functions	Y
Functions	String Functions	Y

The `HELP` statement was added in MySQL 4.1.

### 13.3.3. USE Syntax

USE `db_name`

The `USE db_name` statement tells MySQL to use the `db_name` database as the default (current) database for subsequent statements. The database remains the default until the end of the session or another `USE` statement is issued:

```
USE db1;
SELECT COUNT(*) FROM mytable;    # selects from db1.mytable
USE db2;
SELECT COUNT(*) FROM mytable;    # selects from db2.mytable
```

Making a particular database the default by means of the `USE` statement does not preclude you from accessing tables in other databases. The following example accesses the `author` table from the `db1` database and the `editor` table from the `db2` database:

```
USE db1;
SELECT author_name,editor_name FROM author,db2.editor
WHERE author.editor_id = db2.editor.editor_id;
```

The `USE` statement is provided for compatibility with Sybase.

## 13.4. MySQL Transactional and Locking Statements

### 13.4.1. `START TRANSACTION`, `COMMIT`, and `ROLLBACK` Syntax

```
START TRANSACTION | BEGIN [WORK]
COMMIT
ROLLBACK
SET AUTOCOMMIT = {0 | 1}
```

The `START TRANSACTION` and `BEGIN` statement begin a new transaction. `COMMIT` commits the current transaction, making its changes permanent. `ROLLBACK` rolls back the current transaction, canceling its changes. The `SET AUTOCOMMIT` statement disables or enables the default autocommit mode for the current connection.

By default, MySQL runs with autocommit mode enabled. This means that as soon as you execute a statement that updates (modifies) a table, MySQL stores the update on disk.

If you are using a transaction-safe storage engine (such as `InnoDB`, `BDB`, or `NDB Cluster`), you can disable autocommit mode with the following statement:

```
SET AUTOCOMMIT=0;
```

After disabling autocommit mode by setting the `AUTOCOMMIT` variable to zero, you must use `COMMIT` to store your changes to disk or `ROLLBACK` if you want to ignore the changes you have made since the beginning of your transaction.

To disable autocommit mode for a single series of statements, use the `START TRANSACTION` statement:

```
START TRANSACTION;
SELECT @A:=SUM(salary) FROM table1 WHERE type=1;
UPDATE table2 SET summary=@A WHERE type=1;
COMMIT;
```

With `START TRANSACTION`, autocommit remains disabled until you end the transaction with `COMMIT` or `ROLLBACK`. The autocommit mode then reverts to its previous state.

`BEGIN` and `BEGIN WORK` are supported as aliases of `START TRANSACTION` for initiating a transac-

tion. `START TRANSACTION` was added in MySQL 4.0.11. This is standard SQL syntax and is the recommended way to start an ad-hoc transaction. `BEGIN` and `BEGIN WORK` are available from MySQL 3.23.17 and 3.23.19, respectively.

As of MySQL 4.1.8, you can begin a transaction like this:

```
START TRANSACTION WITH CONSISTENT SNAPSHOT;
```

The `WITH CONSISTENT SNAPSHOT` clause starts a consistent read for storage engines that are capable of it. Currently, this applies only to `InnoDB`. The effect is the same as issuing a `START TRANSACTION` followed by a `SELECT` from any `InnoDB` table. See [Section 14.2.11.4, “Consistent Non-Locking Read”](#).

The `WITH CONSISTENT SNAPSHOT` clause does not change the current transaction isolation level, so it provides a consistent snapshot only if the current isolation level is one that allows consistent read (`REPEATABLE READ` or `SERIALIZABLE`).

Beginning a transaction causes an implicit `UNLOCK TABLES` to be performed.

For best results, transactions should be performed using only tables managed by a single transactional storage engine. Otherwise, the following problems can occur:

- If you use tables from more than one transaction-safe storage engine (such as `InnoDB` and `BDB`), and the transaction isolation level is not `SERIALIZABLE`, it is possible that when one transaction commits, another ongoing transaction that uses the same tables will see only some of the changes made by the first transaction. That is, the atomicity of transactions is not guaranteed with mixed engines and inconsistencies can result. (If mixed-engine transactions are infrequent, you can use `SET TRANSACTION ISOLATION LEVEL` to set the isolation level to `SERIALIZABLE` on a per-transaction basis as necessary.)
- If you use non-transaction-safe tables within a transaction, any changes to those tables are stored at once, regardless of the status of autocommit mode.

If you issue a `ROLLBACK` statement after updating a non-transactional table within a transaction, an `ER_WARNING_NOT_COMPLETE_ROLLBACK` warning occurs. Changes to transaction-safe tables are rolled back, but not changes to non-transaction-safe tables.

If you are using `START TRANSACTION` or `SET AUTOCOMMIT=0`, you should use the MySQL binary log for backups instead of the older update log. Transactions are stored in the binary log in one chunk, upon `COMMIT`. Transactions that are rolled back are not logged. (**Exception:** Modifications to non-transactional tables cannot be rolled back. If a transaction that is rolled back includes modifications to non-transactional tables, the entire transaction is logged with a `ROLLBACK` statement at the end to ensure that the modifications to those tables are replicated. This is true as of MySQL 4.0.15.) See [Section 5.11.4, “The Binary Log”](#).

You can change the isolation level for transactions with `SET TRANSACTION ISOLATION LEVEL`. See [Section 13.4.6, “SET TRANSACTION Syntax”](#).

Rolling back can be a slow operation that may occur without the user having explicitly asked for it (for example, when an error occurs). Because of this, `SHOW PROCESSLIST` displays `Rolling back` in the `State` column for the connection during implicit and explicit (`ROLLBACK` SQL statement) roll-backs, starting from MySQL 4.1.8.

## 13.4.2. Statements That Cannot Be Rolled Back

Some statements cannot be rolled back. In general, these include data definition language (DDL) statements, such as those that create or drop databases or those that create, drop, or alter tables.

You should design your transactions not to include such statements. If you issue a statement early in a transaction that cannot be rolled back, and then another statement later fails, the full effect of the transaction cannot be rolled back in such cases by issuing a `ROLLBACK` statement.

### 13.4.3. Statements That Cause an Implicit Commit

Each of the following statements (and any synonyms for them) implicitly end a transaction, as if you had done a `COMMIT` before executing the statement:

- `ALTER TABLE`, `BEGIN`, `CREATE DATABASE`, `CREATE INDEX`, `CREATE TABLE`, `DROP DATABASE`, `DROP INDEX`, `DROP TABLE`, `LOAD MASTER DATA`, `LOCK TABLES`, `RENAME TABLE`, `SET AUTOCOMMIT=1`, `START TRANSACTION`, `TRUNCATE TABLE`, `UNLOCK TABLES`.
- `UNLOCK TABLES` commits a transaction only if any tables currently are locked.
- Prior to MySQL 4.0.13, `CREATE TABLE` commits a transaction if the binary update log is enabled. The `CREATE TABLE`, `CREATE DATABASE DROP DATABASE`, and `TRUNCATE TABLE` statements cause an implicit commit beginning with MySQL 4.1.13.
- The `CREATE TABLE` statement in `InnoDB` is processed as a single transaction. This means that a `ROLLBACK` from the user does not undo `CREATE TABLE` statements the user made during that transaction.

Transactions cannot be nested. This is a consequence of the implicit `COMMIT` performed for any current transaction when you issue a `START TRANSACTION` statement or one of its synonyms.

### 13.4.4. `SAVEPOINT` and `ROLLBACK TO SAVEPOINT` Syntax

```
SAVEPOINT identifier
ROLLBACK TO SAVEPOINT identifier
```

Starting from MySQL 4.0.14 and 4.1.1, `InnoDB` supports the SQL statements `SAVEPOINT` and `ROLLBACK TO SAVEPOINT`.

The `SAVEPOINT` statement sets a named transaction save point with a name of *identifier*. If the current transaction has a save point with the same name, the old save point is deleted and a new one is set.

The `ROLLBACK TO SAVEPOINT` statement rolls back a transaction to the named save point. Modifications that the current transaction made to rows after the save point was set are undone in the rollback, but `InnoDB` does *not* release the row locks that were stored in memory after the save point. (Note that for a new inserted row, the lock information is carried by the transaction ID stored in the row; the lock is not separately stored in memory. In this case, the row lock is released in the undo.) Savepoints that were set at a later time than the named save point are deleted.

If the `ROLLBACK TO SAVEPOINT` statement returns the following error, it means that no savepoint with the specified name exists:

```
ERROR 1181: Got error 153 during ROLLBACK
```

All save points of the current transaction are deleted if you execute a `COMMIT`, or a `ROLLBACK` that does not name a save point.

### 13.4.5. `LOCK TABLES` and `UNLOCK TABLES` Syntax

```
LOCK TABLES
  tbl_name [AS alias] {READ [LOCAL] | [LOW_PRIORITY] WRITE}
  [, tbl_name [AS alias] {READ [LOCAL] | [LOW_PRIORITY] WRITE}] ...
UNLOCK TABLES
```

`LOCK TABLES` locks tables for the current thread. If any of the tables are locked by other threads, it blocks until all locks can be acquired. `UNLOCK TABLES` releases any locks held by the current thread. All tables that are locked by the current thread are implicitly unlocked when the thread issues another `LOCK TABLES`, or when the connection to the server is closed.

A table lock protects only against inappropriate reads or writes by other clients. The client holding the lock, even a read lock, can perform table-level operations such as `DROP TABLE`. Truncate operations are not transaction-safe, so an error occurs if the client attempts one during an active transaction or while holding a table lock.

Note the following regarding the use of `LOCK TABLES` with transactional tables:

- `LOCK TABLES` is not transaction-safe and implicitly commits any active transactions before attempting to lock the tables. Also, beginning a transaction (for example, with `START TRANSACTION`) implicitly performs an `UNLOCK TABLES`. (See [Section 13.4.3, “Statements That Cause an Implicit Commit”](#).)
- The correct way to use `LOCK TABLES` with transactional tables, such as `InnoDB` tables, is to set `AUTOCOMMIT = 0` and not to call `UNLOCK TABLES` until you commit the transaction explicitly. When you call `LOCK TABLES`, `InnoDB` internally takes its own table lock, and MySQL takes its own table lock. `InnoDB` releases its table lock at the next commit, but for MySQL to release its table lock, you have to call `UNLOCK TABLES`. You should not have `AUTOCOMMIT = 1`, because then `InnoDB` releases its table lock immediately after the call of `LOCK TABLES`, and deadlocks can very easily happen. Starting from 4.1.9, we do not acquire the `InnoDB` table lock at all if `AUTOCOMMIT=1`, to help old applications avoid unnecessary deadlocks.
- `ROLLBACK` does not release MySQL's non-transactional table locks.

As of MySQL 4.0.2, to use `LOCK TABLES` you must have the `LOCK TABLES` privilege and a `SELECT` privilege for the involved tables. In MySQL 3.23, you must have `SELECT`, `INSERT`, `DELETE`, and `UPDATE` privileges for the tables.

The main reasons to use `LOCK TABLES` are to emulate transactions or to get more speed when updating tables. This is explained in more detail later.

If a thread obtains a `READ` lock on a table, that thread (and all other threads) can only read from the table. If a thread obtains a `WRITE` lock on a table, only the thread holding the lock can write to the table. Other threads are blocked from reading or writing the table until the lock has been released.

The difference between `READ LOCAL` and `READ` is that `READ LOCAL` allows non-conflicting `INSERT` statements (concurrent inserts) to execute while the lock is held. However, this cannot be used if you are going to manipulate the database files outside MySQL while you hold the lock. For `InnoDB` tables, `READ LOCAL` is the same as `READ` as of MySQL 4.1.15. (Before that, `READ LOCAL` essentially does nothing: It does not lock the table at all, so for `InnoDB` tables, the use of `READ LOCAL` is deprecated because a plain consistent-read `SELECT` does the same thing, and no locks are needed.)

When you use `LOCK TABLES`, you must lock all tables that you are going to use in your queries. While the locks obtained with a `LOCK TABLES` statement are in effect, you cannot access any tables that were not locked by the statement. Also, you cannot use a locked table multiple times in a single query. Use aliases instead, in which case you must obtain a lock for each alias separately.

```
mysql> LOCK TABLE t WRITE, t AS t1 WRITE;
mysql> INSERT INTO t SELECT * FROM t;
```



```
ERROR 1100: Table 't' was not locked with LOCK TABLES
mysql> INSERT INTO t SELECT * FROM t AS t1;
```

If your queries refer to a table by means of an alias, you must lock the table using that same alias. It does not work to lock the table without specifying the alias:

```
mysql> LOCK TABLE t READ;
mysql> SELECT * FROM t AS myalias;
ERROR 1100: Table 'myalias' was not locked with LOCK TABLES
```

Conversely, if you lock a table using an alias, you must refer to it in your queries using that alias:

```
mysql> LOCK TABLE t AS myalias READ;
mysql> SELECT * FROM t;
ERROR 1100: Table 't' was not locked with LOCK TABLES
mysql> SELECT * FROM t AS myalias;
```

**WRITE** locks normally have higher priority than **READ** locks to ensure that updates are processed as soon as possible. This means that if one thread obtains a **READ** lock and then another thread requests a **WRITE** lock, subsequent **READ** lock requests wait until the **WRITE** thread has gotten the lock and released it. You can use **LOW\_PRIORITY WRITE** locks to allow other threads to obtain **READ** locks while the thread is waiting for the **WRITE** lock. You should use **LOW\_PRIORITY WRITE** locks only if you are sure that eventually there will be a time when no threads have a **READ** lock.

**LOCK TABLES** works as follows:

1. Sort all tables to be locked in an internally defined order. From the user standpoint, this order is undefined.
2. If a table is locked with a read and a write lock, put the write lock before the read lock.
3. Lock one table at a time until the thread gets all locks.

This policy ensures that table locking is deadlock free. There are, however, other things you need to be aware of about this policy: If you are using a **LOW\_PRIORITY WRITE** lock for a table, it means only that MySQL waits for this particular lock until there are no threads that want a **READ** lock. When the thread has gotten the **WRITE** lock and is waiting to get the lock for the next table in the lock table list, all other threads wait for the **WRITE** lock to be released. If this becomes a serious problem with your application, you should consider converting some of your tables to transaction-safe tables.

You can safely use **KILL** to terminate a thread that is waiting for a table lock. See [Section 13.5.5.3, “KILL Syntax”](#).

Note that you should *not* lock any tables that you are using with **INSERT DELAYED** because in that case the **INSERT** is performed by a separate thread.

Normally, you do not need to lock tables, because all single **UPDATE** statements are atomic; no other thread can interfere with any other currently executing SQL statement. However, there are a few cases when locking tables may provide an advantage:

- If you are going to run many operations on a set of **MyISAM** tables, it is much faster to lock the tables you are going to use. Locking **MyISAM** tables speeds up inserting, updating, or deleting on them. The downside is that no thread can update a **READ**-locked table (including the one holding the lock) and no thread can access a **WRITE**-locked table other than the one holding the lock.

The reason some **MyISAM** operations are faster under **LOCK TABLES** is that MySQL does not flush the key cache for the locked tables until **UNLOCK TABLES** is called. Normally, the key cache is

flushed after each SQL statement.

- If you are using a storage engine in MySQL that does not support transactions, you must use `LOCK TABLES` if you want to ensure that no other thread comes between a `SELECT` and an `UPDATE`. The example shown here requires `LOCK TABLES` to execute safely:

```
LOCK TABLES trans READ, customer WRITE;
SELECT SUM(value) FROM trans WHERE customer_id=some_id;
UPDATE customer
  SET total_value=sum_from_previous_statement
  WHERE customer_id=some_id;
UNLOCK TABLES;
```

Without `LOCK TABLES`, it is possible that another thread might insert a new row in the `trans` table between execution of the `SELECT` and `UPDATE` statements.

You can avoid using `LOCK TABLES` in many cases by using relative updates (`UPDATE customer SET value=value+new_value`) or the `LAST_INSERT_ID()` function. See [Section 1.9.5.3, “Transactions and Atomic Operations”](#).

You can also avoid locking tables in some cases by using the user-level advisory lock functions `GET_LOCK()` and `RELEASE_LOCK()`. These locks are saved in a hash table in the server and implemented with `pthread_mutex_lock()` and `pthread_mutex_unlock()` for high speed. See [Section 12.9.4, “Miscellaneous Functions”](#).

See [Section 7.3.1, “Locking Methods”](#), for more information on locking policy.

You can lock all tables in all databases with read locks with the `FLUSH TABLES WITH READ LOCK` statement. See [Section 13.5.5.2, “FLUSH Syntax”](#). This is a very convenient way to get backups if you have a filesystem such as Veritas that can take snapshots in time.

**Note:** If you use `ALTER TABLE` on a locked table, it may become unlocked. See [Section A.7.1, “Problems with ALTER TABLE”](#).

## 13.4.6. SET TRANSACTION Syntax

```
SET [GLOBAL | SESSION] TRANSACTION ISOLATION LEVEL
{ READ UNCOMMITTED | READ COMMITTED | REPEATABLE READ | SERIALIZABLE }
```

This statement sets the transaction isolation level for the next transaction, globally, or for the current session.

The default behavior of `SET TRANSACTION` is to set the isolation level for the next (not yet started) transaction. If you use the `GLOBAL` keyword, the statement sets the default transaction level globally for all new connections created from that point on. Existing connections are unaffected. You need the `SUPER` privilege to do this. Using the `SESSION` keyword sets the default transaction level for all future transactions performed on the current connection.

For descriptions of each InnoDB transaction isolation level, see [Section 14.2.11.3, “InnoDB and TRANSACTION ISOLATION LEVEL”](#). InnoDB supports each of these levels from MySQL 4.0.5 on. The default level is `REPEATABLE READ`.

To set the initial default global isolation level for `mysqld`, use the `--transaction-isolation` option. See [Section 5.2.1, “Command Options”](#).

## 13.5. Database Administration Statements

## 13.5.1. Account Management Statements

MySQL account information is stored in the tables of the `mysql` database. This database and the access control system are discussed extensively in [Chapter 5, Database Administration](#), which you should consult for additional details.

**Important:** Some releases of MySQL introduce changes to the structure of the grant tables to add new privileges or features. Whenever you update to a new version of MySQL, you should update your grant tables to make sure that they have the current structure so that you can take advantage of any new capabilities. See [Section 5.5.1, “mysql\\_fix\\_privilege\\_tables — Upgrade MySQL System Tables”](#).

### 13.5.1.1. DROP USER Syntax

```
DROP USER user [, user] ...
```

The `DROP USER` statement removes one or more MySQL accounts. To use it, you must have the `DELETE` privilege for the `mysql` database. Each account is named using the same format as for the `GRANT` statement; for example, `'jeffrey'@'localhost'`. The user and host parts of the account name correspond to the `User` and `Host` column values of the `user` table row for the account.

`DROP USER` was added in MySQL 4.1.1. In MySQL 4.1, it removes only accounts that have no privileges. This means that the procedure for removing an account depends on your version of MySQL.

From 4.1.1 onward in the MySQL 4.1 release series, `DROP USER` serves only to remove each account record from the `user` table. To remove a MySQL account completely, you should use the following procedure, performing the steps in the order shown:

1. Use `SHOW GRANTS` to determine what privileges the account has. See [Section 13.5.4.10, “SHOW GRANTS Syntax”](#).
2. Use `REVOKE` to revoke the privileges displayed by `SHOW GRANTS`. This removes rows for the account from all the grant tables except the `user` table, and revokes any global privileges listed in the `user` table. See [Section 13.5.1.2, “GRANT Syntax”](#).
3. Delete the account by using `DROP USER` to remove the `user` table record.

**Note:** `DROP USER` does not automatically close any open user sessions. Rather, in the event that a user with an open session is dropped, the statement does not take effect until that user's session is closed. Once the session is closed, the user is dropped, and that user's next attempt to log in will fail. *This is by design.*

Before MySQL 4.1.1, `DROP USER` is not available. You should first revoke the account privileges as just described. Then delete the `user` table record and flush the grant tables as shown here:

```
mysql> DELETE FROM mysql.user
-> WHERE User='user_name' and Host='host_name';
mysql> FLUSH PRIVILEGES;
```

### 13.5.1.2. GRANT Syntax

```
GRANT priv_type [(column_list)] [, priv_type [(column_list))] ...
ON {tbl_name | * | *.* | db_name.*}
TO user [IDENTIFIED BY [PASSWORD] 'password']
    [, user [IDENTIFIED BY [PASSWORD] 'password']] ...
[REQUIRE
    NONE |
    [{SSL | X509}]
    [CIPHER 'cipher' [AND]]
    [ISSUER 'issuer' [AND]]
```

```

    [SUBJECT 'subject']]
    [WITH with_option [with_option] ...]

with_option =
GRANT OPTION
| MAX_QUERIES_PER_HOUR count
| MAX_UPDATES_PER_HOUR count
| MAX_CONNECTIONS_PER_HOUR count

```

The **GRANT** statement enables system administrators to create MySQL user accounts and to grant rights to accounts. To use **GRANT**, you must have the **GRANT OPTION** privilege, and you must have the privileges that you are granting. **GRANT** is implemented in MySQL 3.22.11 or later. For earlier MySQL versions, it does nothing. The **REVOKE** statement is related and enables administrators to remove account privileges. See [Section 13.5.1.3, “REVOKE Syntax”](#).

MySQL account information is stored in the tables of the `mysql` database. This database and the access control system are discussed extensively in [Chapter 5, Database Administration](#), which you should consult for additional details.

**Important:** Some releases of MySQL introduce changes to the structure of the grant tables to add new privileges or features. Whenever you update to a new version of MySQL, you should update your grant tables to make sure that they have the current structure so that you can take advantage of any new capabilities. See [Section 5.5.1, “mysql\\_fix\\_privilege\\_tables — Upgrade MySQL System Tables”](#).

If the grant tables hold privilege rows that contain mixed-case database or table names and the `lower_case_table_names` system variable is set to a non-zero value, **REVOKE** cannot be used to revoke these privileges. It will be necessary to manipulate the grant tables directly. (**GRANT** will not create such rows when `lower_case_table_names` is set, but such rows might have been created prior to setting the variable.)

Privileges can be granted at several levels:

- **Global level**

Global privileges apply to all databases on a given server. These privileges are stored in the `mysql.user` table. **GRANT ALL ON \*.\*** and **REVOKE ALL ON \*.\*** grant and revoke only global privileges.

- **Database level**

Database privileges apply to all tables in a given database. These privileges are stored in the `mysql.db` and `mysql.host` tables. **GRANT ALL ON db\_name.\*** and **REVOKE ALL ON db\_name.\*** grant and revoke only database privileges.

- **Table level**

Table privileges apply to all columns in a given table. These privileges are stored in the `mysql.tables_priv` table. **GRANT ALL ON db\_name.tbl\_name** and **REVOKE ALL ON db\_name.tbl\_name** grant and revoke only table privileges.

- **Column level**

Column privileges apply to single columns in a given table. These privileges are stored in the `mysql.columns_priv` table. When using **REVOKE**, you must specify the same columns that were granted.

For the **GRANT** and **REVOKE** statements, `priv_type` can be specified as any of the following:

Privilege	Meaning
-----------	---------

ALL [ PRIVILEGES ]	Sets all simple privileges except GRANT OPTION
ALTER	Enables use of ALTER TABLE
CREATE	Enables use of CREATE TABLE
CREATE TEMPORARY TABLES	Enables use of CREATE TEMPORARY TABLE
DELETE	Enables use of DELETE
DROP	Enables use of DROP TABLE
EXECUTE	Not implemented
FILE	Enables use of SELECT ... INTO OUTFILE and LOAD DATA INFILE
INDEX	Enables use of CREATE INDEX and DROP INDEX
INSERT	Enables use of INSERT
LOCK TABLES	Enables use of LOCK TABLES on tables for which you have the SELECT privilege
PROCESS	Enables use of SHOW FULL PROCESSLIST
REFERENCES	Not implemented
RELOAD	Enables use of FLUSH
REPLICATION CLIENT	Enables the user to ask where slave or master servers are
REPLICATION SLAVE	Needed for replication slaves (to read binary log events from the master)
SELECT	Enables use of SELECT
SHOW DATABASES	SHOW DATABASES shows all databases
SHUTDOWN	Enables use of mysqladmin shutdown
SUPER	Enables use of CHANGE MASTER, KILL, PURGE MASTER LOGS, and SET GLOBAL statements, the mysqladmin debug command; allows you to connect (once) even if max_connections is reached
UPDATE	Enables use of UPDATE
USAGE	Synonym for “no privileges”
GRANT OPTION	Enables privileges to be granted

The CREATE TEMPORARY TABLES, EXECUTE, LOCK TABLES, REPLICATION CLIENT, REPLICATION SLAVE, SHOW DATABASES, and SUPER privileges were added in MySQL 4.0.2. To use these privileges when upgrading from an earlier version of MySQL that does not have them, you must first upgrade the grant tables. See [Section 5.5.1, “mysql\\_fix\\_privilege\\_tables — Upgrade MySQL System Tables”](#).

The REFERENCES and EXECUTE privileges are unused in MySQL versions up to and including the 4.1 release series.

In older MySQL versions that do not have the SUPER privilege, specify the PROCESS privilege instead.

USAGE can be specified when you want to create a user that has no privileges.

Use SHOW GRANTS to determine what privileges an account has. See [Section 13.5.4.10, “SHOW GRANTS Syntax”](#).

You can assign global privileges by using ON \*.\* syntax or database-level privileges by using ON db\_name.\* syntax. If you specify ON \* and you have selected a default database, the privileges are

granted in that database. (**Warning:** If you specify `ON *` and you have *not* selected a default database, the privileges granted are global.)

The `FILE`, `PROCESS`, `RELOAD`, `REPLICATION CLIENT`, `REPLICATION SLAVE`, `SHOW DATABASES`, `SHUTDOWN`, and `SUPER` privileges are administrative privileges that can only be granted globally (using `ON *.*` syntax).

Other privileges can be granted globally or at more specific levels.

The *priv\_type* values that you can specify for a table are `SELECT`, `INSERT`, `UPDATE`, `DELETE`, `CREATE`, `DROP`, `GRANT OPTION`, `INDEX`, and `ALTER`.

The *priv\_type* values that you can specify for a column (that is, when you use a *column\_list* clause) are `SELECT`, `INSERT`, and `UPDATE`.

For the global, database, and table levels, `GRANT ALL` assigns only the privileges that exist at the level you are granting. For example, if you use `GRANT ALL ON db_name.*`, that is a database-level statement, so none of the global-only privileges such as `FILE` are granted.

MySQL allows you to grant privileges even on databases and tables that do not exist. In such cases, the privileges to be granted must include the `CREATE` privilege. *This behavior is by design*, and is intended to enable the database administrator to prepare user accounts and privileges for databases and tables that are to be created at a later time.

**Important:** *MySQL does not automatically revoke any privileges when you drop a table or database.*

**Note:** the `'_'` and `'%'` wildcards are allowed when specifying database names in `GRANT` statements that grant privileges at the global or database levels. This means, for example, that if you want to use a `'_'` character as part of a database name, you should specify it as `'\_'` in the `GRANT` statement, to prevent the user from being able to access additional databases matching the wildcard pattern; for example, `GRANT ... ON `foo\_bar`.* TO ...`

To accommodate granting rights to users from arbitrary hosts, MySQL supports specifying the *user* value in the form `user_name@host_name`. If a *user\_name* or *host\_name* value is legal as an unquoted identifier, you need not quote it. However, quotes are necessary to specify a *user\_name* string containing special characters (such as `'-'`), or a *host\_name* string containing special characters or wildcard characters (such as `'%'`); for example, `'test-user'@'test-hostname'`. Quote the username and hostname separately.

You can specify wildcards in the hostname. For example, `user_name@'%.loc.gov'` applies to *user\_name* for any host in the `loc.gov` domain, and `user_name@'144.155.166.%'` applies to *user\_name* for any host in the `144.155.166` class C subnet.

The simple form *user\_name* is a synonym for `user_name@'%'`.

*MySQL does not support wildcards in usernames.* Anonymous users are defined by inserting entries with `User=''` into the `mysql.user` table or by creating a user with an empty name with the `GRANT` statement:

```
GRANT ALL ON test.* TO ''@'localhost' ...
```

When specifying quoted values, quote database, table, and column names as identifiers, using backticks (```). Quote hostnames, usernames, and passwords as strings, using single quotes (`'`).

**Warning:** If you allow anonymous users to connect to the MySQL server, you should also grant privileges to all local users as `user_name@localhost`. Otherwise, the anonymous user account for `localhost` in the `mysql.user` table (created during MySQL installation) is used when named users try to log in to the MySQL server from the local machine. For details, see [Section 5.7.5, “Access Control, Stage 1: Connection Verification”](#).

You can determine whether this applies to you by executing the following query, which lists any anonymous users:

```
SELECT Host, User FROM mysql.user WHERE User='';
```

If you want to delete the local anonymous user account to avoid the problem just described, use these statements:

```
DELETE FROM mysql.user WHERE Host='localhost' AND User='';
FLUSH PRIVILEGES;
```

**GRANT** supports hostnames up to 60 characters long. Database, table, and column names can be up to 64 characters. Usernames can be up to 16 characters. **Note:** *The allowable length for usernames cannot be changed by altering the `mysql.user` table, and attempting to do so results in unpredictable behavior which may even make it impossible for users to log in to the MySQL server.* You should never alter any of the tables in the `mysql` database in any manner whatsoever except by means of the procedure prescribed by MySQL AB that is described in [Section 5.5.1, “mysql\\_fix\\_privilege\\_tables — Upgrade MySQL System Tables”](#).

The privileges for a table or column are formed additively as the logical **OR** of the privileges at each of the privilege levels. For example, if the `mysql.user` table specifies that a user has a global **SELECT** privilege, the privilege cannot be denied by an entry at the database, table, or column level.

The privileges for a column can be calculated as follows:

```
global privileges
OR (database privileges AND host privileges)
OR table privileges
OR column privileges
```

In most cases, you grant rights to a user at only one of the privilege levels, so life is not normally this complicated. The details of the privilege-checking procedure are presented in [Section 5.7, “The MySQL Access Privilege System”](#).

If you grant privileges for a username/hostname combination that does not exist in the `mysql.user` table, an entry is added and remains there until deleted with a **DELETE** statement. In other words, **GRANT** may create `user` table entries, but **REVOKE** does not remove them; you must do that explicitly using **DROP USER** or **DELETE**.

**Warning:** If you create a new user but do not specify an **IDENTIFIED BY** clause, the user has no password. This is very insecure. As of MySQL 5.0.2, you can enable the `NO_AUTO_CREATE_USER` SQL mode to keep **GRANT** from creating a new user if it would otherwise do so, unless **IDENTIFIED BY** is given to provide the new user a non-empty password.

In MySQL 3.22.12 or later, if a new user is created or if you have global grant privileges, the user's password is set to the password specified by the **IDENTIFIED BY** clause, if one is given. If the user already had a password, this is replaced by the new one.

**Warning:** If you create a new user but do not specify an **IDENTIFIED BY** clause, the user has no password. This is very insecure.

Passwords can also be set with the **SET PASSWORD** statement. See [Section 13.5.1.4, “SET PASSWORD Syntax”](#).

In the **IDENTIFIED BY** clause, the password should be given as the literal password value. It is unnecessary to use the `PASSWORD()` function as it is for the **SET PASSWORD** statement. For example:

```
GRANT ... IDENTIFIED BY 'mypass';
```

If you do not want to send the password in clear text and you know the hashed value that `PASSWORD()` would return for the password, you can specify the hashed value preceded by the keyword `PASSWORD`:

```
GRANT ...
IDENTIFIED BY PASSWORD '*6C8989366EAF75BB670AD8EA7A7FC1176A95CEF4';
```

In a C program, you can get the hashed value by using the `make_scrambled_password()` C API function.

If you grant privileges for a database, an entry in the `mysql.db` table is created if needed. If all privileges for the database are removed with `REVOKE`, this entry is deleted.

The `SHOW DATABASES` privilege enables the account to see database names by issuing the `SHOW DATABASE` statement. Accounts that do not have this privilege see only databases for which they have some privileges, and cannot use the statement at all if the server was started with the `-skip-show-database` option.

If a user has no privileges for a table, the table name is not displayed when the user requests a list of tables (for example, with a `SHOW TABLES` statement).

The `WITH GRANT OPTION` clause gives the user the ability to give to other users any privileges the user has at the specified privilege level. You should be careful to whom you give the `GRANT OPTION` privilege, because two users with different privileges may be able to join privileges!

You cannot grant another user a privilege which you yourself do not have; the `GRANT OPTION` privilege enables you to assign only those privileges which you yourself possess.

Be aware that when you grant a user the `GRANT OPTION` privilege at a particular privilege level, any privileges the user possesses (or may be given in the future) at that level can also be granted by that user to other users. Suppose that you grant a user the `INSERT` privilege on a database. If you then grant the `SELECT` privilege on the database and specify `WITH GRANT OPTION`, that user can give to other users not only the `SELECT` privilege, but also `INSERT`. If you then grant the `UPDATE` privilege to the user on the database, the user can grant `INSERT`, `SELECT`, and `UPDATE`.

For a non-administrative user, you should not grant the `ALTER` privilege globally or for the `mysql` database. If you do that, the user can try to subvert the privilege system by renaming tables!

The `MAX_QUERIES_PER_HOUR count`, `MAX_UPDATES_PER_HOUR count`, and `MAX_CONNECTIONS_PER_HOUR count` options are implemented in MySQL 4.0.2. They limit the number of queries, updates, and logins a user can perform during one hour. If `count` is 0 (the default), this means that there is no limitation for that user.

Note: To specify any of these resource-limit options for an existing user without affecting existing privileges, use `GRANT USAGE ON *.* ... WITH MAX_...`

See [Section 5.8.4, “Limiting Account Resources”](#).

MySQL can check X509 certificate attributes in addition to the usual authentication that is based on the username and password. To specify SSL-related options for a MySQL account, use the `REQUIRE` clause of the `GRANT` statement. (For background information on the use of SSL with MySQL, see [Section 5.8.7, “Using Secure Connections”](#).)

There are a number of different possibilities for limiting connection types for a given account:

- If the account has no SSL or X509 requirements, unencrypted connections are allowed if the username and password are valid. However, encrypted connections can also be used, at the client's option, if the client has the proper certificate and key files.



- The `REQUIRE SSL` option tells the server to allow only SSL-encrypted connections for the account. Note that this option can be omitted if there are any access-control rows that allow non-SSL connections.

```
GRANT ALL PRIVILEGES ON test.* TO 'root'@'localhost'
  IDENTIFIED BY 'goodsecret' REQUIRE SSL;
```

- `REQUIRE X509` means that the client must have a valid certificate but that the exact certificate, issuer, and subject do not matter. The only requirement is that it should be possible to verify its signature with one of the CA certificates.

```
GRANT ALL PRIVILEGES ON test.* TO 'root'@'localhost'
  IDENTIFIED BY 'goodsecret' REQUIRE X509;
```

- `REQUIRE ISSUER 'issuer'` places the restriction on connection attempts that the client must present a valid X509 certificate issued by CA `'issuer'`. If the client presents a certificate that is valid but has a different issuer, the server rejects the connection. Use of X509 certificates always implies encryption, so the `SSL` option is unnecessary in this case.

```
GRANT ALL PRIVILEGES ON test.* TO 'root'@'localhost'
  IDENTIFIED BY 'goodsecret'
  REQUIRE ISSUER '/C=FI/ST=Some-State/L=Helsinki/
  O=MySQL Finland AB/CN=Tonu Samuel/Email=tonu@example.com';
```

Note that the `'issuer'` value should be entered as a single string.

- `REQUIRE SUBJECT 'subject'` places the restriction on connection attempts that the client must present a valid X509 certificate containing the subject `subject`. If the client presents a certificate that is valid but has a different subject, the server rejects the connection.

```
GRANT ALL PRIVILEGES ON test.* TO 'root'@'localhost'
  IDENTIFIED BY 'goodsecret'
  REQUIRE SUBJECT '/C=EE/ST=Some-State/L=Tallinn/
  O=MySQL demo client certificate/
  CN=Tonu Samuel/Email=tonu@example.com';
```

Note that the `'subject'` value should be entered as a single string.

- `REQUIRE CIPHER 'cipher'` is needed to ensure that ciphers and key lengths of sufficient strength are used. SSL itself can be weak if old algorithms using short encryption keys are used. Using this option, you can ask that a specific cipher method is used to allow a connection.

```
GRANT ALL PRIVILEGES ON test.* TO 'root'@'localhost'
  IDENTIFIED BY 'goodsecret'
  REQUIRE CIPHER 'EDH-RSA-DES-CBC3-SHA';
```

The `SUBJECT`, `ISSUER`, and `CIPHER` options can be combined in the `REQUIRE` clause like this:

```
GRANT ALL PRIVILEGES ON test.* TO 'root'@'localhost'
  IDENTIFIED BY 'goodsecret'
  REQUIRE SUBJECT '/C=EE/ST=Some-State/L=Tallinn/
  O=MySQL demo client certificate/
  CN=Tonu Samuel/Email=tonu@example.com'
  AND ISSUER '/C=FI/ST=Some-State/L=Helsinki/
  O=MySQL Finland AB/CN=Tonu Samuel/Email=tonu@example.com'
  AND CIPHER 'EDH-RSA-DES-CBC3-SHA';
```

Starting from MySQL 4.0.4, the `AND` keyword is optional between `REQUIRE` options.

The order of the options does not matter, but no option can be specified twice.

When `mysqld` starts, all privileges are read into memory. For details, see [Section 5.7.7, “When Privilege Changes Take Effect”](#).

Note that if you are using table or column privileges for even one user, the server examines table and column privileges for all users and this slows down MySQL a bit. Similarly, if you limit the number of queries, updates, or connections for any users, the server must monitor these values.

The biggest differences between the standard SQL and MySQL versions of `GRANT` are:

- In MySQL, privileges are associated with the combination of a hostname and username and not with only a username.
- Standard SQL does not have global or database-level privileges, nor does it support all the privilege types that MySQL supports.
- MySQL does not support the standard SQL `TRIGGER` or `UNDER` privileges.
- Standard SQL privileges are structured in a hierarchical manner, which means that if you remove a user, all privileges the user has been granted are revoked. This is not the case in MySQL 4.1 and earlier versions, where the granted privileges are not automatically revoked and you must revoke them explicitly. See [Section 13.5.1.1, “DROP USER Syntax”](#).
- In standard SQL, when you drop a table, all privileges for the table are revoked. In standard SQL, when you revoke a privilege, all privileges that were granted based on that privilege are also revoked. In MySQL, privileges can be dropped only with explicit `REVOKE` statements or by manipulating values stored in the MySQL grant tables.
- In MySQL, it is possible to have the `INSERT` privilege for only some of the columns in a table. In this case, you can still execute `INSERT` statements on the table, provided that you omit those columns for which you do not have the `INSERT` privilege. The omitted columns are set to their implicit default values. (Standard SQL requires you to have the `INSERT` privilege on all columns.) [Section 11.1.4, “Data Type Default Values”](#), discusses implicit default values.

### 13.5.1.3. REVOKE Syntax

```
REVOKE priv_type [(column_list)] [, priv_type [(column_list)]] ...
ON {tbl_name | * | *.* | db_name.*}
FROM user [, user] ...

REVOKE ALL PRIVILEGES, GRANT OPTION FROM user [, user] ...
```

The `REVOKE` statement enables system administrators to revoke privileges from MySQL accounts. To use `REVOKE`, you must have the `GRANT OPTION` privilege, and you must have the privileges that you are revoking. `REVOKE` is implemented in MySQL 3.22.11 or later. For earlier MySQL versions, it does nothing.

For details on the levels at which privileges exist, the allowable *priv\_type* values, and the syntax for specifying users and passwords, see [Section 13.5.1.2, “GRANT Syntax”](#)

If the grant tables hold privilege rows that contain mixed-case database or table names and the `lower_case_table_names` system variable is set to a non-zero value, `REVOKE` cannot be used to revoke these privileges. It will be necessary to manipulate the grant tables directly. (`GRANT` will not create such rows when `lower_case_table_names` is set, but such rows might have been created prior to setting the variable.)

To make it easy to revoke all privileges, MySQL 4.1.2 has added the following syntax, which drops all global, database-, table-, and column-level privileges for the named users:

```
REVOKE ALL PRIVILEGES, GRANT OPTION FROM user [, user] ...
```

To use this `REVOKE` syntax, you must have the `UPDATE` privilege for the `mysql` database.

Before MySQL 4.1.2, all privileges cannot be dropped at once. Two statements are necessary:

```
REVOKE ALL PRIVILEGES ON *.* FROM user [, user] ...
REVOKE GRANT OPTION ON *.* FROM user [, user] ...
```

### 13.5.1.4. SET PASSWORD Syntax

```
SET PASSWORD [FOR user] = PASSWORD('some password')
```

The `SET PASSWORD` statement assigns a password to an existing MySQL user account.

With no `FOR` clause, this statement sets the password for the current user. Any client that has connected to the server using a non-anonymous account can change the password for that account.

With a `FOR` clause, this statement sets the password for a specific account on the current server host. Only clients that have the `UPDATE` privilege for the `mysql` database can do this. The `user` value should be given in `user_name@host_name` format, where `user_name` and `host_name` are exactly as they are listed in the `User` and `Host` columns of the `mysql.user` table entry. For example, if you had an entry with `User` and `Host` column values of `'bob'` and `'%.loc.gov'`, you would write the statement like this:

```
SET PASSWORD FOR 'bob'@'%.loc.gov' = PASSWORD('newpass');
```

That is equivalent to the following statements:

```
UPDATE mysql.user SET Password=PASSWORD('newpass')
  WHERE User='bob' AND Host='%.loc.gov';
FLUSH PRIVILEGES;
```

**Note:** If you are connecting to a MySQL 4.1 or later server using a pre-4.1 client program, do not use the preceding `SET PASSWORD` or `UPDATE` statement without reading [Section 5.7.9, “Password Hashing as of MySQL 4.1”](#), first. The password format changed in MySQL 4.1, and under certain circumstances it is possible that if you change your password, you might not be able to connect to the server afterward.

Starting from MySQL 4.1, you can see which account the server authenticated you as by executing `SELECT CURRENT_USER()`.

## 13.5.2. Table Maintenance Statements

### 13.5.2.1. ANALYZE TABLE Syntax

```
ANALYZE [LOCAL | NO_WRITE_TO_BINLOG] TABLE tbl_name [, tbl_name] ...
```

`ANALYZE TABLE` analyzes and stores the key distribution for a table. During the analysis, the table is locked with a read lock for `MyISAM` and `BDB`. For `InnoDB` the table is locked with a write lock. This statement works with `MyISAM`, `BDB`, and (as of MySQL 4.0.13) `InnoDB` tables. For `MyISAM` tables, this statement is equivalent to using `myisamchk --analyze`.

For more information on how the analysis works within `InnoDB`, see [Section 14.2.17, “Restrictions on InnoDB Tables”](#).

MySQL uses the stored key distribution to decide the order in which tables should be joined when you perform a join on something other than a constant.

This statement requires `SELECT` and `INSERT` privileges for the table.

`ANALYZE TABLE` returns a result set with the following columns:

Column	Value
Table	The table name
Op	Always <code>analyze</code>
Msg_type	One of <code>status</code> , <code>error</code> , <code>info</code> , or <code>warning</code>
Msg_text	The message

You can check the stored key distribution with the `SHOW INDEX` statement. See [Section 13.5.4.11, “SHOW INDEX Syntax”](#).

If the table has not changed since the last `ANALYZE TABLE` statement, the table is not analyzed again.

Before MySQL 4.1.1, `ANALYZE TABLE` statements are not written to the binary log. As of MySQL 4.1.1, they are written to the binary log unless the optional `NO_WRITE_TO_BINLOG` keyword (or its alias `LOCAL`) is used. This is done so that `ANALYZE TABLE` statements used on a MySQL server acting as a replication master will be replicated by default to the replication slave.

### 13.5.2.2. BACKUP TABLE Syntax

```
BACKUP TABLE tbl_name [, tbl_name] ... TO '/path/to/backup/directory'
```

**Note:** This statement is deprecated. We are working on a better replacement for it that will provide on-line backup capabilities. In the meantime, the `mysqlhotcopy` script can be used instead.

`BACKUP TABLE` copies to the backup directory the minimum number of table files needed to restore the table, after flushing any buffered changes to disk. The statement works only for `MyISAM` tables. It copies the `.frm` definition and `.MYD` data files. The `.MYI` index file can be rebuilt from those two files. The directory should be specified as a full pathname. To restore the table, use `RESTORE TABLE`.

During the backup, a read lock is held for each table, one at a time, as they are being backed up. If you want to back up several tables as a snapshot (preventing any of them from being changed during the backup operation), issue a `LOCK TABLES` statement first, to obtain a read lock for all tables in the group.

`BACKUP TABLE` returns a result set with the following columns:

Column	Value
Table	The table name
Op	Always <code>backup</code>
Msg_type	One of <code>status</code> , <code>error</code> , <code>info</code> , or <code>warning</code>
Msg_text	The message

`BACKUP TABLE` is available in MySQL 3.23.25 and later.

### 13.5.2.3. CHECK TABLE Syntax

```
CHECK TABLE tbl_name [, tbl_name] ... [option] ...
option = {QUICK | FAST | MEDIUM | EXTENDED | CHANGED}
```

**CHECK TABLE** checks a table or tables for errors. **CHECK TABLE** works for **MyISAM** and **InnoDB** tables. For **MyISAM** tables, the key statistics are updated as well.

**CHECK TABLE** returns a result set with the following columns:

Column	Value
Table	The table name
Op	Always <code>check</code>
Msg_type	One of <code>status</code> , <code>error</code> , <code>info</code> , or <code>warning</code>
Msg_text	The message

Note that the statement might produce many rows of information for each checked table. The last row has a `Msg_type` value of `status` and the `Msg_text` normally should be `OK`. If you don't get `OK`, or `Table is already up to date` you should normally run a repair of the table. See [Section 5.9.4, “Table Maintenance and Crash Recovery”](#). `Table is already up to date` means that the storage engine for the table indicated that there was no need to check the table.

The different check options that can be given are shown in the following table. These options apply only to checking **MyISAM** tables and are ignored for **InnoDB** tables.

Type	Meaning
<code>QUICK</code>	Do not scan the rows to check for incorrect links.
<code>FAST</code>	Check only tables that have not been closed properly.
<code>CHANGED</code>	Check only tables that have been changed since the last check or that have not been closed properly.
<code>MEDIUM</code>	Scan rows to verify that deleted links are valid. This also calculates a key checksum for the rows and verifies this with a calculated checksum for the keys.
<code>EXTENDED</code>	Do a full key lookup for all keys for each row. This ensures that the table is 100% consistent, but takes a long time.

If none of the options `QUICK`, `MEDIUM`, or `EXTENDED` are specified, the default check type for dynamic-format **MyISAM** tables is `MEDIUM`. This has the same result as running `myisamchk -medium-check tbl_name` on the table. The default check type also is `MEDIUM` for static-format **MyISAM** tables, unless `CHANGED` or `FAST` is specified. In that case, the default is `QUICK`. The row scan is skipped for `CHANGED` and `FAST` because the rows are very seldom corrupted.

You can combine check options, as in the following example that does a quick check on the table to determine whether it was closed properly:

```
CHECK TABLE test_table FAST QUICK;
```

**Note:** In some cases, **CHECK TABLE** changes the table. This happens if the table is marked as “corrupted” or “not closed properly” but **CHECK TABLE** does not find any problems in the table. In this case, **CHECK TABLE** marks the table as okay.

If a table is corrupted, it is most likely that the problem is in the indexes and not in the data part. All of the preceding check types check the indexes thoroughly and should thus find most errors.

If you just want to check a table that you assume is okay, you should use no check options or the `QUICK` option. The latter should be used when you are in a hurry and can take the very small risk that `QUICK` does not find an error in the data file. (In most cases, under normal usage, MySQL should find any error in the data file. If this happens, the table is marked as “corrupted” and cannot be used until it is repaired.)

`FAST` and `CHANGED` are mostly intended to be used from a script (for example, to be executed from `cron`) if you want to check tables from time to time. In most cases, `FAST` is to be preferred over `CHANGED`. (The only case when it is not preferred is when you suspect that you have found a bug in the `MyISAM` code.)

`EXTENDED` is to be used only after you have run a normal check but still get strange errors from a table when MySQL tries to update a row or find a row by key. This is very unlikely if a normal check has succeeded.

Some problems reported by `CHECK TABLE` cannot be corrected automatically:

- Found row where the `auto_increment` column has the value 0.

This means that you have a row in the table where the `AUTO_INCREMENT` index column contains the value 0. (It is possible to create a row where the `AUTO_INCREMENT` column is 0 by explicitly setting the column to 0 with an `UPDATE` statement.)

This is not an error in itself, but could cause trouble if you decide to dump the table and restore it or do an `ALTER TABLE` on the table. In this case, the `AUTO_INCREMENT` column changes value according to the rules of `AUTO_INCREMENT` columns, which could cause problems such as a duplicate-key error.

To get rid of the warning, simply execute an `UPDATE` statement to set the column to some value other than 0.

### 13.5.2.4. CHECKSUM TABLE Syntax

```
CHECKSUM TABLE tbl_name [, tbl_name] ... [ QUICK | EXTENDED ]
```

`CHECKSUM TABLE` reports a table checksum.

With `QUICK`, the live table checksum is reported if it is available, or `NULL` otherwise. This is very fast. A live checksum is enabled by specifying the `CHECKSUM=1` table option when you create the table; currently, this is supported only for `MyISAM` tables. See [Section 13.1.5, “CREATE TABLE Syntax”](#).

With `EXTENDED`, the entire table is read row by row and the checksum is calculated. This can be very slow for large tables.

If neither `QUICK` nor `EXTENDED` is specified, MySQL returns a live checksum if the table storage engine supports it and scans the table otherwise.

For a non-existent table, `CHECKSUM TABLE` returns `NULL`.

This statement is implemented in MySQL 4.1.1.

### 13.5.2.5. OPTIMIZE TABLE Syntax

```
OPTIMIZE [LOCAL | NO_WRITE_TO_BINLOG] TABLE tbl_name [, tbl_name] ...
```

`OPTIMIZE TABLE` should be used if you have deleted a large part of a table or if you have made many changes to a table with variable-length rows (tables that have `VARCHAR`, `VARBINARY`, `BLOB`, or `TEXT` columns). Deleted rows are maintained in a linked list and subsequent `INSERT` operations reuse old row positions. You can use `OPTIMIZE TABLE` to reclaim the unused space and to defragment the data file.

This statement requires `SELECT` and `INSERT` privileges for the table.

In most setups, you need not run `OPTIMIZE TABLE` at all. Even if you do a lot of updates to variable-length rows, it is not likely that you need to do this more than once a week or month and only on certain tables.

`OPTIMIZE TABLE` works only for `MyISAM`, `BDB`, and `InnoDB` tables.

For `MyISAM` tables, `OPTIMIZE TABLE` works as follows:

1. If the table has deleted or split rows, repair the table.
2. If the index pages are not sorted, sort them.
3. If the table's statistics are not up to date (and the repair could not be accomplished by sorting the index), update them.

For `BDB` tables, `OPTIMIZE TABLE` currently is mapped to `ANALYZE TABLE`. See [Section 13.5.2.1, “ANALYZE TABLE Syntax”](#).

That was also the case for `InnoDB` tables before MySQL 4.1.3. As of 4.1.3, `OPTIMIZE TABLE` is mapped to `ALTER TABLE`, which rebuilds the table to update index statistics and free unused space in the clustered index.

You can make `OPTIMIZE TABLE` work on other storage engines by starting `mysqld` with the `-skip-new` or `--safe-mode` option. In this case, `OPTIMIZE TABLE` is just mapped to `ALTER TABLE`.

`OPTIMIZE TABLE` returns a result set with the following columns:

Column	Value
<code>Table</code>	The table name
<code>Op</code>	Always <code>optimize</code>
<code>Msg_type</code>	One of <code>status</code> , <code>error</code> , <code>info</code> , or <code>warning</code>
<code>Msg_text</code>	The message

Note that MySQL locks the table during the time `OPTIMIZE TABLE` is running.

Before MySQL 4.1.1, `OPTIMIZE TABLE` statements are not written to the binary log. As of MySQL 4.1.1, they are written to the binary log unless the optional `NO_WRITE_TO_BINLOG` keyword (or its alias `LOCAL`) is used. This is done so that `OPTIMIZE TABLE` statements used on a MySQL server acting as a replication master will be replicated by default to the replication slave.

### 13.5.2.6. REPAIR TABLE Syntax

```
REPAIR [LOCAL | NO_WRITE_TO_BINLOG] TABLE
      tbl_name [, tbl_name] ... [QUICK] [EXTENDED] [USE_FRM]
```

`REPAIR TABLE` repairs a possibly corrupted table. By default, it has the same effect as `myisamchk`

`--recover tbl_name`. `REPAIR TABLE` works for `MyISAM` and for `ARCHIVE` tables. See [Section 14.1](#), “The `MyISAM` Storage Engine”, and [Section 14.7](#), “The `ARCHIVE` Storage Engine”.

This statement requires `SELECT` and `INSERT` privileges for the table.

Normally, you should never have to run this statement. However, if disaster strikes, `REPAIR TABLE` is very likely to get back all your data from a `MyISAM` table. If tables become corrupted often, you should try to find the reason for it and so to eliminate the need to use `REPAIR TABLE`. See [Section A.4.2](#), “What to Do If MySQL Keeps Crashing”, and [Section 14.1.4](#), “`MyISAM` Table Problems”.

**Warning:** If the server dies during a `REPAIR TABLE` operation, it is essential after restarting it that you immediately execute another `REPAIR TABLE` statement for the table before performing any other operations on it. (It is always a good idea to start by making a backup.) In the worst case, you might have a new clean index file without information about the data file, and then the next operation you perform could overwrite the data file. This is an unlikely but possible scenario.

`REPAIR TABLE` returns a result set with the following columns:

Column	Value
Table	The table name
Op	Always <code>repair</code>
Msg_type	One of <code>status</code> , <code>error</code> , <code>info</code> , or <code>warning</code>
Msg_text	The message

The `REPAIR TABLE` statement might produce many rows of information for each repaired table. The last row has a `Msg_type` value of `status` and `Msg_text` normally should be `OK`. If you do not get `OK`, you should try repairing the table with `myisamchk --safe-recover`. (`REPAIR TABLE` does not yet implement all the options of `myisamchk`.) With `myisamchk --safe-recover`, you can also use options that `REPAIR TABLE` does not support, such as `--max-record-length`.

If `QUICK` is given, `REPAIR TABLE` tries to repair only the index tree. This type of repair is like that done by `myisamchk --recover --quick`.

If you use `EXTENDED`, MySQL creates the index row by row instead of creating one index at a time with sorting. (Before MySQL 4.1, this might be better than sorting on fixed-length keys if you have long `CHAR` keys that compress very well.) This type of repair is like that done by `myisamchk --safe-recover`.

As of MySQL 4.0.2, there is a `USE_FRM` mode for `REPAIR TABLE`. Use this if the `.MYI` index file is missing or if its header is corrupted. In this mode, MySQL re-creates the `.MYI` file using information from the `.frm` file. This kind of repair cannot be done with `myisamchk`. **Note:** Use this mode *only* if you cannot use regular `REPAIR` modes. The `.MYI` header contains important table metadata (in particular, current `AUTO_INCREMENT` value and `Delete link`) that are lost in `REPAIR ... USE_FRM`. Don't use `USE_FRM` if the table is compressed because this information is also stored in the `.MYI` file.

Before MySQL 4.1.1, `REPAIR TABLE` statements are not written to the binary log. As of MySQL 4.1.1, they are written to the binary log unless the optional `NO_WRITE_TO_BINLOG` keyword (or its alias `LOCAL`) is used. This is done so that `REPAIR TABLE` statements used on a MySQL server acting as a replication master will be replicated by default to the replication slave.

### 13.5.2.7. `RESTORE TABLE` Syntax

```
RESTORE TABLE tbl_name [, tbl_name] ... FROM '/path/to/backup/directory'
```

`RESTORE TABLE` restores the table or tables from a backup that was made with `BACKUP TABLE`. The directory should be specified as a full pathname.



Existing tables are not overwritten; if you try to restore over an existing table, an error occurs. Just as for `BACKUP TABLE`, `RESTORE TABLE` currently works only for `MyISAM` tables. Restored tables are not replicated from master to slave.

The backup for each table consists of its `.frm` format file and `.MYD` data file. The restore operation restores those files, and then uses them to rebuild the `.MYI` index file. Restoring takes longer than backing up due to the need to rebuild the indexes. The more indexes the table has, the longer it takes.

`RESTORE TABLE` returns a result set with the following columns:

Column	Value
Table	The table name
Op	Always <code>restore</code>
Msg_type	One of <code>status</code> , <code>error</code> , <code>info</code> , or <code>warning</code>
Msg_text	The message

### 13.5.3. SET Syntax

```
SET variable_assignment [, variable_assignment] ...
variable_assignment:
  user_var_name = expr
  | [GLOBAL | SESSION] system_var_name = expr
  | [@@global. | @@session. | @@]system_var_name = expr
```

The `SET` statement assigns values to different types of variables that affect the operation of the server or your client. Older versions of MySQL employed `SET OPTION`, but this syntax is deprecated in favor of `SET` without `OPTION`.

This section describes use of `SET` for assigning values to system variables or user variables. For general information about these types of variables, see [Section 9.3, “User-Defined Variables”](#). System variables also can be set at server startup, as described in [Section 5.2.3, “Using System Variables”](#).

Some variants of `SET` syntax are used in other contexts:

- `SET PASSWORD` assigns account passwords. See [Section 13.5.1.4, “SET PASSWORD Syntax”](#).
- `SET TRANSACTION ISOLATION LEVEL` sets the isolation level for transaction processing. See [Section 13.4.6, “SET TRANSACTION Syntax”](#).

The following discussion shows the different `SET` syntaxes that you can use to set variables. The examples use the `=` assignment operator, but the `:=` operator also is allowable.

A user variable is written as `@var_name` and can be set as follows:

```
SET @var_name = expr;
```

As of MySQL 4.0.3, many system variables are dynamic and can be changed while the server runs by using the `SET` statement. For a list, see [Section 5.2.3.2, “Dynamic System Variables”](#). To change a system variable with `SET`, refer to it as `var_name`, optionally preceded by a modifier:

- To indicate explicitly that a variable is a global variable, precede its name by `GLOBAL` or `@@global..` The `SUPER` privilege is required to set global variables.

- To indicate explicitly that a variable is a session variable, precede its name by `SESSION`, `@@session.`, or `@@`. Setting a session variable requires no special privilege, but a client can change only its own session variables, not those of any other client.
- `LOCAL` and `@@local.` are synonyms for `SESSION` and `@@session.`
- If no modifier is present, `SET` changes the session variable.

A `SET` statement can contain multiple variable assignments, separated by commas. If you set several system variables, the most recent `GLOBAL` or `SESSION` modifier in the statement is used for following variables that have no modifier specified.

Examples:

```
SET sort_buffer_size=10000;
SET @@local.sort_buffer_size=10000;
SET GLOBAL sort_buffer_size=1000000, SESSION sort_buffer_size=1000000;
SET @@sort_buffer_size=1000000;
SET @@global.sort_buffer_size=1000000, @@local.sort_buffer_size=1000000;
```

When you assign a value to a system variable with `SET`, you cannot use suffix letters in the value (as can be done with startup options). However, the value can take the form of an expression:

```
SET sort_buffer_size = 10 * 1024 * 1024;
```

The `@@var_name` syntax for system variables is supported for compatibility with some other database systems.

If you change a session system variable, the value remains in effect until your session ends or until you change the variable to a different value. The change is not visible to other clients.

If you change a global system variable, the value is remembered and used for new connections until the server restarts. (To make a global system variable setting permanent, you should set it in an option file.) The change is visible to any client that accesses that global variable. However, the change affects the corresponding session variable only for clients that connect after the change. The global variable change does not affect the session variable for any client that is currently connected (not even that of the client that issues the `SET GLOBAL` statement).

To prevent incorrect usage, MySQL produces an error if you use `SET GLOBAL` with a variable that can only be used with `SET SESSION` or if you do not specify `GLOBAL` (or `@@global.`) when setting a global variable.

To set a `SESSION` variable to the `GLOBAL` value or a `GLOBAL` value to the compiled-in MySQL default value, use the `DEFAULT` keyword. For example, the following two statements are identical in setting the session value of `max_join_size` to the global value:

```
SET max_join_size=DEFAULT;
SET @@session.max_join_size=@@global.max_join_size;
```

Not all system variables can be set to `DEFAULT`. In such cases, use of `DEFAULT` results in an error.

You can refer to the values of specific global or session system variables in expressions by using one of the `@@`-modifiers. For example, you can retrieve values in a `SELECT` statement like this:

```
SELECT @@global.sql_mode, @@session.sql_mode, @@sql_mode;
```

When you refer to a system variable in an expression as `@@var_name` (that is, when you do not specify `@@global.` or `@@session.`), MySQL returns the session value if it exists and the global value oth-

erwise. (This differs from `SET @@var_name = value`, which always refers to the session value.)

To display system variables names and values, use the `SHOW VARIABLES` statement. (See [Section 13.5.4.20](#), “`SHOW VARIABLES Syntax`”.)

The following list describes options that have non-standard syntax or that are not described in the list of system variables found in [Section 5.2.2](#), “`System Variables`”. Although the options described here are not displayed by `SHOW VARIABLES`, you can obtain their values with `SELECT` (with the exception of `CHARACTER SET` and `SET NAMES`). For example:

```
mysql> SELECT @@AUTOCOMMIT;
+-----+
| @@AUTOCOMMIT |
+-----+
|              1 |
+-----+
```

The lettercase of these options does not matter.

- `AUTOCOMMIT = {0 | 1}`

Set the autocommit mode. If set to 1, all changes to a table take effect immediately. If set to 0 you have to use `COMMIT` to accept a transaction or `ROLLBACK` to cancel it. By default, client connections begin with `AUTOCOMMIT` set to 1. If you change `AUTOCOMMIT` mode from 0 to 1, MySQL performs an automatic `COMMIT` of any open transaction. Another way to begin a transaction is to use a `START TRANSACTION` or `BEGIN` statement. See [Section 13.4.1](#), “`START TRANSACTION, COMMIT, and ROLLBACK Syntax`”.

- `BIG_TABLES = {0 | 1}`

If set to 1, all temporary tables are stored on disk rather than in memory. This is a little slower, but the error `The table tbl_name is full` does not occur for `SELECT` operations that require a large temporary table. The default value for a new connection is 0 (use in-memory temporary tables). As of MySQL 4.0, you should normally never need to set this variable, because MySQL automatically converts in-memory tables to disk-based tables as necessary. (*Note:* This variable previously was named `SQL_BIG_TABLES`.)

- `CHARACTER SET {charset_name | DEFAULT}`

This maps all strings from and to the client with the given mapping. Before MySQL 4.1, the only allowable value for `charset_name` is `cp1251_koi8`, but you can add new mappings by editing the `sql/convert.cc` file in the MySQL source distribution. As of MySQL 4.1.1, `SET CHARACTER SET` sets three session system variables: `character_set_client` and `character_set_results` are set to the given character set, and `character_set_connection` to the value of `character_set_database`. See [Section 10.4](#), “`Connection Character Sets and Collations`”.

The default mapping can be restored by using the value `DEFAULT`. The default depends on the server configuration.

Note that the syntax for `SET CHARACTER SET` differs from that for setting most other options.

- `FOREIGN_KEY_CHECKS = {0 | 1}`

If set to 1 (the default), foreign key constraints for `InnoDB` tables are checked. If set to 0, they are ignored. Disabling foreign key checking can be useful for reloading `InnoDB` tables in an order different from that required by their parent/child relationships. This variable was added in MySQL 3.23.52. See [Section 14.2.7.4](#), “`FOREIGN KEY Constraints`”.

- `IDENTITY = value`

This variable is a synonym for the `LAST_INSERT_ID` variable. It exists for compatibility with other database systems. As of MySQL 3.23.25, you can read its value with `SELECT @@IDENTITY`. As of MySQL 4.0.3, you can also set its value with `SET IDENTITY`.

- `INSERT_ID = value`

Set the value to be used by the following `INSERT` or `ALTER TABLE` statement when inserting an `AUTO_INCREMENT` value. This is mainly used with the binary log.

- `LAST_INSERT_ID = value`

Set the value to be returned from `LAST_INSERT_ID()`. This is stored in the binary log when you use `LAST_INSERT_ID()` in a statement that updates a table. Setting this variable does not update the value returned by the `mysql_insert_id()` C API function.

- `NAMES {'charset_name' [COLLATE 'collation_name'] | DEFAULT}`

`SET NAMES` sets the three session system variables `character_set_client`, `character_set_connection`, and `character_set_results` to the given character set. Setting `character_set_connection` to `charset_name` also sets `collation_connection` to the default collation for `charset_name`. The optional `COLLATE` clause may be used to specify a collation explicitly. See [Section 10.4, “Connection Character Sets and Collations”](#).

The default mapping can be restored by using a value of `DEFAULT`. The default depends on the server configuration.

Note that the syntax for `SET NAMES` differs from that for setting most other options. This statement is available as of MySQL 4.1.0.

- `ONE_SHOT`

This option is a modifier, not a variable. It can be used to influence the effect of variables that set the character set, the collation, and the time zone. `ONE_SHOT` is primarily used for replication purposes: `mysqlbinlog` uses `SET ONE_SHOT` to modify temporarily the values of character set, collation, and time zone variables to reflect at rollforward what they were originally. `ONE_SHOT` is available as of MySQL 4.1.3.

You cannot use `ONE_SHOT` with other than the allowed set of variables; if you try, you get an error like this:

```
mysql> SET ONE_SHOT max_allowed_packet = 1;
ERROR 1382 (HY000): The 'SET ONE_SHOT' syntax is reserved for purposes
internal to the MySQL server
```

If `ONE_SHOT` is used with the allowed variables, it changes the variables as requested, but only for the next non-`SET` statement. After that, the server resets all character set, collation, and time zone-related system variables to their previous values. Example:

```
mysql> SET ONE_SHOT character_set_connection = latin5;
mysql> SET ONE_SHOT collation_connection = latin5_turkish_ci;
mysql> SHOW VARIABLES LIKE '%_connection';
+-----+-----+
| Variable_name | Value |
+-----+-----+
| character_set_connection | latin5 |
| collation_connection | latin5_turkish_ci |
+-----+-----+
mysql> SHOW VARIABLES LIKE '%_connection';
+-----+-----+
```

Variable_name	Value
character_set_connection	latin1
collation_connection	latin1_swedish_ci

- `SQL_AUTO_IS_NULL = {0 | 1}`

If set to 1 (the default), you can find the last inserted row for a table that contains an `AUTO_INCREMENT` column by using the following construct:

```
WHERE auto_increment_column IS NULL
```

This behavior is used by some ODBC programs, such as Access.

- `SQL_BIG_SELECTS = {0 | 1}`

If set to 0, MySQL aborts `SELECT` statements that are likely to take a very long time to execute (that is, statements for which the optimizer estimates that the number of examined rows exceeds the value of `max_join_size`). This is useful when an inadvisable `WHERE` statement has been issued. The default value for a new connection is 1, which allows all `SELECT` statements.

If you set the `max_join_size` system variable to a value other than `DEFAULT`, `SQL_BIG_SELECTS` is set to 0.

- `SQL_BUFFER_RESULT = {0 | 1}`

If set to 1, `SQL_BUFFER_RESULT` forces results from `SELECT` statements to be put into temporary tables. This helps MySQL free the table locks early and can be beneficial in cases where it takes a long time to send results to the client. The default value is 0. This variable was added in MySQL 3.23.13.

- `SQL_LOG_BIN = {0 | 1}`

If set to 0, no logging is done to the binary log for the client. The client must have the `SUPER` privilege to set this option. The default value is 1. This variable was added in MySQL 3.23.16.

- `SQL_LOG_OFF = {0 | 1}`

If set to 1, no logging is done to the general query log for this client. The client must have the `SUPER` privilege to set this option. The default value is 0.

- `SQL_LOG_UPDATE = {0 | 1}`

If set to 0, no logging is done to the update log for the client. The client must have the `SUPER` privilege to set this option. The default value is 1. This variable was added in MySQL 3.22.5.

- `SQL_NOTES = {0 | 1}`

If set to 1 (the default), warnings of `Note` level are recorded. If set to 0, `Note` warnings are suppressed. `mysqldump` includes output to set this variable to 0 so that reloading the dump file does not produce warnings for events that do not affect the integrity of the reload operation. `SQL_NOTES` was added in MySQL 4.1.11.

- `SQL_QUOTE_SHOW_CREATE = {0 | 1}`

If set to 1 (the default), the server quotes identifiers for `SHOW CREATE TABLE` and `SHOW CREATE DATABASE` statements. If set to 0, quoting is disabled. This option is enabled by default so that replication works for identifiers that require quoting. See [Section 13.5.4.5, “SHOW CREATE](#)

TABLE Syntax”, and Section 13.5.4.4, “SHOW CREATE DATABASE Syntax”. This variable was added in MySQL 3.23.26.

- `SQL_SAFE_UPDATES = {0 | 1}`

If set to 1, MySQL aborts `UPDATE` or `DELETE` statements that do not use a key in the `WHERE` clause or a `LIMIT` clause. This makes it possible to catch `UPDATE` or `DELETE` statements where keys are not used properly and that would probably change or delete a large number of rows. The default value is 0. This variable was added in MySQL 3.22.32.

- `SQL_SELECT_LIMIT = {value | DEFAULT}`

The maximum number of rows to return from `SELECT` statements. The default value for a new connection is “unlimited.” If you have changed the limit, the default value can be restored by using a `SQL_SELECT_LIMIT` value of `DEFAULT`.

If a `SELECT` has a `LIMIT` clause, the `LIMIT` takes precedence over the value of `SQL_SELECT_LIMIT`.

`SQL_SELECT_LIMIT` does not apply to `SELECT` statements executed within stored routines. It also does not apply to `SELECT` statements that do not produce a result set to be returned to the client. These include `SELECT` statements in subqueries, `CREATE TABLE ... SELECT`, and `INSERT INTO ... SELECT`.

- `SQL_WARNINGS = {0 | 1}`

This variable controls whether single-row `INSERT` statements produce an information string if warnings occur. The default is 0. Set the value to 1 to produce an information string. This variable was added in MySQL 3.22.11.

- `TIMESTAMP = {timestamp_value | DEFAULT}`

Set the time for this client. This is used to get the original timestamp if you use the binary log to restore rows. `timestamp_value` should be a Unix epoch timestamp, not a MySQL timestamp.

- `UNIQUE_CHECKS = {0 | 1}`

If set to 1 (the default), uniqueness checks for secondary indexes in `InnoDB` tables are performed. If set to 0, storage engines are allowed to assume that duplicate keys are not present in input data. If you know for certain that your data does not contain uniqueness violations, you can set this to 0 to speed up large table imports to `InnoDB`. This variable was added in MySQL 3.23.52.

Note that setting this variable to 0 does not *require* storage engines to ignore duplicate keys. An engine is still allowed to check for them and issue duplicate-key errors if it detects them.

## 13.5.4. SHOW Syntax

`SHOW` has many forms that provide information about databases, tables, columns, or status information about the server. This section describes those following:

```
SHOW [FULL] COLUMNS FROM tbl_name [FROM db_name] [LIKE 'pattern']
SHOW CREATE DATABASE db_name
SHOW CREATE TABLE tbl_name
SHOW DATABASES [LIKE 'pattern']
SHOW ENGINE engine_name {LOGS | STATUS }
SHOW [STORAGE] ENGINES
SHOW ERRORS [LIMIT [offset,] row_count]
SHOW GRANTS FOR user
SHOW INDEX FROM tbl_name [FROM db_name]
SHOW INNODB STATUS
SHOW [BDB] LOGS
```

```
SHOW PRIVILEGES
SHOW [FULL] PROCESSLIST
SHOW [GLOBAL | SESSION] STATUS [LIKE 'pattern']
SHOW TABLE STATUS [FROM db_name] [LIKE 'pattern']
SHOW [OPEN] TABLES [FROM db_name] [LIKE 'pattern']
SHOW [GLOBAL | SESSION] VARIABLES [LIKE 'pattern']
SHOW WARNINGS [LIMIT [offset,] row_count]
```

The `SHOW` statement also has forms that provide information about replication master and slave servers and are described in [Section 13.6, “Replication Statements”](#):

```
SHOW BINARY LOGS
SHOW BINLOG EVENTS
SHOW MASTER STATUS
SHOW SLAVE HOSTS
SHOW SLAVE STATUS
```

If the syntax for a given `SHOW` statement includes a `LIKE 'pattern'` part, `pattern` is a string that can contain the SQL `'%'` and `'_'` wildcard characters. The pattern is useful for restricting statement output to matching values.

### 13.5.4.1. SHOW CHARACTER SET Syntax

```
SHOW CHARACTER SET [LIKE 'pattern']
```

The `SHOW CHARACTER SET` statement shows all available character sets. It takes an optional `LIKE` clause that indicates which character set names to match. For example:

```
mysql> SHOW CHARACTER SET LIKE 'latin%';
```

Charset	Description	Default collation	Maxlen
latin1	cp1252 West European	latin1_swedish_ci	1
latin2	ISO 8859-2 Central European	latin2_general_ci	1
latin5	ISO 8859-9 Turkish	latin5_turkish_ci	1
latin7	ISO 8859-13 Baltic	latin7_general_ci	1

The `Maxlen` column shows the maximum number of bytes required to store one character.

`SHOW CHARACTER SET` is available as of MySQL 4.1.0.

### 13.5.4.2. SHOW COLLATION Syntax

```
SHOW COLLATION [LIKE 'pattern']
```

The output from `SHOW COLLATION` includes all available character sets. It takes an optional `LIKE` clause whose `pattern` indicates which collation names to match. For example:

```
mysql> SHOW COLLATION LIKE 'latin1%';
```

Collation	Charset	Id	Default	Compiled	Sortlen
latin1_german1_ci	latin1	5			0
latin1_swedish_ci	latin1	8	Yes	Yes	0
latin1_danish_ci	latin1	15			0
latin1_german2_ci	latin1	31		Yes	2
latin1_bin	latin1	47		Yes	0
latin1_general_ci	latin1	48			0
latin1_general_cs	latin1	49			0
latin1_spanish_ci	latin1	94			0

The `Default` column indicates whether a collation is the default for its character set. `Compiled in-`

icates whether the character set is compiled into the server. `Sortlen` is related to the amount of memory required to sort strings expressed in the character set.

`SHOW COLLATION` is available as of MySQL 4.1.0.

### 13.5.4.3. SHOW COLUMNS Syntax

```
SHOW [FULL] COLUMNS FROM tbl_name [FROM db_name] [LIKE 'pattern']
```

`SHOW COLUMNS` displays information about the columns in a given table.

If the data types differ from what you expect them to be based on your `CREATE TABLE` statement, note that MySQL sometimes changes data types when you create or alter a table. The conditions for which this occurs are described in [Section 13.1.5.1, “Silent Column Specification Changes”](#).

The `FULL` keyword can be used from MySQL 3.23.32 on. It causes the output to include the privileges you have for each column. As of MySQL 4.1, `FULL` also causes any per-column comments to be displayed.

You can use `db_name.tbl_name` as an alternative to the `tbl_name FROM db_name` syntax. In other words, these two statements are equivalent:

```
mysql> SHOW COLUMNS FROM mytable FROM mydb;
mysql> SHOW COLUMNS FROM mydb.mytable;
```

`SHOW FIELDS` is a synonym for `SHOW COLUMNS`. You can also list a table's columns with the `mysqlshow db_name tbl_name` command.

The `DESCRIBE` statement provides information similar to `SHOW COLUMNS`. See [Section 13.3.1, “DESCRIBE Syntax”](#).

### 13.5.4.4. SHOW CREATE DATABASE Syntax

```
SHOW CREATE DATABASE db_name
```

Shows the `CREATE DATABASE` statement that creates the given database. It was added in MySQL 4.1.

```
mysql> SHOW CREATE DATABASE test\G
***** 1. row *****
      Database: test
Create Database: CREATE DATABASE `test`
                  /*!40100 DEFAULT CHARACTER SET latin1 */
```

`SHOW CREATE DATABASE` quotes table and column names according to the value of the `SQL_QUOTE_SHOW_CREATE` option. See [Section 13.5.3, “SET Syntax”](#).

### 13.5.4.5. SHOW CREATE TABLE Syntax

```
SHOW CREATE TABLE tbl_name
```

Shows the `CREATE TABLE` statement that creates the given table. It was added in MySQL 3.23.20.

```
mysql> SHOW CREATE TABLE t\G
***** 1. row *****
      Table: t
Create Table: CREATE TABLE t (
  id INT(11) default NULL auto_increment,
  s char(60) default NULL,
  PRIMARY KEY (id)
) ENGINE=MyISAM
```



`SHOW CREATE TABLE` quotes table and column names according to the value of the `SQL_QUOTE_SHOW_CREATE` option. See [Section 13.5.3, “SET Syntax”](#).

### 13.5.4.6. SHOW DATABASES Syntax

```
SHOW DATABASES [LIKE 'pattern']
```

`SHOW DATABASES` lists the databases on the MySQL server host. As of MySQL 4.0.2, you see only those databases for which you have some kind of privilege, if you do not have the global `SHOW DATABASES` privilege. You can also get this list using the `mysqlshow` command.

If the server was started with the `--skip-show-database` option, you cannot use this statement at all unless you have the `SHOW DATABASES` privilege.

### 13.5.4.7. SHOW ENGINE Syntax

```
SHOW ENGINE engine_name {LOGS | STATUS }
```

`SHOW ENGINE` displays log or status information about a storage engine. The following statements currently are supported:

```
SHOW ENGINE BDB LOGS
SHOW ENGINE INNODB STATUS
```

`SHOW ENGINE BDB LOGS` displays status information about existing `BDB` log files. It returns the following fields:

- `File`  
The full path to the log file.
- `Type`  
The log file type (`BDB` for Berkeley DB log files).
- `Status`  
The status of the log file (`FREE` if the file can be removed, or `IN USE` if the file is needed by the transaction subsystem)

`SHOW ENGINE INNODB STATUS` displays extensive information about the state of the `InnoDB` storage engine.

The `InnoDB` Monitors provide additional information about `InnoDB` processing. See [Section 14.2.12.1, “SHOW ENGINE INNODB STATUS and the InnoDB Monitors”](#).

Older (and now deprecated) synonyms for these statements are `SHOW [BDB] LOGS` and `SHOW INNODB STATUS`.

`SHOW ENGINE` can be used as of MySQL 4.1.2.

### 13.5.4.8. SHOW ENGINES Syntax

```
SHOW [STORAGE] ENGINES
```

**SHOW ENGINES** displays status information about the the server's storage engines. This is particularly useful for checking whether a storage engine is supported, or to see what the default engine is. This statement is implemented in MySQL 4.1.2. **SHOW TABLE TYPES** is a deprecated synonym.

```
mysql> SHOW ENGINES\G
***** 1. row *****
Engine: MyISAM
Support: DEFAULT
Comment: Default engine as of MySQL 3.23 with great performance
***** 2. row *****
Engine: HEAP
Support: YES
Comment: Alias for MEMORY
***** 3. row *****
Engine: MEMORY
Support: YES
Comment: Hash based, stored in memory, useful for temporary tables
***** 4. row *****
Engine: MERGE
Support: YES
Comment: Collection of identical MyISAM tables
***** 5. row *****
Engine: MRG_MYISAM
Support: YES
Comment: Alias for MERGE
***** 6. row *****
Engine: ISAM
Support: NO
Comment: Obsolete storage engine, now replaced by MyISAM
***** 7. row *****
Engine: MRG_ISAM
Support: NO
Comment: Obsolete storage engine, now replaced by MERGE
***** 8. row *****
Engine: InnoDB
Support: YES
Comment: Supports transactions, row-level locking, and foreign keys
***** 9. row *****
Engine: INNODB
Support: YES
Comment: Alias for INNODB
***** 10. row *****
Engine: BDB
Support: YES
Comment: Supports transactions and page-level locking
***** 11. row *****
Engine: BERKELEYDB
Support: YES
Comment: Alias for BDB
***** 12. row *****
Engine: NDBCLUSTER
Support: NO
Comment: Clustered, fault-tolerant, memory-based tables
***** 13. row *****
Engine: NDB
Support: NO
Comment: Alias for NDBCLUSTER
***** 14. row *****
Engine: EXAMPLE
Support: NO
Comment: Example storage engine
***** 15. row *****
Engine: ARCHIVE
Support: NO
Comment: Archive storage engine
***** 16. row *****
Engine: CSV
Support: NO
Comment: CSV storage engine
***** 17. row *****
Engine: BLACKHOLE
Support: NO
Comment: Storage engine designed to act as null storage
```

The **Support** value indicates whether the particular storage engine is supported, and which is the default engine. For example, if the server is started with the `--default-table-type=InnoDB` op-

tion, the `Support` value for the `InnoDB` row has the value `DEFAULT`. See [Chapter 14, Storage Engines and Table Types](#).

All MySQL servers beginning with the 3.23 release series support `MyISAM` tables, because `MyISAM` is the default storage engine.

### 13.5.4.9. SHOW ERRORS Syntax

```
SHOW ERRORS [LIMIT [offset,] row_count]
SHOW COUNT(*) ERRORS
```

This statement is similar to `SHOW WARNINGS`, except that instead of displaying errors, warnings, and notes, it displays only errors. `SHOW ERRORS` is available as of MySQL 4.1.0.

The `LIMIT` clause has the same syntax as for the `SELECT` statement. See [Section 13.2.7, “SELECT Syntax”](#).

The `SHOW COUNT(*) ERRORS` statement displays the number of errors. You can also retrieve this number from the `error_count` variable:

```
SHOW COUNT(*) ERRORS;
SELECT @@error_count;
```

For more information, see [Section 13.5.4.21, “SHOW WARNINGS Syntax”](#).

### 13.5.4.10. SHOW GRANTS Syntax

```
SHOW GRANTS FOR user
```

This statement lists the `GRANT` statement or statements that must be issued to duplicate the privileges that are granted to a MySQL user account. The account is named using the same format as for the `GRANT` statement; for example, `'jeffrey'@'localhost'`. The user and host parts of the account name correspond to the `User` and `Host` column values of the `user` table row for the account.

```
mysql> SHOW GRANTS FOR 'root'@'localhost';
+-----+-----+
| Grants for root@localhost |
+-----+-----+
| GRANT ALL PRIVILEGES ON *.* TO 'root'@'localhost' WITH GRANT OPTION |
+-----+-----+
```

As of MySQL 4.1.2, to list the privileges granted to the account that you are using to connect to the server, you can use any of the following statements:

```
SHOW GRANTS;
SHOW GRANTS FOR CURRENT_USER;
SHOW GRANTS FOR CURRENT_USER();
```

Before MySQL 4.1.2, you can find out what user the session was authenticated as by selecting the value of the `CURRENT_USER()` function (new in MySQL 4.0.6). Then use that value in the `SHOW GRANTS` statement. See [Section 12.9.3, “Information Functions”](#).

`SHOW GRANTS` displays only the privileges granted explicitly to the named account. Other privileges might be available to the account, but they are not displayed. For example, if an anonymous account exists, the named account might be able to use its privileges, but `SHOW GRANTS` will not display them.

`SHOW GRANTS` is available as of MySQL 3.23.4.

### 13.5.4.11. SHOW INDEX Syntax

```
SHOW INDEX FROM tbl_name [FROM db_name]
```

`SHOW INDEX` returns table index information. The format resembles that of the `SQLStatistics` call in ODBC.

`SHOW INDEX` returns the following fields:

- `Table`  
The name of the table.
- `Non_unique`  
0 if the index cannot contain duplicates, 1 if it can.
- `Key_name`  
The name of the index.
- `Seq_in_index`  
The column sequence number in the index, starting with 1.
- `Column_name`  
The column name.
- `Collation`  
How the column is sorted in the index. In MySQL, this can have values 'A' (Ascending) or `NULL` (Not sorted).
- `Cardinality`  
An estimate of the number of unique values in the index. This is updated by running `ANALYZE TABLE` or `myisamchk -a`. `Cardinality` is counted based on statistics stored as integers, so the value is not necessarily exact even for small tables. The higher the cardinality, the greater the chance that MySQL uses the index when doing joins.
- `Sub_part`  
The number of indexed characters if the column is only partly indexed, `NULL` if the entire column is indexed.
- `Packed`  
Indicates how the key is packed. `NULL` if it is not.
- `Null`  
Contains `YES` if the column may contain `NULL`. If not, in MySQL 4.1 and earlier, the column contains an empty string ('').
- `Index_type`  
The index method used (`BTREE`, `FULLTEXT`, `HASH`, `RTREE`).
- `Comment`

Various remarks. Before MySQL 4.0.2 when the `Index_type` column was added, `Comment` indicates whether an index is `FULLTEXT`.

The `Packed` and `Comment` columns were added in MySQL 3.23.0. The `Null` and `Index_type` columns were added in MySQL 4.0.2.

You can use `db_name.tbl_name` as an alternative to the `tbl_name FROM db_name` syntax. These two statements are equivalent:

```
SHOW INDEX FROM mytable FROM mydb;  
SHOW INDEX FROM mydb.mytable;
```

`SHOW KEYS` is a synonym for `SHOW INDEX`. You can also list a table's indexes with the `mysqlshow -k db_name tbl_name` command.

### 13.5.4.12. SHOW INNODB STATUS Syntax

```
SHOW INNODB STATUS
```

This statement shows extensive information about the state of the `InnoDB` storage engine. As of MySQL 4.1.2, it is deprecated and `SHOW ENGINE INNODB STATUS` should be used instead. See [Section 13.5.4.7, “SHOW ENGINE Syntax”](#).

### 13.5.4.13. SHOW LOGS Syntax

```
SHOW [BDB] LOGS
```

`SHOW LOGS` displays status information about existing `BDB` log files. It was implemented in MySQL 3.23.29. An alias for it (available as of MySQL 4.1.1) is `SHOW BDB LOGS`. As of MySQL 4.1.2, this statement is deprecated and `SHOW ENGINE BDB LOGS` should be used instead. See [Section 13.5.4.7, “SHOW ENGINE Syntax”](#).

### 13.5.4.14. SHOW OPEN TABLES Syntax

```
SHOW OPEN TABLES
```

`SHOW OPEN TABLES` lists the non-`TEMPORARY` tables that are currently open in the table cache. See [Section 7.4.8, “How MySQL Opens and Closes Tables”](#).

`SHOW OPEN TABLES` returns the following fields:

- `Database`

The database containing the table.

- `Table`

The table name.

- `In_use`

The number of times the table currently is in use by queries. If the count is zero, the table is open, but not currently being used.

- `Name_locked`

Whether the table name is locked. Name locking is used for operations such as dropping or renaming tables.

`SHOW OPEN TABLES` was added in MySQL 3.23.33.

### 13.5.4.15. `SHOW PRIVILEGES` Syntax

```
SHOW PRIVILEGES
```

`SHOW PRIVILEGES` shows the list of system privileges that the MySQL server supports. This statement is implemented as of MySQL 4.1.0. The exact list of privileges depends on the version of your server.

```
mysql> SHOW PRIVILEGES\G
***** 1. row *****
Privilege: Alter
Context: Tables
Comment: To alter the table
***** 2. row *****
Privilege: Create temporary tables
Context: Databases
Comment: To use CREATE TEMPORARY TABLE
***** 3. row *****
Privilege: Create
Context: Databases,Tables,Indexes
Comment: To create new databases and tables
***** 4. row *****
Privilege: Delete
Context: Tables
Comment: To delete existing rows
***** 5. row *****
Privilege: Drop
Context: Databases,Tables
Comment: To drop databases and tables
...
```

### 13.5.4.16. `SHOW PROCESSLIST` Syntax

```
SHOW [FULL] PROCESSLIST
```

`SHOW PROCESSLIST` shows you which threads are running. You can also get this information using the `mysqladmin processlist` command. If you have the `PROCESS` privilege, you can see all threads. Otherwise, you can see only your own threads (that is, threads associated with the MySQL account that you are using). See [Section 13.5.5.3, “KILL Syntax”](#). If you do not use the `FULL` keyword, only the first 100 characters of each statement are shown in the `Info` field.

This statement is very useful if you get the “too many connections” error message and want to find out what is going on. MySQL reserves one extra connection to be used by accounts that have the `SUPER` privilege, to ensure that administrators should always be able to connect and check the system (assuming that you are not giving this privilege to all your users).

The output of `SHOW PROCESSLIST` may look like this:

```
mysql> SHOW FULL PROCESSLIST\G
***** 1. row *****
Id: 1
User: system user
Host:
db: NULL
Command: Connect
Time: 1030455
State: Waiting for master to send event
```

```

Info: NULL
***** 2. row *****
Id: 2
User: system user
Host:
db: NULL
Command: Connect
Time: 1004
State: Has read all relay log; waiting for the slave I/O thread to update it
Info: NULL
***** 3. row *****
Id: 3112
User: replikator
Host: artemis:2204
db: NULL
Command: Binlog Dump
Time: 2144
State: Has sent all binlog to slave; waiting for binlog to be updated
Info: NULL
***** 4. row *****
Id: 3113
User: replikator
Host: iconnect2:45781
db: NULL
Command: Binlog Dump
Time: 2086
State: Has sent all binlog to slave; waiting for binlog to be updated
Info: NULL
***** 5. row *****
Id: 3123
User: stefan
Host: localhost
db: apollon
Command: Query
Time: 0
State: NULL
Info: SHOW FULL PROCESSLIST
5 rows in set (0.00 sec)

```

The columns have the following meaning:

- **Id**

The connection identifier.

- **User**

The MySQL user who issued the statement. If this is `system user`, it refers to a non-client thread spawned by the server to handle tasks internally. This could be the I/O or SQL thread used on replication slaves or a delayed-row handler. For `system user`, there is no host specified in the `Host` column.

- **Host**

The hostname of the client issuing the statement (except for `system user` where there is no host). As of MySQL 4.0.12, `SHOW PROCESSLIST` reports the hostname for TCP/IP connections in `host_name:client_port` format to make it easier to determine which client is doing what.

- **db**

The default database, if one is selected, otherwise `NULL`.

- **Command**

The value of that column corresponds to the `COM_xxx` commands of the client/server protocol. See [Section 5.2.4, “Status Variables”](#)

The `Command` value may be one of the following: `Binlog Dump`, `Change user`, `Close`

stmt, Connect, Connect Out, Create DB, Daemon, Debug, Delayed insert, Drop DB, Error, Execute, Fetch, Field List, Init DB, Kill, Long Data, Ping, Prepare, Processlist, Query, Quit, Refresh, Register Slave, Reset stmt, Set option, Shutdown, Sleep, Statistics, Table Dump, Time

- Time

The time in seconds between the start of the statement or command and now.

- State

An action, event, or state, which can be one of the following: After create, Analyzing, Changing master, Checking master version, Checking table, Connecting to master, Copying to group table, Copying to tmp table, Creating delayed handler, Creating index, Creating sort index, Creating table from master dump, Creating tmp table, Execution of init\_command, FULLTEXT initialization, Finished reading one binlog; switching to next binlog, Flushing tables, Killed, Killing slave, Locked, Making temp file , Opening master dump table, Opening table, Opening tables, Processing request, Purging old relay logs, Queueing master event to the relay log, Reading event from the relay log, Reading from net, Reading master dump table data, Rebuilding the index on master dump table, Reconnecting after a failed binlog dump request, Reconnecting after a failed master event read, Registering slave on master, Removing duplicates, Reopen tables, Repair by sorting, Repair done, Repair with keycache, Requesting binlog dump, Rolling back, Saving state, Searching rows for update, Sending binlog event to slave, Sending data, Sorting for group, Sorting for order, Sorting index, Sorting result, System lock, Table lock, Thread initialized, Updating, User lock, Waiting for INSERT, Waiting for master to send event, Waiting for master update, Waiting for slave mutex on exit, Waiting for table, Waiting for tables, Waiting for the next event in relay log, Waiting on cond, Waiting to finalize termination, Waiting to reconnect after a failed binlog dump request, Waiting to reconnect after a failed master event read, Writing to net, allocating local table, cleaning up, closing tables, converting HEAP to MyISAM, copy to tmp table, creating table, deleting from main table, deleting from reference tables, discard\_or\_import\_tablespace, end, freeing items, got handler lock, got old table, info, init, insert, logging slow query, login, preparing, purging old relay logs, query end, removing tmp table, rename, rename result table, reschedule, setup, starting slave, statistics, storing row into queue, unauthenticated user, update, updating, updating main table, updating reference tables, upgrading lock, waiting for delay\_list, waiting for handler insert, waiting for handler lock, waiting for handler open, Waiting for event from ndbcluster

The most common `State` values are described in the rest of this section. Most of the other `State` values are useful only for finding bugs in the server. See also [Section 6.3, “Replication Implementation Details”](#), for additional information about process states for replication servers.

For the `SHOW PROCESSLIST` statement, the value of `State` is `NULL`.

- Info

The statement that the thread is executing, or `NULL` if it is not executing any statement.

Some `State` values commonly seen in the output from `SHOW PROCESSLIST`:



- `Checking table`

The thread is performing a table check operation.

- `Closing tables`

Means that the thread is flushing the changed table data to disk and closing the used tables. This should be a fast operation. If not, you should verify that you do not have a full disk and that the disk is not in very heavy use.

- `Connect Out`

A replication slave is connecting to its master.

- `Copying to group table`

If a statement has different `ORDER BY` and `GROUP BY` criteria, the rows are sorted by group and copied to a temporary table.

- `Copying to tmp table`

The server is copying to a temporary table in memory.

- `Copying to tmp table on disk`

The server is copying to a temporary table on disk. The temporary result set was larger than `tmp_table_size` and the thread is changing the temporary table from in-memory to disk-based format to save memory.

- `Creating tmp table`

The thread is creating a temporary table to hold a part of the result for the query.

- `deleting from main table`

The server is executing the first part of a multiple-table delete. It is deleting only from the first table, and saving fields and offsets to be used for deleting from the other (reference) tables.

- `deleting from reference tables`

The server is executing the second part of a multiple-table delete and deleting the matched rows from the other tables.

- `Flushing tables`

The thread is executing `FLUSH TABLES` and is waiting for all threads to close their tables.

- `FULLTEXT initialization`

The server is preparing to perform a natural-language full-text search.

- `Killed`

Someone has sent a `KILL` statement to the thread and it should abort next time it checks the kill flag. The flag is checked in each major loop in MySQL, but in some cases it might still take a short time for the thread to die. If the thread is locked by some other thread, the kill takes effect as soon as the other thread releases its lock.

- `Locked`

The query is locked by another query.

- `Sending data`

The thread is processing rows for a `SELECT` statement and also is sending data to the client.

- `Sorting for group`

The thread is doing a sort to satisfy a `GROUP BY`.

- `Sorting for order`

The thread is doing a sort to satisfy a `ORDER BY`.

- `Opening tables`

The thread is trying to open a table. This should be very fast procedure, unless something prevents opening. For example, an `ALTER TABLE` or a `LOCK TABLE` statement can prevent opening a table until the statement is finished.

- `Reading from net`

The server is reading a packet from the network.

- `Removing duplicates`

The query was using `SELECT DISTINCT` in such a way that MySQL could not optimize away the distinct operation at an early stage. Because of this, MySQL requires an extra stage to remove all duplicated rows before sending the result to the client.

- `Reopen table`

The thread got a lock for the table, but noticed after getting the lock that the underlying table structure changed. It has freed the lock, closed the table, and is trying to reopen it.

- `Repair by sorting`

The repair code is using a sort to create indexes.

- `Repair with keycache`

The repair code is using creating keys one by one through the key cache. This is much slower than `Repair by sorting`.

- `Searching rows for update`

The thread is doing a first phase to find all matching rows before updating them. This has to be done if the `UPDATE` is changing the index that is used to find the involved rows.

- `Sleeping`

The thread is waiting for the client to send a new statement to it.

- `statistics`

The server is calculating statistics to develop a query execution plan.

- `System lock`

The thread is waiting to get an external system lock for the table. If you are not using multiple `mysqld` servers that are accessing the same tables, you can disable system locks with the `-skip-external-locking` option.

- `unauthenticated user`

The state of a thread that has become associated with a client connection but for which authentication of the client user has not yet been done.

- `Upgrading lock`

The `INSERT DELAYED` handler is trying to get a lock for the table to insert rows.

- `Updating`

The thread is searching for rows to update and is updating them.

- `updating main table`

The server is executing the first part of a multiple-table update. It is updating only the first table, and saving fields and offsets to be used for updating the other (reference) tables.

- `updating reference tables`

The server is executing the second part of a multiple-table update and updating the matched rows from the other tables.

- `User Lock`

The thread is waiting on a `GET_LOCK()`.

- `Waiting for event from ndbcluster`

The server is acting as an SQL node in a MySQL Cluster, and is connected to a cluster management node.

- `Waiting for tables`

The thread got a notification that the underlying structure for a table has changed and it needs to reopen the table to get the new structure. However, to reopen the table, it must wait until all other threads have closed the table in question.

This notification takes place if another thread has used `FLUSH TABLES` or one of the following statements on the table in question: `FLUSH TABLES tbl_name`, `ALTER TABLE`, `RENAME TABLE`, `REPAIR TABLE`, `ANALYZE TABLE`, or `OPTIMIZE TABLE`.

- `waiting for handler insert`

The `INSERT DELAYED` handler has processed all pending inserts and is waiting for new ones.

- `Writing to net`

The server is writing a packet to the network.

Most states correspond to very quick operations. If a thread stays in any of these states for many seconds, there might be a problem that needs to be investigated.

### 13.5.4.17. SHOW STATUS Syntax

```
SHOW STATUS [LIKE 'pattern']
```

`SHOW STATUS` provides server status information. This information also can be obtained using the `mysqladmin extended-status` command.

Partial output is shown here. The list of names and values may be different for your server. The meaning of each variable is given in [Section 5.2.4, “Status Variables”](#).

```
mysql> SHOW STATUS;
+-----+-----+
| Variable_name | Value |
+-----+-----+
| Aborted_clients | 0 |
| Aborted_connects | 0 |
| Bytes_received | 155372598 |
| Bytes_sent | 1176560426 |
| Connections | 30023 |
| Created_tmp_disk_tables | 0 |
| Created_tmp_tables | 8340 |
| Created_tmp_files | 60 |
| ... | ... |
| Open_tables | 1 |
| Open_files | 2 |
| Open_streams | 0 |
| Opened_tables | 44600 |
| Questions | 2026873 |
| ... | ... |
| Table_locks_immediate | 1920382 |
| Table_locks_waited | 0 |
| Threads_cached | 0 |
| Threads_created | 30022 |
| Threads_connected | 1 |
| Threads_running | 1 |
| Uptime | 80380 |
+-----+-----+
```

With a `LIKE` clause, the statement displays only rows for those variables with names that match the pattern:

```
mysql> SHOW STATUS LIKE 'Key%';
+-----+-----+
| Variable_name | Value |
+-----+-----+
| Key_blocks_used | 14955 |
| Key_read_requests | 96854827 |
| Key_reads | 162040 |
| Key_write_requests | 7589728 |
| Key_writes | 3813196 |
+-----+-----+
```

### 13.5.4.18. SHOW TABLE STATUS Syntax

```
SHOW TABLE STATUS [FROM db_name] [LIKE 'pattern']
```

`SHOW TABLE STATUS` works like `SHOW TABLE`, but provides a lot of information about each table. You can also get this list using the `mysqlshow --status db_name` command. This statement was added in MySQL 3.23.

`SHOW TABLE STATUS` returns the following fields:

- Name

The name of the table.

- `Engine`

The storage engine for the table. See [Chapter 14, Storage Engines and Table Types](#). Before MySQL 4.1.2, this value is labeled as `Type`.
- `Version`

The version number of the table's `.frm` file.
- `Row_format`

The row storage format (`Fixed`, `Dynamic`, `Compressed`, `Redundant`, `Compact`). InnoDB tables are always in the `Redundant` format.
- `Rows`

The number of rows. Some storage engines, such as `MyISAM` and `ISAM`, store the exact count. For other storage engines, such as `InnoDB`, this value is an approximation, and may vary from the actual value by as much as 40 to 50%. In such cases, use `SELECT COUNT(*)` to obtain an accurate count.
- `Avg_row_length`

The average row length.
- `Data_length`

The length of the data file.
- `Max_data_length`

The maximum length of the data file. This is the total number of bytes of data that can be stored in the table, given the data pointer size used.
- `Index_length`

The length of the index file.
- `Data_free`

The number of allocated but unused bytes.
- `Auto_increment`

The next `AUTO_INCREMENT` value.
- `Create_time`

When the table was created.
- `Update_time`

When the data file was last updated. For some storage engines, this value is `NULL`. For example, `InnoDB` stores multiple tables in its tablespace and the data file timestamp does not apply.
- `Check_time`

When the table was last checked. Not all storage engines update this time, in which case the value is always `NULL`.

- `Collation`  
The table's character set and collation. (Implemented in 4.1.1)
- `Checksum`  
The live checksum value (if any). (Implemented in 4.1.1)
- `Create_options`  
Extra options used with `CREATE TABLE`.
- `Comment`  
The comment used when creating the table (or information as to why MySQL could not access the table information).

In the table comment, `InnoDB` tables report the free space of the tablespace to which the table belongs. For a table located in the shared tablespace, this is the free space of the shared tablespace. If you are using multiple tablespaces and the table has its own tablespace, the free space is for only that table.

For `MEMORY` (`HEAP`) tables, the `Data_length`, `Max_data_length`, and `Index_length` values approximate the actual amount of allocated memory. The allocation algorithm reserves memory in large amounts to reduce the number of allocation operations.

For views, all the fields displayed by `SHOW TABLE STATUS` are `NULL` except that `Name` indicates the view name and `Comment` says `view`.

### 13.5.4.19. `SHOW TABLES` Syntax

```
SHOW TABLES [FROM db_name] [LIKE 'pattern']
```

`SHOW TABLES` lists the non-`TEMPORARY` tables in a given database. You can also get this list using the `mysqlshow db_name` command.

The output from `SHOW TABLES` contains a single column of table names.

**Note:** If you have no privileges for a table, the table does not show up in the output from `SHOW TABLES` or `mysqlshow db_name`.

### 13.5.4.20. `SHOW VARIABLES` Syntax

```
SHOW [GLOBAL | SESSION] VARIABLES [LIKE 'pattern']
```

`SHOW VARIABLES` shows the values of MySQL system variables. This information also can be obtained using the `mysqladmin variables` command.

The `GLOBAL` and `SESSION` modifiers are new in MySQL 4.0.3. With the `GLOBAL` modifier, `SHOW VARIABLES` displays the values that are used for new connections to MySQL. With `SESSION`, it displays the values that are in effect for the current connection. If no modifier is present, the default is `SESSION`. `LOCAL` is a synonym for `SESSION`.

If the default system variable values are unsuitable, you can set them using command options when `mysqld` starts, and most can be changed at runtime with the `SET` statement. See [Section 5.2.3, “Using System Variables”](#), and [Section 13.5.3, “SET Syntax”](#).

Partial output is shown here. The list of names and values may be different for your server. [Sec-](#)

tion 5.2.2, “System Variables”, describes the meaning of each variable, and Section 7.5.2, “Tuning Server Parameters”, provides information about tuning them.

```
mysql> SHOW VARIABLES;
```

Variable_name	Value
back_log	50
basedir	/usr/local/mysql
bdb_cache_size	8388572
bdb_log_buffer_size	32768
bdb_home	/usr/local/mysql
...	
max_connections	100
max_connect_errors	10
max_delayed_threads	20
max_error_count	64
max_heap_table_size	16777216
max_join_size	4294967295
max_relay_log_size	0
max_sort_length	1024
...	
timezone	EEST
tmp_table_size	33554432
tmpdir	/tmp/:/mnt/hd2/tmp/
version	4.1.18
wait_timeout	28800

With a `LIKE` clause, the statement displays only rows for those variables with names that match the pattern. To obtain the row for a specific variable, use a `LIKE` clause as shown:

```
SHOW VARIABLES LIKE 'max_join_size';
SHOW SESSION VARIABLES LIKE 'max_join_size';
```

To get a list of variables whose name match a pattern, use the `'%'` wildcard character in a `LIKE` clause:

```
SHOW VARIABLES LIKE '%size%';
SHOW GLOBAL VARIABLES LIKE '%size%';
```

Wildcard characters can be used in any position within the pattern to be matched. Strictly speaking, because `'_'` is a wildcard that matches any single character, you should escape it as `'\_'` to match it literally. In practice, this is rarely necessary.

### 13.5.4.21. SHOW WARNINGS Syntax

```
SHOW WARNINGS [LIMIT [offset,] row_count]
SHOW COUNT(*) WARNINGS
```

`SHOW WARNINGS` shows the error, warning, and note messages that resulted from the last statement that generated messages, or nothing if the last statement that used a table generated no messages. This statement is implemented as of MySQL 4.1.0. A related statement, `SHOW ERRORS`, shows only the errors. See Section 13.5.4.9, “`SHOW ERRORS Syntax`”.

The list of messages is reset for each new statement that uses a table.

The `SHOW COUNT(*) WARNINGS` statement displays the total number of errors, warnings, and notes. You can also retrieve this number from the `warning_count` variable:

```
SHOW COUNT(*) WARNINGS;
SELECT @@warning_count;
```

The value of `warning_count` might be greater than the number of messages displayed by `SHOW WARNINGS` if the `max_error_count` system variable is set so low that not all messages are stored. An example shown later in this section demonstrates how this can happen.

The `LIMIT` clause has the same syntax as for the `SELECT` statement. See [Section 13.2.7, “SELECT Syntax”](#).

The MySQL server sends back the total number of errors, warnings, and notes resulting from the last statement. If you are using the C API, this value can be obtained by calling `mysql_warning_count()`. See [Section 17.2.3.70, “mysql\\_warning\\_count\(\)”](#).

Note that the framework for warnings was added in MySQL 4.1.0, at which point many statements did not generate warnings. In 4.1.1, the situation is much improved, with warnings generated for statements such as `LOAD DATA INFILE` and DML statements such as `INSERT`, `UPDATE`, `CREATE TABLE`, and `ALTER TABLE`.

The following `DROP TABLE` statement results in a note:

```
mysql> DROP TABLE IF EXISTS no_such_table;
mysql> SHOW WARNINGS;
+-----+-----+-----+
| Level | Code | Message                                     |
+-----+-----+-----+
| Note  | 1051 | Unknown table 'no_such_table'            |
+-----+-----+-----+
```

Here is a simple example that shows a syntax warning for `CREATE TABLE` and conversion warnings for `INSERT`:

```
mysql> CREATE TABLE t1 (a TINYINT NOT NULL, b CHAR(4)) TYPE=MyISAM;
Query OK, 0 rows affected, 1 warning (0.00 sec)
mysql> SHOW WARNINGS\G
***** 1. row *****
Level: Warning
Code: 1287
Message: 'TYPE=storage_engine' is deprecated, use
'ENGINE=storage_engine' instead
1 row in set (0.00 sec)

mysql> INSERT INTO t1 VALUES(10,'mysql'),(NULL,'test'),
-> (300,'Open Source');
Query OK, 3 rows affected, 4 warnings (0.01 sec)
Records: 3 Duplicates: 0 Warnings: 4

mysql> SHOW WARNINGS\G
***** 1. row *****
Level: Warning
Code: 1265
Message: Data truncated for column 'b' at row 1
***** 2. row *****
Level: Warning
Code: 1263
Message: Data truncated, NULL supplied to NOT NULL column 'a' at row 2
***** 3. row *****
Level: Warning
Code: 1264
Message: Data truncated, out of range for column 'a' at row 3
***** 4. row *****
Level: Warning
Code: 1265
Message: Data truncated for column 'b' at row 3
4 rows in set (0.00 sec)
```

The maximum number of error, warning, and note messages to store is controlled by the `max_error_count` system variable. By default, its value is 64. To change the number of messages you want stored, change the value of `max_error_count`. In the following example, the `ALTER TABLE` statement produces three warning messages, but only one is stored because `max_error_count` has been set to 1:

```
mysql> SHOW VARIABLES LIKE 'max_error_count';
+-----+-----+
| Variable_name | Value |
+-----+-----+
| max_error_count | 64   |
+-----+-----+
```



```

+-----+
1 row in set (0.00 sec)

mysql> SET max_error_count=1;
Query OK, 0 rows affected (0.00 sec)

mysql> ALTER TABLE t1 MODIFY b CHAR;
Query OK, 3 rows affected, 3 warnings (0.00 sec)
Records: 3 Duplicates: 0 Warnings: 3

mysql> SELECT @@warning_count;
+-----+
| @@warning_count |
+-----+
|                 3 |
+-----+
1 row in set (0.01 sec)

mysql> SHOW WARNINGS;
+-----+-----+-----+
| Level      | Code | Message
+-----+-----+-----+
| Warning    | 1263 | Data truncated for column 'b' at row 1
+-----+-----+-----+
1 row in set (0.00 sec)

```

To disable warnings, set `max_error_count` to 0. In this case, `warning_count` still indicates how many warnings have occurred, but none of the messages are stored.

As of MySQL 4.1.11, you can set the `SQL_NOTES` session variable to 0 to cause `Note`-level warnings not to be recorded.

## 13.5.5. Other Administrative Statements

### 13.5.5.1. CACHE INDEX Syntax

```

CACHE INDEX
tbl_index_list [, tbl_index_list] ...
IN key_cache_name

tbl_index_list:
tbl_name [[INDEX|KEY] (index_name[, index_name] ...)]

```

The `CACHE INDEX` statement assigns table indexes to a specific key cache. It is used only for `MyISAM` tables.

The following statement assigns indexes from the tables `t1`, `t2`, and `t3` to the key cache named `hot_cache`:

```

mysql> CACHE INDEX t1, t2, t3 IN hot_cache;
+-----+-----+-----+-----+
| Table  | Op           | Msg_type | Msg_text
+-----+-----+-----+-----+
| test.t1 | assign_to_keycache | status   | OK
| test.t2 | assign_to_keycache | status   | OK
| test.t3 | assign_to_keycache | status   | OK
+-----+-----+-----+-----+

```

The syntax of `CACHE INDEX` enables you to specify that only particular indexes from a table should be assigned to the cache. The current implementation assigns all the table's indexes to the cache, so there is no reason to specify anything other than the table name.

The key cache referred to in a `CACHE INDEX` statement can be created by setting its size with a parameter setting statement or in the server parameter settings. For example:

```

mysql> SET GLOBAL keycache1.key_buffer_size=128*1024;

```

Key cache parameters can be accessed as members of a structured system variable. See [Section 5.2.3.1, “Structured System Variables”](#).

A key cache must exist before you can assign indexes to it:

```
mysql> CACHE INDEX t1 IN non_existent_cache;
ERROR 1284 (HY000): Unknown key cache 'non_existent_cache'
```

By default, table indexes are assigned to the main (default) key cache created at the server startup. When a key cache is destroyed, all indexes assigned to it become assigned to the default key cache again.

Index assignment affects the server globally: If one client assigns an index to a given cache, this cache is used for all queries involving the index, no matter which client issues the queries.

`CACHE INDEX` was added in MySQL 4.1.1.

### 13.5.5.2. FLUSH Syntax

```
FLUSH [LOCAL | NO_WRITE_TO_BINLOG] flush_option [, flush_option] ...
```

The `FLUSH` statement clears or reloads various internal caches used by MySQL. To execute `FLUSH`, you must have the `RELOAD` privilege.

The `RESET` statement is similar to `FLUSH`. See [Section 13.5.5.5, “RESET Syntax”](#).

*flush\_option* can be any of the following:

- `HOSTS`

Empties the host cache tables. You should flush the host tables if some of your hosts change IP number or if you get the error message `Host 'host_name' is blocked`. When more than `max_connect_errors` errors occur successively for a given host while connecting to the MySQL server, MySQL assumes that something is wrong and blocks the host from further connection requests. Flushing the host tables allows the host to attempt to connect again. See [Section A.2.5, “Host 'host\\_name' is blocked”](#). You can start `mysqld` with `--max_connect_errors=999999999` to avoid this error message.

- `DES_KEY_FILE`

Reloads the DES keys from the file that was specified with the `--des-key-file` option at server startup time.

- `LOGS`

Closes and reopens all log files. If you have specified an update log file or a binary log file without an extension, the extension number of the log file is incremented by one relative to the previous file. If you have used an extension in the file name, MySQL closes and reopens the update log or binary log file. See [Section 5.11.4, “The Binary Log”](#). On Unix, this is the same thing as sending a `SIGHUP` signal to the `mysqld` server (except on some Mac OS X 10.3 versions where `mysqld` ignores `SIGHUP` and `SIGQUIT`).

Beginning with MySQL 4.0.10, if the server was started with the `--log-error` option, `FLUSH LOGS` causes the error log is renamed with a suffix of `-old` and `mysqld` creates a new empty log file. No renaming occurs if the `--log-error` option was not given.

- `MASTER (DEPRECATED)`. Deletes all binary logs, resets the binary log index file and creates a new binary log. Deprecated in favor of `RESET MASTER`, supported for backwards compatibility only

See [Section 13.6.1.2, “RESET MASTER Syntax”](#).

- `PRIVILEGES`

Reloads the privileges from the grant tables in the `mysql` database.

- `QUERY CACHE`

Defragment the query cache to better utilize its memory. `FLUSH QUERY CACHE` does not remove any queries from the cache, unlike `RESET QUERY CACHE`.

- `SLAVE (DEPRECATED)`. Resets all replication slave parameters, including relay log files and replication position in the master's binary logs. Deprecated in favor of `RESET SLAVE`, supported for backwards compatibility only See [Section 13.6.2.5, “RESET SLAVE Syntax”](#).

- `STATUS`

Resets most status variables to zero. This is something you should use only when debugging a query. See [Section 1.8, “How to Report Bugs or Problems”](#).

- `{TABLE | TABLES} [tbl_name [, tbl_name] ...]`

When no tables are named, closes all open tables and forces all tables in use to be closed. This also flushes the query cache. With one or more table names, flushes only the given tables. `FLUSH TABLES` also removes all query results from the query cache, like the `RESET QUERY CACHE` statement.

- `TABLES WITH READ LOCK`

Closes all open tables and locks all tables for all databases with a read lock until you execute `UNLOCK TABLES`. This is very convenient way to get backups if you have a filesystem such as Veritas that can take snapshots in time.

- `USER_RESOURCES`

Resets all per-hour user resources to zero. This enables clients that have reached their hourly connection, query, or update limits to resume activity immediately. `FLUSH USER_RESOURCES` does not apply to the limit on maximum simultaneous connections. See [Section 13.5.1.2, “GRANT Syntax”](#).

Before MySQL 4.1.1, `FLUSH` statements are not written to the binary log. As of MySQL 4.1.1, they are written to the binary log unless the optional `NO_WRITE_TO_BINLOG` keyword (or its alias `LOCAL`) is used. This is done so that `FLUSH` statements used on a MySQL server acting as a replication master will be replicated by default to the replication slave.

**Note:** `FLUSH LOGS`, `FLUSH MASTER`, `FLUSH SLAVE`, and `FLUSH TABLES WITH READ LOCK` are not logged in any case because they would cause problems if replicated to a slave.

You can also access some of these statements with the `mysqladmin` utility, using the `flush-hosts`, `flush-logs`, `flush-privileges`, `flush-status`, or `flush-tables` commands.

See also [Section 13.5.5.5, “RESET Syntax”](#), for information about how the `RESET` statement is used with replication.

### 13.5.5.3. `KILL` Syntax

```
KILL thread_id
```

Each connection to `mysqld` runs in a separate thread. You can see which threads are running with the `SHOW PROCESSLIST` statement and kill a thread with the `KILL thread_id` statement.

If you have the `PROCESS` privilege, you can see all threads. If you have the `SUPER` privilege, you can kill all threads and statements. Otherwise, you can see and kill only your own threads and statements.

You can also use the `mysqladmin processlist` and `mysqladmin kill` commands to examine and kill threads.

**Note:** You cannot use `KILL` with the Embedded MySQL Server library, because the embedded server merely runs inside the threads of the host application. It does not create any connection threads of its own.

When you use `KILL`, a thread-specific kill flag is set for the thread. In most cases, it might take some time for the thread to die, because the kill flag is checked only at specific intervals:

- In `SELECT`, `ORDER BY` and `GROUP BY` loops, the flag is checked after reading a block of rows. If the kill flag is set, the statement is aborted.
- During `ALTER TABLE`, the kill flag is checked before each block of rows are read from the original table. If the kill flag was set, the statement is aborted and the temporary table is deleted.
- During `UPDATE` or `DELETE` operations, the kill flag is checked after each block read and after each updated or deleted row. If the kill flag is set, the statement is aborted. Note that if you are not using transactions, the changes are not rolled back.
- `GET_LOCK()` aborts and returns `NULL`.
- An `INSERT DELAYED` thread quickly flushes (inserts) all rows it has in memory and then terminates.
- If the thread is in the table lock handler (state: `Locked`), the table lock is quickly aborted.
- If the thread is waiting for free disk space in a write call, the write is aborted with a “disk full” error message.
- Some threads might refuse to be killed. For example, `REPAIR TABLE`, `CHECK TABLE`, and `OPTIMIZE TABLE` cannot be killed before MySQL 4.1 and run to completion. This is changed: `REPAIR TABLE` and `OPTIMIZE TABLE` can be killed as of MySQL 4.1.0, as can `CHECK TABLE` as of MySQL 4.1.3. However, killing a `REPAIR TABLE` or `OPTIMIZE TABLE` operation on a `MyISAM` table results in a table that *is* corrupted and is unusable (reads and writes to it fail) until you optimize or repair it again (without interruption).

#### 13.5.5.4. LOAD INDEX INTO CACHE Syntax

```
LOAD INDEX INTO CACHE
  tbl_index_list [, tbl_index_list] ...

tbl_index_list:
  tbl_name
  [[INDEX|KEY] (index_name[, index_name] ...)]
  [IGNORE LEAVES]
```

The `LOAD INDEX INTO CACHE` statement preloads a table index into the key cache to which it has been assigned by an explicit `CACHE INDEX` statement, or into the default key cache otherwise. `LOAD INDEX INTO CACHE` is used only for `MyISAM` tables.

The `IGNORE LEAVES` modifier causes only blocks for the non-leaf nodes of the index to be preloaded.

The following statement preloads nodes (index blocks) of indexes for the tables `t1` and `t2`:

```
mysql> LOAD INDEX INTO CACHE t1, t2 IGNORE LEAVES;
+-----+-----+-----+-----+
| Table | Op      | Msg_type | Msg_text |
+-----+-----+-----+-----+
| test.t1 | preload_keys | status   | OK       |
| test.t2 | preload_keys | status   | OK       |
+-----+-----+-----+-----+
```

This statement preloads all index blocks from `t1`. It preloads only blocks for the non-leaf nodes from `t2`.

The syntax of `LOAD INDEX INTO CACHE` enables you to specify that only particular indexes from a table should be preloaded. The current implementation preloads all the table's indexes into the cache, so there is no reason to specify anything other than the table name.

`LOAD INDEX INTO CACHE` was added in MySQL 4.1.1.

### 13.5.5.5. RESET Syntax

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

The `RESET` statement is used to clear the state of various server operations. You must have the `RELOAD` privilege to execute `RESET`.

`RESET` acts as a stronger version of the `FLUSH` statement. See [Section 13.5.5.2, “FLUSH Syntax”](#).

*reset\_option* can be any of the following:

- `MASTER`

Deletes all binary logs listed in the index file, resets the binary log index file to be empty, and creates a new binary log file. Previously named `FLUSH MASTER`. See [Section 13.6.1, “SQL Statements for Controlling Master Servers”](#).
- `QUERY CACHE`

Removes all query results from the query cache.
- `SLAVE`

Makes the slave forget its replication position in the master binary logs. Also resets the relay log by deleting any existing relay log files and beginning a new one. Previously named `FLUSH SLAVE`. See [Section 13.6.2, “SQL Statements for Controlling Slave Servers”](#).

## 13.6. Replication Statements

This section describes SQL statements related to replication. One group of statements is used for controlling master servers. The other is used for controlling slave servers.

### 13.6.1. SQL Statements for Controlling Master Servers

Replication can be controlled through the SQL interface. This section discusses statements for managing master replication servers. [Section 13.6.2, “SQL Statements for Controlling Slave Servers”](#), discusses statements for managing slave servers.

### 13.6.1.1. PURGE MASTER LOGS Syntax

```
PURGE {MASTER | BINARY} LOGS TO 'log_name'
PURGE {MASTER | BINARY} LOGS BEFORE 'date'
```

Deletes all the binary logs listed in the log index prior to the specified log or date. The logs also are removed from the list recorded in the log index file, so that the given log becomes the first.

Example:

```
PURGE MASTER LOGS TO 'mysql-bin.010';
PURGE MASTER LOGS BEFORE '2003-04-02 22:46:26';
```

The **BEFORE** variant is available as of MySQL 4.1. Its date argument can be in 'YYYY-MM-DD hh:mm:ss' format. **MASTER** and **BINARY** are synonyms, but **BINARY** can be used only as of MySQL 4.1.1.

This statement is safe to run while slaves are replicating. You do not need to stop them. If you have an active slave that currently is reading one of the logs you are trying to delete, this statement does nothing and fails with an error. However, if a slave is dormant and you happen to purge one of the logs it has yet to read, the slave will be unable to replicate after it comes up.

To safely purge logs, follow this procedure:

1. On each slave server, use **SHOW SLAVE STATUS** to check which log it is reading.
2. Obtain a listing of the binary logs on the master server with **SHOW BINARY LOGS**.
3. Determine the earliest log among all the slaves. This is the target log. If all the slaves are up to date, this is the last log on the list.
4. Make a backup of all the logs you are about to delete. (This step is optional, but always advisable.)
5. Purge all logs up to but not including the target log.

You can also set the `expire_logs_days` system variable to expire binary log files automatically after a given number of days (see [Section 5.2.2, “System Variables”](#)). If you are using replication, you should set the variable no lower than the maximum number of days your slaves might lag behind the master.

### 13.6.1.2. RESET MASTER Syntax

```
RESET MASTER
```

Deletes all binary logs listed in the index file, resets the binary log index file to be empty, and creates a new binary log file.

This statement was named **FLUSH MASTER** before MySQL 3.23.26.

### 13.6.1.3. SET SQL\_LOG\_BIN Syntax

```
SET SQL_LOG_BIN = {0|1}
```

Disables or enables binary logging for the current connection (`SQL_LOG_BIN` is a session variable) if the client that has the **SUPER** privilege. The statement is refused with an error if the client does not have

that privilege. (Before MySQL 4.1.2, the statement was simply ignored in that case.)

### 13.6.1.4. SHOW BINLOG EVENTS Syntax

```
SHOW BINLOG EVENTS
  [IN 'log_name'] [FROM pos] [LIMIT [offset,] row_count]
```

Shows the events in the binary log. If you do not specify 'log\_name', the first binary log is displayed.

The `LIMIT` clause has the same syntax as for the `SELECT` statement. See [Section 13.2.7, “SELECT Syntax”](#).

This statement is available as of MySQL 4.0.

### 13.6.1.5. SHOW BINARY LOGS Syntax

```
SHOW BINARY LOGS
SHOW MASTER LOGS
```

Lists the binary log files on the server. This statement is used as part of the procedure described in [Section 13.6.1.1, “PURGE MASTER LOGS Syntax”](#), that shows how to determine which logs can be purged.

```
mysql> SHOW BINARY LOGS;
+-----+-----+
| Log_name      | File_size |
+-----+-----+
| binlog.000015 | 724935    |
| binlog.000016 | 733481    |
+-----+-----+
```

`SHOW MASTER LOGS` was added in MySQL 3.23.38. As of MySQL 4.1.1, you can also use `SHOW BINARY LOGS`, which is equivalent. The `File_size` column is displayed as of MySQL 5.0.7.

### 13.6.1.6. SHOW MASTER STATUS Syntax

```
SHOW MASTER STATUS
```

Provides status information about the binary log files of the master. Example:

```
mysql > SHOW MASTER STATUS;
+-----+-----+-----+-----+
| File          | Position | Binlog_Do_DB | Binlog_Ignore_DB |
+-----+-----+-----+-----+
| mysql-bin.003 | 73       | test         | manual,mysql      |
+-----+-----+-----+-----+
```

### 13.6.1.7. SHOW SLAVE HOSTS Syntax

```
SHOW SLAVE HOSTS
```

Displays a list of replication slaves currently registered with the master. Only slaves started with the `--report-host=slave_name` option are visible in this list.

The list is displayed on any server (not just the master server). The output looks like this:

```
mysql> SHOW SLAVE HOSTS;
+-----+-----+-----+-----+
| Server_id | Host      | Port | Master_id |
+-----+-----+-----+-----+
```

192168010	iconnect2	3306	192168011
1921680101	athena	3306	192168011

- **Server\_id**: The unique server ID of the slave server, as configured in the server's option file, or on the command line with `--server-id=value`.
- **Host**: The host name of the slave server, as configured in the server's option file, or on the command line with `--report-host=value`. Note that this can differ from the machine name as configured in the operating system.
- **Port**: The port the slave server is listening on.
- **Master\_id**: The unique server ID of the master server that the slave server is replicating from.

Some MySQL versions report another variable, `Rpl_recovery_rank`. This variable was never used, and was eventually removed.

## 13.6.2. SQL Statements for Controlling Slave Servers

Replication can be controlled through the SQL interface. This section discusses statements for managing slave replication servers. [Section 13.6.1, “SQL Statements for Controlling Master Servers”](#), discusses statements for managing master servers.

### 13.6.2.1. CHANGE MASTER TO Syntax

```
CHANGE MASTER TO master_def [, master_def] ...
```

```
master_def:
MASTER_HOST = 'host_name'
MASTER_USER = 'user_name'
MASTER_PASSWORD = 'password'
MASTER_PORT = port_num
MASTER_CONNECT_RETRY = count
MASTER_LOG_FILE = 'master_log_name'
MASTER_LOG_POS = master_log_pos
RELAY_LOG_FILE = 'relay_log_name'
RELAY_LOG_POS = relay_log_pos
MASTER_SSL = {0|1}
MASTER_SSL_CA = 'ca_file_name'
MASTER_SSL_CAPATH = 'ca_directory_name'
MASTER_SSL_CERT = 'cert_file_name'
MASTER_SSL_KEY = 'key_file_name'
MASTER_SSL_CIPHER = 'cipher_list'
```

`CHANGE MASTER TO` changes the parameters that the slave server uses for connecting to and communicating with the master server. It also updates the contents of the `master.info` and `relay-log.info` files.

`MASTER_USER`, `MASTER_PASSWORD`, `MASTER_SSL`, `MASTER_SSL_CA`, `MASTER_SSL_CAPATH`, `MASTER_SSL_CERT`, `MASTER_SSL_KEY`, and `MASTER_SSL_CIPHER` provide information to the slave about how to connect to its master.

The relay log options (`RELAY_LOG_FILE` and `RELAY_LOG_POS`) are available beginning with MySQL 4.0.

The SSL options (`MASTER_SSL`, `MASTER_SSL_CA`, `MASTER_SSL_CAPATH`, `MASTER_SSL_CERT`, `MASTER_SSL_KEY`, and `MASTER_SSL_CIPHER`) are available beginning with MySQL 4.1.1. You can change these options even on slaves that are compiled without SSL support. They are saved to the `master.info` file, but are ignored unless you use a server that has SSL



support enabled.

If you don't specify a given parameter, it keeps its old value, except as indicated in the following discussion. For example, if the password to connect to your MySQL master has changed, you just need to issue these statements to tell the slave about the new password:

```
STOP SLAVE; -- if replication was running
CHANGE MASTER TO MASTER_PASSWORD='new3cret';
START SLAVE; -- if you want to restart replication
```

There is no need to specify the parameters that do not change (host, port, user, and so forth).

`MASTER_HOST` and `MASTER_PORT` are the hostname (or IP address) of the master host and its TCP/IP port. Note that if `MASTER_HOST` is equal to `localhost`, then, like in other parts of MySQL, the port number might be ignored (if Unix socket files can be used, for example).

If you specify `MASTER_HOST` or `MASTER_PORT`, the slave assumes that the master server is different from before (even if you specify a host or port value that is the same as the current value.) In this case, the old values for the master binary log name and position are considered no longer applicable, so if you do not specify `MASTER_LOG_FILE` and `MASTER_LOG_POS` in the statement, `MASTER_LOG_FILE=''` and `MASTER_LOG_POS=4` are silently appended to it.

`MASTER_LOG_FILE` and `MASTER_LOG_POS` are the coordinates at which the slave I/O thread should begin reading from the master the next time the thread starts. If you specify either of them, you cannot specify `RELAY_LOG_FILE` or `RELAY_LOG_POS`. If neither of `MASTER_LOG_FILE` or `MASTER_LOG_POS` are specified, the slave uses the last coordinates of the *slave SQL thread* before `CHANGE MASTER` was issued. This ensures that replication has no discontinuity, even if the slave SQL thread was late compared to the slave I/O thread, when you just want to change, say, the password to use. This safe behavior was introduced starting from MySQL 4.0.17 and 4.1.1. (Before these versions, the coordinates used were the last coordinates of the slave I/O thread before `CHANGE MASTER` was issued. This caused the SQL thread to possibly lose some events from the master, thus breaking replication.)

`CHANGE MASTER` deletes all relay log files and starts a new one, unless you specify `RELAY_LOG_FILE` or `RELAY_LOG_POS`. In that case, relay logs are kept; as of MySQL 4.1.1, the `relay_log_purge` global variable is set silently to 0.

`CHANGE MASTER` is useful for setting up a slave when you have the snapshot of the master and have recorded the log and the offset corresponding to it. After loading the snapshot into the slave, you can run `CHANGE MASTER TO MASTER_LOG_FILE='log_name_on_master', MASTER_LOG_POS=log_offset_on_master` on the slave.

The following example changes the master and master's binary log coordinates. This is used when you want to set up the slave to replicate the master:

```
CHANGE MASTER TO
  MASTER_HOST='master2.mycompany.com',
  MASTER_USER='replication',
  MASTER_PASSWORD='big3cret',
  MASTER_PORT=3306,
  MASTER_LOG_FILE='master2-bin.001',
  MASTER_LOG_POS=4,
  MASTER_CONNECT_RETRY=10;
```

The next example shows an operation that is less frequently employed. It is used when the slave has relay logs that you want it to execute again for some reason. To do this, the master need not be reachable. You need only use `CHANGE MASTER TO` and start the SQL thread (`START SLAVE SQL_THREAD`):

```
CHANGE MASTER TO
  RELAY_LOG_FILE='slave-relay-bin.006',
  RELAY_LOG_POS=4025;
```

You can even use the second operation in a non-replication setup with a standalone, non-slave server for recovery following a crash. Suppose that your server has crashed and you have restored a backup. You want to replay the server's own binary logs (not relay logs, but regular binary logs), named (for example) `myhost-bin.*`. First, make a backup copy of these binary logs in some safe place, in case you don't exactly follow the procedure below and accidentally have the server purge the binary logs. If using MySQL 4.1.1 or newer, use `SET GLOBAL relay_log_purge=0` for additional safety. Then start the server without the `--log-bin` option. Before MySQL 4.0.19, start it with a new server ID; in newer versions there is no need; simply use the `--replicate-same-server-id` option. Start it with `--relay-log=myhost-bin` (to make the server believe that these regular binary logs are relay logs) and `--skip-slave-start` options. After the server starts, issue these statements:

```
CHANGE MASTER TO
  RELAY_LOG_FILE='myhost-bin.153',
  RELAY_LOG_POS=410,
  MASTER_HOST='some_dummy_string';
START SLAVE SQL_THREAD;
```

The server reads and executes its own binary logs, thus achieving crash recovery. Once the recovery is finished, run `STOP SLAVE`, shut down the server, delete the `master.info` and `relay-log.info` files, and restart the server with its original options.

Specifying the `MASTER_HOST` option (even with a dummy value) is required to make the server think it is a slave.

### 13.6.2.2. LOAD DATA FROM MASTER Syntax

```
LOAD DATA FROM MASTER
```

**This feature is deprecated. We recommend not using it anymore. It is subject to removal in a future version of MySQL.**

Since the current implementation of `LOAD DATA FROM MASTER` and `LOAD TABLE FROM MASTER` is very limited, these statements are deprecated in versions 4.1 of MySQL and above. We will introduce a more advanced technique (called “online backup”) in a future version. That technique will have the additional advantage of working with more storage engines.

For MySQL 5.1 and earlier, the recommended alternative solution to using `LOAD DATA FROM MASTER` or `LOAD TABLE FROM MASTER` is using `mysqldump` or `mysqlhotcopy`. The latter requires Perl and two Perl modules (`DBI` and `DBD:mysql`) and works for `MyISAM` and `ARCHIVE` tables only. With `mysqldump`, you can create SQL dumps on the master and pipe (or copy) these to a `mysql` client on the slave. This has the advantage of working for all storage engines, but can be quite slow, since it works using `SELECT`.

This statement takes a snapshot of the master and copies it to the slave. It updates the values of `MASTER_LOG_FILE` and `MASTER_LOG_POS` so that the slave starts replicating from the correct position. Any table and database exclusion rules specified with the `--replicate-*-do-*` and `--replicate-*-ignore-*` options are honored. `--replicate-rewrite-db` is *not* taken into account because a user could use this option to set up a non-unique mapping such as `--replicate-rewrite-db="db1->db3"` and `--replicate-rewrite-db="db2->db3"`, which would confuse the slave when loading tables from the master.

Use of this statement is subject to the following conditions:

- It works only for `MyISAM` tables. Attempting to load a non-`MyISAM` table results in the following error:

```
ERROR 1189 (08S01): Net error reading from master
```

- It acquires a global read lock on the master while taking the snapshot, which prevents updates on the master during the load operation.

If you are loading large tables, you might have to increase the values of `net_read_timeout` and `net_write_timeout` on both the master and slave servers. See [Section 5.2.2, “System Variables”](#).

Note that `LOAD DATA FROM MASTER` does *not* copy any tables from the `mysql` database. This makes it easy to have different users and privileges on the master and the slave.

To use `LOAD DATA FROM MASTER`, the replication account that is used to connect to the master must have the `RELOAD` and `SUPER` privileges on the master and the `SELECT` privilege for all master tables you want to load. All master tables for which the user does not have the `SELECT` privilege are ignored by `LOAD DATA FROM MASTER`. This is because the master hides them from the user: `LOAD DATA FROM MASTER` calls `SHOW DATABASES` to know the master databases to load, but `SHOW DATABASES` returns only databases for which the user has some privilege. See [Section 13.5.4.6, “SHOW DATABASES Syntax”](#). On the slave side, the user that issues `LOAD DATA FROM MASTER` must have privileges for dropping and creating the databases and tables that are copied.

### 13.6.2.3. `LOAD TABLE tbl_name FROM MASTER` Syntax

```
LOAD TABLE tbl_name FROM MASTER
```

**This feature is deprecated. We recommend not using it anymore. It is subject to removal in a future version of MySQL.**

Since the current implementation of `LOAD DATA FROM MASTER` and `LOAD TABLE FROM MASTER` is very limited, these statements are deprecated in versions 4.1 of MySQL and above. We will introduce a more advanced technique (called “online backup”) in a future version. That technique will have the additional advantage of working with more storage engines.

For MySQL 5.1 and earlier, the recommended alternative solution to using `LOAD DATA FROM MASTER` or `LOAD TABLE FROM MASTER` is using `mysqldump` or `mysqlhotcopy`. The latter requires Perl and two Perl modules (`DBI` and `DBD:mysql`) and works for `MyISAM` and `ARCHIVE` tables only. With `mysqldump`, you can create SQL dumps on the master and pipe (or copy) these to a `mysql` client on the slave. This has the advantage of working for all storage engines, but can be quite slow, since it works using `SELECT`.

Transfers a copy of the table from the master to the slave. This statement is implemented mainly debugging `LOAD DATA FROM MASTER` operations. To use `LOAD TABLE`, the account used for connecting to the master server must have the `RELOAD` and `SUPER` privileges on the master and the `SELECT` privilege for the master table to load. On the slave side, the user that issues `LOAD TABLE FROM MASTER` must have privileges for dropping and creating the table.

The conditions for `LOAD DATA FROM MASTER` apply here as well. For example, `LOAD TABLE FROM MASTER` works only for `MyISAM` tables. The timeout notes for `LOAD DATA FROM MASTER` apply as well.

### 13.6.2.4. `MASTER_POS_WAIT( )` Syntax

```
SELECT MASTER_POS_WAIT('master_log_file', master_log_pos)
```

This is actually a function, not a statement. It is used to ensure that the slave has read and executed events up to a given position in the master's binary log. See [Section 12.9.4, “Miscellaneous Functions”](#), for a full description.

### 13.6.2.5. `RESET SLAVE` Syntax

```
RESET SLAVE
```

`RESET SLAVE` makes the slave forget its replication position in the master's binary logs. This statement is meant to be used for a clean start: It deletes the `master.info` and `relay-log.info` files, all the relay logs, and starts a new relay log.

**Note:** All relay logs are deleted, even if they have not been completely executed by the slave SQL thread. (This is a condition likely to exist on a replication slave if you have issued a `STOP SLAVE` statement or if the slave is highly loaded.)

Connection information stored in the `master.info` file is immediately reset using any values specified in the corresponding startup options. This information includes values such as master host, master port, master user, and master password. If the slave SQL thread was in the middle of replicating temporary tables when it was stopped, and `RESET SLAVE` is issued, these replicated temporary tables are deleted on the slave.

This statement was named `FLUSH SLAVE` before MySQL 3.23.26.

### 13.6.2.6. SET GLOBAL SQL\_SLAVE\_SKIP\_COUNTER Syntax

```
SET GLOBAL SQL_SLAVE_SKIP_COUNTER = N
```

This statement skips the next `N` events from the master. This is useful for recovering from replication stops caused by a statement.

This statement is valid only when the slave thread is not running. Otherwise, it produces an error.

Before MySQL 4.0, omit the `GLOBAL` keyword from the statement.

### 13.6.2.7. SHOW SLAVE STATUS Syntax

```
SHOW SLAVE STATUS
```

This statement provides status information on essential parameters of the slave threads. If you issue this statement using the `mysql` client, you can use a `\G` statement terminator rather than a semicolon to obtain a more readable vertical layout:

```
mysql> SHOW SLAVE STATUS\G
***** 1. row *****
      Slave_IO_State: Waiting for master to send event
      Master_Host: localhost
      Master_User: root
      Master_Port: 3306
      Connect_Retry: 3
      Master_Log_File: gbichot-bin.005
      Read_Master_Log_Pos: 79
      Relay_Log_File: gbichot-relay-bin.005
      Relay_Log_Pos: 548
      Relay_Master_Log_File: gbichot-bin.005
      Slave_IO_Running: Yes
      Slave_SQL_Running: Yes
      Replicate_Do_DB:
      Replicate_Ignore_DB:
           Last_Errno: 0
           Last_Error:
           Skip_Counter: 0
      Exec_Master_Log_Pos: 79
      Relay_Log_Space: 552
      Until_Condition: None
      Until_Log_File:
      Until_Log_Pos: 0
      Master_SSL_Allowed: No
      Master_SSL_CA_File:
      Master_SSL_CA_Path:
```

```
Master_SSL_Cert:  
Master_SSL_Cipher:  
Master_SSL_Key:  
Seconds_Behind_Master: 8
```

Depending on your version of MySQL, you may not see all the fields just shown. In particular, several fields are present only as of MySQL 4.1.1.

`SHOW SLAVE STATUS` returns the following fields:

- `Slave_IO_State`

A copy of the `State` field of the output of `SHOW PROCESSLIST` for the slave I/O thread. This tells you what the thread is doing: trying to connect to the master, waiting for events from the master, reconnecting to the master, and so on. Possible states are listed in [Section 6.3, “Replication Implementation Details”](#). It is necessary to check this field for older versions of MySQL (prior to 4.1.14) because in these versions the thread could be running while unsuccessfully trying to connect to the master; only this field makes you aware of the connection problem. The state of the SQL thread is not copied because it is simpler. If it is running, there is no problem; if it is not, you can find the error in the `Last_Error` field (described below).

This field is present beginning with MySQL 4.1.1.

- `Master_Host`

The current master host.

- `Master_User`

The current user used to connect to the master.

- `Master_Port`

The current master port.

- `Connect_Retry`

The current value of the `--master-connect-retry` option.

- `Master_Log_File`

The name of the master binary log file from which the I/O thread is currently reading.

- `Read_Master_Log_Pos`

The position up to which the I/O thread has read in the current master binary log.

- `Relay_Log_File`

The name of the relay log file from which the SQL thread is currently reading and executing.

- `Relay_Log_Pos`

The position up to which the SQL thread has read and executed in the current relay log.

- `Relay_Master_Log_File`

The name of the master binary log file containing the most recent event executed by the SQL thread.

- `Slave_IO_Running`

Whether the I/O thread is started and has connected successfully to the master. For older versions of MySQL (prior to 4.1.14 and 5.0.12) `Slave_IO_Running` is `YES` if the I/O thread is started, even if the slave hasn't connected to the master yet.

- `Slave_SQL_Running`

Whether the SQL thread is started.

- `Replicate_Do_DB`, `Replicate_Ignore_DB`

The lists of databases that were specified with the `--replicate-do-db` and `--replicate-ignore-db` options, if any.

These fields are present beginning with MySQL 4.1.1.

- `Replicate_Do_Table`, `Replicate_Ignore_Table`, `Replicate_Wild_Do_Table`, `Replicate_Wild_Ignore_Table`

The lists of tables that were specified with the `--replicate-do-table`, `--replicate-ignore-table`, `--replicate-wild-do-table`, and `--replicate-wild-ignore-table` options, if any.

These fields are present beginning with MySQL 4.1.1.

- `Last_Errno`, `Last_Error`

The error number and error message returned by the most recently executed query. An error number of 0 and message of the empty string mean “no error.” If the `Last_Error` value is not empty, it also appears as a message in the slave's error log. For example:

```
Last_Errno: 1051
Last_Error: error 'Unknown table 'z'' on query 'drop table z'
```

The message indicates that the table `z` existed on the master and was dropped there, but it did not exist on the slave, so `DROP TABLE` failed on the slave. (This might occur, for example, if you forget to copy the table to the slave when setting up replication.)

- `Skip_Counter`

The most recently used value for `SQL_SLAVE_SKIP_COUNTER`.

- `Exec_Master_Log_Pos`

The position of the last event executed by the SQL thread from the master's binary log (`Relay_Master_Log_File`). (`Relay_Master_Log_File`, `Exec_Master_Log_Pos`) in the master's binary log corresponds to (`Relay_Log_File`, `Relay_Log_Pos`) in the relay log.

- `Relay_Log_Space`

The total combined size of all existing relay logs.

- `Until_Condition`, `Until_Log_File`, `Until_Log_Pos`

The values specified in the `UNTIL` clause of the `START SLAVE` statement.

`Until_Condition` has these values:

- `None` if no `UNTIL` clause was specified

- `Master` if the slave is reading until a given position in the master's binary logs
- `Relay` if the slave is reading until a given position in its relay logs

`Until_Log_File` and `Until_Log_Pos` indicate the log filename and position values that define the point at which the SQL thread stops executing.

These fields are present beginning with MySQL 4.1.1.

- `Master_SSL_Allowed`, `Master_SSL_CA_File`, `Master_SSL_CA_Path`, `Master_SSL_Cert`, `Master_SSL_Cipher`, `Master_SSL_Key`

These fields show the SSL parameters used by the slave to connect to the master, if any.

`Master_SSL_Allowed` has these values:

- `Yes` if an SSL connection to the master is permitted
- `No` if an SSL connection to the master is not permitted
- `Ignored` if an SSL connection is permitted but the slave server does not have SSL support enabled

The values of the other SSL-related fields correspond to the values of the `--master-ca`, `--master-capath`, `--master-cert`, `--master-cipher`, and `--master-key` options.

These fields are present beginning with MySQL 4.1.1.

- `Seconds_Behind_Master`

This field is present beginning with MySQL 4.1.1. It is been experimental and has been changed in MySQL 4.1.9. The following applies to slaves running MySQL 4.1.9 or newer. This field is an indication of how “late” the slave is:

- When the slave SQL thread is actively running (processing updates), this field is the number of seconds that have elapsed since the timestamp of the most recent event on the master executed by that thread.
- When the SQL thread thread has caught up to the slave I/O thread and goes idle waiting for more events from the I/O thread, this field is zero.

In essence, this field measures the time difference in seconds between the slave SQL thread and the slave I/O thread.

If the network connection between master and slave is fast, the slave I/O thread is very close to the master, so this field is a good approximation of how late the slave SQL thread is compared to the master. If the network is slow, this is *not* a good approximation; the slave SQL thread may quite often be caught up with the slow-reading slave I/O thread, so `Seconds_Behind_Master` often shows a value of 0, even if the I/O thread is late compared to the master. In other words, *this column is useful only for fast networks*.

This time difference computation works even though the master and slave do not have identical clocks (the clock difference is computed when the slave I/O thread starts, and assumed to remain constant from then on). `Seconds_Behind_Master` is `NULL` (which means “unknown”) if the slave SQL thread is not running, or if the slave I/O thread is not running or not connected to master. For example if the slave I/O thread is sleeping for the number of seconds given by the `--master-connect-retry` option before reconnecting, `NULL` is shown, as the slave cannot know what the master is doing, and so cannot say reliably how late it is.

This field has one limitation. The timestamp is preserved through replication, which means that, if a master M1 is itself a slave of M0, any event from M1's binlog which originates in replicating an event from M0's binlog has the timestamp of that event. This enables MySQL to replicate `TIMESTAMP` successfully. However, the drawback for `Seconds_Behind_Master` is that if M1 also receives direct updates from clients, the value randomly deviates, because sometimes the last M1's event is from M0 and sometimes it is the most recent timestamp from a direct update.

### 13.6.2.8. `START SLAVE` Syntax

```
START SLAVE [thread_type [, thread_type] ... ]
START SLAVE [SQL_THREAD] UNTIL
  MASTER_LOG_FILE = 'log_name', MASTER_LOG_POS = log_pos
START SLAVE [SQL_THREAD] UNTIL
  RELAY_LOG_FILE = 'log_name', RELAY_LOG_POS = log_pos

thread_type: IO_THREAD | SQL_THREAD
```

`START SLAVE` with no `thread_type` options starts both of the slave threads. The I/O thread reads queries from the master server and stores them in the relay log. The SQL thread reads the relay log and executes the queries. `START SLAVE` requires the `SUPER` privilege.

If `START SLAVE` succeeds in starting the slave threads, it returns without any error. However, even in that case, it might be that the slave threads start and then later stop (for example, because they do not manage to connect to the master or read its binary logs, or some other problem). `START SLAVE` does not warn you about this. You must check the slave's error log for error messages generated by the slave threads, or check that they are running satisfactorily with `SHOW SLAVE STATUS`.

As of MySQL 4.0.2, you can add `IO_THREAD` and `SQL_THREAD` options to the statement to name which of the threads to start.

As of MySQL 4.1.1, an `UNTIL` clause may be added to specify that the slave should start and run until the SQL thread reaches a given point in the master binary logs or in the slave relay logs. When the SQL thread reaches that point, it stops. If the `SQL_THREAD` option is specified in the statement, it starts only the SQL thread. Otherwise, it starts both slave threads. If the SQL thread is running, the `UNTIL` clause is ignored and a warning is issued.

For an `UNTIL` clause, you must specify both a log filename and position. Do not mix master and relay log options.

Any `UNTIL` condition is reset by a subsequent `STOP SLAVE` statement, a `START SLAVE` statement that includes no `UNTIL` clause, or a server restart.

The `UNTIL` clause can be useful for debugging replication, or to cause replication to proceed until just before the point where you want to avoid having the slave replicate a statement. For example, if an unwise `DROP TABLE` statement was executed on the master, you can use `UNTIL` to tell the slave to execute up to that point but no farther. To find what the event is, use `mysqlbinlog` with the master logs or slave relay logs, or by using a `SHOW BINLOG EVENTS` statement.

If you are using `UNTIL` to have the slave process replicated queries in sections, it is recommended that you start the slave with the `--skip-slave-start` option to prevent the SQL thread from running when the slave server starts. It is probably best to use this option in an option file rather than on the command line, so that an unexpected server restart does not cause it to be forgotten.

The `SHOW SLAVE STATUS` statement includes output fields that display the current values of the `UNTIL` condition.

This statement is called `SLAVE START` before MySQL 4.0.5. `SLAVE START` is still accepted for backward compatibility, but is now deprecated.



### 13.6.2.9. STOP SLAVE Syntax

```
STOP SLAVE [thread_type [, thread_type] ... ]  
thread_type: IO_THREAD | SQL_THREAD
```

Stops the slave threads. `STOP SLAVE` requires the `SUPER` privilege.

Like `START SLAVE`, as of MySQL 4.0.2, this statement may be used with the `IO_THREAD` and `SQL_THREAD` options to name the thread or threads to be stopped.

This statement is called `SLAVE STOP` before MySQL 4.0.5. `SLAVE STOP` is still accepted for backward compatibility, but is deprecated.

## 13.7. SQL Syntax for Prepared Statements

Support for server-side prepared statements was added in MySQL 4.1. This support takes advantage of the efficient client/server binary protocol, provided that you use an appropriate client programming interface. Candidate interfaces include the MySQL C API client library (for C programs), MySQL Connector/J (for Java programs), and MySQL Connector/NET. For example, the C API provides a set of function calls that make up its prepared statement API. See [Section 17.2.4, “C API Prepared Statements”](#). Other language interfaces can provide support for prepared statements that use the binary protocol by linking in the C client library, one example being the [mysql\\_i extension](#) [<http://php.net/mysqli>], available in PHP 5.0 and later.

Beginning with MySQL 4.1.3, an alternative SQL interface to prepared statements is available. This interface is not as efficient as using the binary protocol through a prepared statement API, but requires no programming because it is available directly at the SQL level:

- You can use it when no programming interface is available to you.
- You can use it from any program that allows you to send SQL statements to the server to be executed, such as the `mysql` client program.
- You can use it even if the client is using an old version of the client library. The only requirement is that you be able to connect to a server that is recent enough to support SQL syntax for prepared statements.

SQL syntax for prepared statements is intended to be used for situations such as these:

- You want to test how prepared statements work in your application before coding it.
- An application has problems executing prepared statements and you want to determine interactively what the problem is.
- You want to create a test case that describes a problem you are having with prepared statements, so that you can file a bug report.
- You need to use prepared statements but do not have access to a programming API that supports them.

SQL syntax for prepared statements is based on three SQL statements:

- `PREPARE stmt_name FROM preparable_stmt`

The `PREPARE` statement prepares a statement and assigns it a name, `stmt_name`, by which to refer to the statement later. Statement names are not case sensitive. `preparable_stmt` is either a string literal or a user variable that contains the text of the statement. The text must represent a single SQL statement, not multiple statements. Within the statement, ‘?’ characters can be used as parameter markers to indicate where data values are to be bound to the query later when you execute it. The ‘?’ characters should not be enclosed within quotes, even if you intend to bind them to string values. Parameter markers can be used only where data values should appear, not for SQL keywords, identifiers, and so forth.

If a prepared statement with the given name already exists, it is deallocated implicitly before the new statement is prepared. This means that if the new statement contains an error and cannot be prepared, an error is returned and no statement with the given name exists.

The scope of a prepared statement is the client session within which it is created. Other clients cannot see it.

- `EXECUTE stmt_name [USING @var_name [, @var_name] ...]`

After preparing a statement, you execute it with an `EXECUTE` statement that refers to the prepared statement name. If the prepared statement contains any parameter markers, you must supply a `USING` clause that lists user variables containing the values to be bound to the parameters. Parameter values can be supplied only by user variables, and the `USING` clause must name exactly as many variables as the number of parameter markers in the statement.

You can execute a given prepared statement multiple times, passing different variables to it or setting the variables to different values before each execution.

- `{DEALLOCATE | DROP} PREPARE stmt_name`

To deallocate a prepared statement, use the `DEALLOCATE PREPARE` statement. Attempting to execute a prepared statement after deallocating it results in an error.

If you terminate a client session without deallocating a previously prepared statement, the server deallocates it automatically.

The following SQL statements can be used in prepared statements: `CREATE TABLE`, `DELETE`, `DO`, `INSERT`, `REPLACE`, `SELECT`, `SET`, `UPDATE`, and most `SHOW` statements. Other statements are not yet supported.

The following examples show two equivalent ways of preparing a statement that computes the hypotenuse of a triangle given the lengths of the two sides.

The first example shows how to create a prepared statement by using a string literal to supply the text of the statement:

```
mysql> PREPARE stmt1 FROM 'SELECT SQRT(POW(?,2) + POW(?,2)) AS hypotenuse';
mysql> SET @a = 3;
mysql> SET @b = 4;
mysql> EXECUTE stmt1 USING @a, @b;
+-----+
| hypotenuse |
+-----+
|          5 |
+-----+
mysql> DEALLOCATE PREPARE stmt1;
```

The second example is similar, but supplies the text of the statement as a user variable:

```
mysql> SET @s = 'SELECT SQRT(POW(?,2) + POW(?,2)) AS hypotenuse';
mysql> PREPARE stmt2 FROM @s;
mysql> SET @a = 6;
mysql> SET @b = 8;
mysql> EXECUTE stmt2 USING @a, @b;
+-----+
| hypotenuse |
+-----+
|          10 |
+-----+
mysql> DEALLOCATE PREPARE stmt2;
```

SQL syntax for prepared statements cannot be used in nested fashion. That is, a statement passed to `PREPARE` cannot itself be a `PREPARE`, `EXECUTE`, or `DEALLOCATE PREPARE` statement.

SQL syntax for prepared statements is distinct from using prepared statement API calls. For example, you cannot use the `mysql_stmt_prepare()` C API function to prepare a `PREPARE`, `EXECUTE`, or `DEALLOCATE PREPARE` statement.

SQL syntax for prepared statements does not support multi-statements (that is, multiple statements within a single string separated by `;` characters).

---

# Chapter 14. Storage Engines and Table Types

MySQL supports several storage engines that act as handlers for different table types. MySQL storage engines include both those that handle transaction-safe tables and those that handle non-transaction-safe tables:

- The original storage engine was [ISAM](#), which managed non-transactional tables. This engine has been replaced by [MyISAM](#) and should no longer be used. It is deprecated in MySQL 4.1, and is removed in subsequent MySQL release series.
- In MySQL 3.23.0, the [MyISAM](#) and [HEAP](#) storage engines were introduced. [MyISAM](#) is an improved replacement for [ISAM](#). The [HEAP](#) storage engine provides in-memory tables. The [MERGE](#) storage engine was added in MySQL 3.23.25. It allows a collection of identical [MyISAM](#) tables to be handled as a single table. All three of these storage engines handle non-transactional tables, and all are included in MySQL by default. Note that the [HEAP](#) storage engine has been renamed the [MEMORY](#) engine.
- The [InnoDB](#) and [BDB](#) storage engines that handle transaction-safe tables were introduced in later versions of MySQL 3.23. Both are available in source distributions as of MySQL 3.23.34a. [BDB](#) is included in MySQL-Max binary distributions on those operating systems that support it. [InnoDB](#) also is included in MySQL-Max binary distributions for MySQL 3.23. Beginning with MySQL 4.0, [InnoDB](#) is included by default in all MySQL binary distributions. In source distributions, you can enable or disable either engine by configuring MySQL as you like.
- The [EXAMPLE](#) storage engine was added in MySQL 4.1.3. It is a “stub” engine that does nothing. You can create tables with this engine, but no data can be stored in them or retrieved from them. The purpose of this engine is to serve as an example in the MySQL source code that illustrates how to begin writing new storage engines. As such, it is primarily of interest to developers.
- [NDB Cluster](#) is the storage engine used by MySQL Cluster to implement tables that are partitioned over many computers. It is available in source code distributions as of MySQL 4.1.2 and binary distributions as of MySQL-Max 4.1.3.

MySQL Cluster is covered in a separate chapter of this Manual. See [Chapter 15, MySQL Cluster](#), for more information.

- The [ARCHIVE](#) storage engine was added in MySQL 4.1.3. It is used for storing large amounts of data without indexes in a very small footprint.
- The [CSV](#) storage engine was added in MySQL 4.1.4. This engine stores data in text files using comma-separated values format.
- The [BLACKHOLE](#) storage engine was added in MySQL 4.1.11. This engine accepts but does not store data and retrievals always return an empty set.

This chapter describes each of the MySQL storage engines except for [NDB Cluster](#), which is covered in [Chapter 15, MySQL Cluster](#).

When you create a new table, you can specify which storage engine to use by adding an [ENGINE](#) or [TYPE](#) table option to the [CREATE TABLE](#) statement:

```
CREATE TABLE t (i INT) ENGINE = INNODB;  
CREATE TABLE t (i INT) TYPE = MEMORY;
```

[ENGINE](#) is the preferred term, but cannot be used before MySQL 4.0.18. [TYPE](#) is available beginning

with MySQL 3.23.0, the first version of MySQL for which multiple storage engines were available. `TYPE` is supported for backward compatibility but is deprecated.

If you omit the `ENGINE` or `TYPE` option, the default storage engine is used. Normally, this is `MyISAM`, but you can change it by using the `--default-storage-engine` or `--default-table-type` server startup option, or by setting the `default-storage-engine` or `default-table-type` option in the `my.cnf` configuration file.

You can set the default storage engine to be used during the current session by setting the `storage_engine` or `table_type` variable:

```
SET storage_engine=MYISAM;  
SET table_type=BDB;
```

When MySQL is installed on Windows using the MySQL Configuration Wizard, the `InnoDB` storage engine can be selected as the default instead of `MyISAM`. See [Section 2.3.4.6, “The Database Usage Dialog”](#).

To convert a table from one storage engine to another, use an `ALTER TABLE` statement that indicates the new engine:

```
ALTER TABLE t ENGINE = MYISAM;  
ALTER TABLE t TYPE = BDB;
```

See [Section 13.1.5, “CREATE TABLE Syntax”](#), and [Section 13.1.2, “ALTER TABLE Syntax”](#).

If you try to use a storage engine that is not compiled in or that is compiled in but deactivated, MySQL instead creates a table using the default storage engine, usually `MyISAM`. (Before MySQL, `MyISAM` is always used for unavailable storage engines.) type `MyISAM`. This behavior is convenient when you want to copy tables between MySQL servers that support different storage engines. (For example, in a replication setup, perhaps your master server supports transactional storage engines for increased safety, but the slave servers use only non-transactional storage engines for greater speed.)

This automatic substitution of the default storage engine for unavailable engines can be confusing for new MySQL users. In MySQL 4.1, a warning is generated when a storage engine is automatically changed.

For new tables, MySQL always creates an `.frm` file to hold the table and column definitions. The table's index and data may be stored in one or more other files, depending on the storage engine. The server creates the `.frm` file above the storage engine level. Individual storage engines create any additional files required for the tables that they manage.

A database may contain tables of different types. That is, tables need not all be created with the same storage engine.

Transaction-safe tables (TSTs) have several advantages over non-transaction-safe tables (NTSTs):

- They are safer. Even if MySQL crashes or you get hardware problems, you can get your data back, either by automatic recovery or from a backup plus the transaction log.
- You can combine many statements and accept them all at the same time with the `COMMIT` statement (if autocommit is disabled).
- You can execute `ROLLBACK` to ignore your changes (if autocommit is disabled).
- If an update fails, all of your changes are reverted. (With non-transaction-safe tables, all changes that have taken place are permanent.)
- Transaction-safe storage engines can provide better concurrency for tables that get many updates

concurrently with reads.

You can combine transaction-safe and non-transaction-safe tables in the same statements to get the best of both worlds. However, although MySQL supports several transaction-safe storage engines, for best results, you should not mix different storage engines within a transaction with autocommit disabled. For example, if you do this, changes to non-transaction-safe tables still are committed immediately and cannot be rolled back. For information about this and other problems that can occur in transactions that use mixed storage engines, see [Section 13.4.1, “START TRANSACTION, COMMIT, and ROLLBACK Syntax”](#).

Note that to use the `InnoDB` storage engine in MySQL 3.23, you must configure at least the `innodb_data_file_path` startup option. In 4.0 and up, `InnoDB` uses default configuration values if you specify none. See [Section 14.2.4, “InnoDB Configuration”](#).

Non-transaction-safe tables have several advantages of their own, all of which occur because there is no transaction overhead:

- Much faster
- Lower disk space requirements
- Less memory required to perform updates

## 14.1. The `MyISAM` Storage Engine

`MyISAM` is the default storage engine as of MySQL 3.23. It is based on the `ISAM` code but has many useful extensions.

Each `MyISAM` table is stored on disk in three files. The files have names that begin with the table name and have an extension to indicate the file type. An `.frm` file stores the table format. The data file has an `.MYD` (`MYData`) extension. The index file has an `.MYI` (`MYIndex`) extension.

To specify explicitly that you want a `MyISAM` table, indicate that with an `ENGINE` table option:

```
CREATE TABLE t (i INT) ENGINE = MYISAM;
```

The older term `TYPE` is supported as a synonym for `ENGINE` for backward compatibility, but `ENGINE` is the preferred term from MySQL 4.0.18 on and `TYPE` is deprecated.

Normally, the `ENGINE` or `TYPE` option is unnecessary; `MyISAM` is the default storage engine unless the default has been changed. To ensure that `MyISAM` is used in situations where the default might have been changed, specify the storage engine explicitly.

You can check or repair `MyISAM` tables with the `mysqlcheck` client or `myisamchk` utility. You can also compress `MyISAM` tables with `myisampack` to take up much less space. See [Section 8.11, “mysqlcheck — A Table Maintenance and Repair Program”](#), [Section 5.9.4.1, “Using myisamchk for Crash Recovery”](#), and [Section 8.5, “myisampack — Generate Compressed, Read-Only MyISAM Tables”](#).

The following characteristics of the `MyISAM` storage engine are improvements over the older `ISAM` engine:

- All data values are stored with the low byte first. This makes the data machine and operating system independent. The only requirements for binary portability are that the machine uses

two's-complement signed integers and IEEE floating-point format. These requirements are widely used among mainstream machines. Binary compatibility might not be applicable to embedded systems, which sometimes have peculiar processors.

There is no significant speed penalty for storing data low byte first; the bytes in a table row normally are unaligned and it takes little more processing to read an unaligned byte in order than in reverse order. Also, the code in the server that fetches column values is not time critical compared to other code.

- All numeric key values are stored with the high byte first to allow better index compression.
- Large files (up to 63-bit file length) are supported on filesystems and operating systems that support large files.
- The maximum number of indexes per table is 64 (32 before MySQL 4.1.2). This can be changed by changing the source and recompiling. The maximum number of columns per index is 16.
- The maximum key length is 1000 bytes (500 before MySQL 4.1.2). This can be changed by changing the source and recompiling. For the case of a key longer than 250 bytes, a larger key block size than the default of 1024 bytes is used.
- Index files are usually much smaller with `MyISAM` than with `ISAM`. This means that `MyISAM` normally uses less system resources than `ISAM`, but needs more CPU time when inserting data into a compressed index.
- When rows are inserted in sorted order (as when you are using an `AUTO_INCREMENT` column), the index tree is split so that the high node only contains one key. This improves space utilization in the index tree.
- Internal handling of one `AUTO_INCREMENT` column per table is supported. `MyISAM` automatically updates this column for `INSERT/UPDATE`. This makes `AUTO_INCREMENT` columns faster (at least 10%). Values at the top of the sequence are not reused after being deleted as they are with `ISAM`. (When an `AUTO_INCREMENT` column is defined as the last column of a multiple-column index, reuse of values deleted from the top of a sequence does occur.) The `AUTO_INCREMENT` value can be reset with `ALTER TABLE` or `myisamchk`.
- Dynamic-sized rows are much less fragmented when mixing deletes with updates and inserts. This is done by automatically combining adjacent deleted blocks and by extending blocks if the next block is deleted.
- If a table has no free blocks in the middle of the data file, you can `INSERT` new rows into it at the same time that other threads are reading from the table. (These are known as concurrent inserts.) A free block can occur as a result of deleting rows or an update of a dynamic length row with more data than its current contents. When all free blocks are used up (filled in), future inserts become concurrent again. See [Section 7.3.3, “Concurrent Inserts”](#).
- You can put the data file and index file on different directories to get more speed with the `DATA DIRECTORY` and `INDEX DIRECTORY` table options to `CREATE TABLE`. See [Section 13.1.5, “CREATE TABLE Syntax”](#).
- `BLOB` and `TEXT` columns can be indexed.
- `NULL` values are allowed in indexed columns. This takes 0-1 bytes per key.
- As of MySQL 4.1, each character column can have a different character set.
- There is a flag in the `MyISAM` index file that indicates whether the table was closed correctly. If `mysqld` is started with the `--myisam-recover` option, `MyISAM` tables are automatically checked when opened, and are repaired if the table wasn't closed properly.

- `myisamchk` marks tables as checked if you run it with the `--update-state` option. `myisamchk --fast` checks only those tables that don't have this mark.
- `myisamchk --analyze` stores statistics for portions of keys, not only for whole keys as in `ISAM`.
- `myisampack` can pack `BLOB` and `VARCHAR` columns; `pack_isam` cannot.

`MyISAM` also supports the following features, which MySQL will be able to use in the near future:

- Support for a true `VARCHAR` type; a `VARCHAR` column starts with a length stored in one or two bytes.
- Tables with `VARCHAR` columns may have fixed or dynamic row length.
- The sum of the lengths of the `VARCHAR` and `CHAR` columns in a table may be up to 64KB.
- A hashed computed index can be used for `UNIQUE`. This allows you to have `UNIQUE` on any combination of columns in a table. (However, you cannot search on a `UNIQUE` computed index.)

#### Additional resources

- A forum dedicated to the `MyISAM` storage engine is available at <http://forums.mysql.com/list.php?21>.

### 14.1.1. `MyISAM` Startup Options

The following options to `mysqld` can be used to change the behavior of `MyISAM` tables. For additional information, see [Section 5.2.1, “Command Options”](#).

- `--myisam-recover=mode`

Set the mode for automatic recovery of crashed `MyISAM` tables.

- `--delay-key-write=ALL`

Don't flush key buffers between writes for any `MyISAM` table.

**Note:** If you do this, you should not access `MyISAM` tables from another program (such as from another MySQL server or with `myisamchk`) when the tables are in use. Doing so risks index corruption. Using `--external-locking` does not eliminate this risk.

The following system variables affect the behavior of `MyISAM` tables. For additional information, see [Section 5.2.2, “System Variables”](#).

- `bulk_insert_buffer_size`

The size of the tree cache used in bulk insert optimization. **Note:** This is a limit *per thread*!

- `myisam_max_extra_sort_file_size`



Used to help MySQL to decide when to use the slow but safe key cache index creation method.

**Note:** This parameter is given in megabytes before MySQL 4.0.3, and in bytes as of 4.0.3.

- `myisam_max_sort_file_size`

The maximum size of the temporary file that MySQL is allowed to use while re-creating a **MyISAM** index (during `REPAIR TABLE`, `ALTER TABLE`, or `LOAD DATA INFILE`). If the file size would be larger than this value, the index is created using the key cache instead, which is slower. This variable was added in MySQL 3.23.37. **Note:** The value is given in megabytes before 4.0.3 and in bytes thereafter.

- `myisam_sort_buffer_size`

Set the size of the buffer used when recovering tables.

Automatic recovery is activated if you start `mysqld` with the `--myisam-recover` option. In this case, when the server opens a **MyISAM** table, it checks whether the table is marked as crashed or whether the open count variable for the table is not 0 and you are running the server with external locking disabled. If either of these conditions is true, the following happens:

- The server checks the table for errors.
- If the server finds an error, it tries to do a fast table repair (with sorting and without re-creating the data file).
- If the repair fails because of an error in the data file (for example, a duplicate-key error), the server tries again, this time re-creating the data file.
- If the repair still fails, the server tries once more with the old repair option method (write row by row without sorting). This method should be able to repair any type of error and has low disk space requirements.

If the recovery wouldn't be able to recover all rows from previously completed statements and you didn't specify `FORCE` in the value of the `--myisam-recover` option, automatic repair aborts with an error message in the error log:

```
Error: Couldn't repair table: test.g00pages
```

If you specify `FORCE`, a warning like this is written instead:

```
Warning: Found 344 of 354 rows when repairing ./test/g00pages
```

Note that if the automatic recovery value includes `BACKUP`, the recovery process creates files with names of the form `tbl_name-datetime.BAK`. You should have a `cron` script that automatically moves these files from the database directories to backup media.

## 14.1.2. Space Needed for Keys

**MyISAM** tables use B-tree indexes. You can roughly calculate the size for the index file as  $(key\_length+4)/0.67$ , summed over all keys. This is for the worst case when all keys are inserted in sorted order and the table doesn't have any compressed keys.

String indexes are space compressed. If the first index part is a string, it is also prefix compressed. Space compression makes the index file smaller than the worst-case figure if a string column has a lot of trail-

ing space or is a `VARCHAR` column that is not always used to the full length. Prefix compression is used on keys that start with a string. Prefix compression helps if there are many strings with an identical prefix.

In `MyISAM` tables, you can also prefix compress numbers by specifying the `PACK_KEYS=1` table option when you create the table. Numbers are stored with the high byte first, so this helps when you have many integer keys that have an identical prefix.

### 14.1.3. `MyISAM` Table Storage Formats

`MyISAM` supports three different storage formats. Two of them, fixed and dynamic format, are chosen automatically depending on the type of columns you are using. The third, compressed format, can be created only with the `myisampack` utility.

When you use `CREATE TABLE` or `ALTER TABLE` for a table that has no `BLOB` or `TEXT` columns, you can force the table format to `FIXED` or `DYNAMIC` with the `ROW_FORMAT` table option. This causes `CHAR` and `VARCHAR` columns to become `CHAR` for `FIXED` format, or `VARCHAR` for `DYNAMIC` format. See [Section 13.1.5, “CREATE TABLE Syntax”](#), for information about `ROW_FORMAT`.

#### 14.1.3.1. Static (Fixed-Length) Table Characteristics

Static format is the default for `MyISAM` tables. It is used when the table contains no variable-length columns (`VARCHAR`, `VARBINARY`, `BLOB`, or `TEXT`). Each row is stored using a fixed number of bytes.

Of the three `MyISAM` storage formats, static format is the simplest and most secure (least subject to corruption). It is also the fastest of the on-disk formats due to the ease with which rows in the data file can be found on disk: To look up a row based on a row number in the index, multiply the row number by the row length to calculate the row position. Also, when scanning a table, it is very easy to read a constant number of rows with each disk read operation.

The security is evidenced if your computer crashes while the MySQL server is writing to a fixed-format `MyISAM` file. In this case, `myisamchk` can easily determine where each row starts and ends, so it can usually reclaim all rows except the partially written one. Note that `MyISAM` table indexes can always be reconstructed based on the data rows.

Static-format tables have these characteristics:

- `CHAR` and `BINARY` columns are space-padded to the column width. This is also true for `NUMERIC` and `DECIMAL` columns.
- Very quick.
- Easy to cache.
- Easy to reconstruct after a crash, because rows are located in fixed positions.
- Reorganization is unnecessary unless you delete a huge number of rows and want to return free disk space to the operating system. To do this, use `OPTIMIZE TABLE` or `myisamchk -r`.
- Usually require more disk space than dynamic-format tables.

#### 14.1.3.2. Dynamic Table Characteristics

Dynamic storage format is used if a `MyISAM` table contains any variable-length columns (`VARCHAR`, `VARBINARY`, `BLOB`, or `TEXT`), or if the table was created with the `ROW_FORMAT=DYNAMIC` table option.

Dynamic format is a little more complex than static format because each row has a header that indicates how long it is. A row can become fragmented (stored in non-contiguous pieces) when it is made longer as a result of an update.

You can use `OPTIMIZE TABLE` or `myisamchk -r` to defragment a table. If you have fixed-length columns that you access or change frequently in a table that also contains some variable-length columns, it might be a good idea to move the variable-length columns to other tables just to avoid fragmentation.

Dynamic-format tables have these characteristics:

- All string columns are dynamic except those with a length less than four.
- Each row is preceded by a bitmap that indicates which columns contain the empty string (for string columns) or zero (for numeric columns). Note that this does not include columns that contain `NULL` values. If a string column has a length of zero after trailing space removal, or a numeric column has a value of zero, it is marked in the bitmap and not saved to disk. Non-empty strings are saved as a length byte plus the string contents.
- Much less disk space usually is required than for fixed-length tables.
- Each row uses only as much space as is required. However, if a row becomes larger, it is split into as many pieces as are required, resulting in row fragmentation. For example, if you update a row with information that extends the row length, the row becomes fragmented. In this case, you may have to run `OPTIMIZE TABLE` or `myisamchk -r` from time to time to improve performance. Use `myisamchk -ei` to obtain table statistics.
- More difficult than static-format tables to reconstruct after a crash, because rows may be fragmented into many pieces and links (fragments) may be missing.
- The expected row length for dynamic-sized rows is calculated using the following expression:

```

3
+ (number of columns + 7) / 8
+ (number of char columns)
+ (packed size of numeric columns)
+ (length of strings)
+ (number of NULL columns + 7) / 8

```

There is a penalty of 6 bytes for each link. A dynamic row is linked whenever an update causes an enlargement of the row. Each new link is at least 20 bytes, so the next enlargement probably goes in the same link. If not, another link is created. You can find the number of links using `myisamchk -ed`. All links may be removed with `OPTIMIZE TABLE` or `myisamchk -r`.

### 14.1.3.3. Compressed Table Characteristics

Compressed storage format is a read-only format that is generated with the `myisampack` tool.

All MySQL distributions as of version 3.23.19 include `myisampack` by default. (This version is when MySQL was placed under the GPL.) For earlier versions, `myisampack` was included only with licenses or support agreements, but the server still can read tables that were compressed with `myisampack`. Compressed tables can be uncompressed with `myisamchk`. (For the `ISAM` storage engine, compressed tables can be created with `pack_isam` and uncompressed with `isamchk`.)

Compressed tables have the following characteristics:

- Compressed tables take very little disk space. This minimizes disk usage, which is helpful when using slow disks (such as CD-ROMs).

- Each row is compressed separately, so there is very little access overhead. The header for a row takes up one to three bytes depending on the biggest row in the table. Each column is compressed differently. There is usually a different Huffman tree for each column. Some of the compression types are:
  - Suffix space compression.
  - Prefix space compression.
  - Numbers with a value of zero are stored using one bit.
  - If values in an integer column have a small range, the column is stored using the smallest possible type. For example, a `BIGINT` column (eight bytes) can be stored as a `TINYINT` column (one byte) if all its values are in the range from `-128` to `127`.
  - If a column has only a small set of possible values, the data type is converted to `ENUM`.
  - A column may use any combination of the preceding compression types.
- Can be used for fixed-length or dynamic-length rows.

**Note.** While a compressed table is read-only, and you cannot therefore update or add rows in the table, DDL (Data Definition Language) operations are still valid. For example, you may still use `DROP` to drop the table, and `TRUNCATE` to empty the table.

## 14.1.4. `MyISAM` Table Problems

The file format that MySQL uses to store data has been extensively tested, but there are always circumstances that may cause database tables to become corrupted. The following discussion describes how this can happen and how to handle it.

### 14.1.4.1. Corrupted `MyISAM` Tables

Even though the `MyISAM` table format is very reliable (all changes to a table made by an SQL statement are written before the statement returns), you can still get corrupted tables if any of the following events occur:

- The `mysqld` process is killed in the middle of a write.
- An unexpected computer shutdown occurs (for example, the computer is turned off).
- Hardware failures.
- You are using an external program (such as `myisamchk`) to modify a table that is being modified by the server at the same time.
- A software bug in the MySQL or `MyISAM` code.

Typical symptoms of a corrupt table are:

- You get the following error while selecting data from the table:

```
Incorrect key file for table: '...'. Try to repair it
```

- Queries don't find rows in the table or return incomplete results.

You can check the health of a `MyISAM` table using the `CHECK TABLE` statement, and repair a corrupted `MyISAM` table with `REPAIR TABLE`. When `mysqld` is not running, you can also check or repair a table with the `myisamchk` command. See [Section 13.5.2.3, “CHECK TABLE Syntax”](#), [Section 13.5.2.6, “REPAIR TABLE Syntax”](#), and [Section 8.3, “myisamchk — MyISAM Table-Maintenance Utility”](#).

If your tables become corrupted frequently, you should try to determine why this is happening. The most important thing to know is whether the table became corrupted as a result of a server crash. You can verify this easily by looking for a recent `restarted mysqld` message in the error log. If there is such a message, it is likely that table corruption is a result of the server dying. Otherwise, corruption may have occurred during normal operation. This is a bug. You should try to create a reproducible test case that demonstrates the problem. See [Section A.4.2, “What to Do If MySQL Keeps Crashing”](#), and [Section E.1.6, “Making a Test Case If You Experience Table Corruption”](#).

### 14.1.4.2. Problems from Tables Not Being Closed Properly

Each `MyISAM` index file (`.MYI` file) has a counter in the header that can be used to check whether a table has been closed properly. If you get the following warning from `CHECK TABLE` or `myisamchk`, it means that this counter has gone out of sync:

```
clients are using or haven't closed the table properly
```

This warning doesn't necessarily mean that the table is corrupted, but you should at least check the table.

The counter works as follows:

- The first time a table is updated in MySQL, a counter in the header of the index files is incremented.
- The counter is not changed during further updates.
- When the last instance of a table is closed (because a `FLUSH TABLES` operation was performed or because there is no room in the table cache), the counter is decremented if the table has been updated at any point.
- When you repair the table or check the table and it is found to be okay, the counter is reset to zero.
- To avoid problems with interaction with other processes that might check the table, the counter is not decremented on close if it was zero.

In other words, the counter can become incorrect only under these conditions:

- A `MyISAM` table is copied without first issuing `LOCK TABLES` and `FLUSH TABLES`.
- MySQL has crashed between an update and the final close. (Note that the table may still be okay, because MySQL always issues writes for everything between each statement.)
- A table was modified by `myisamchk --recover` or `myisamchk --update-state` at the same time that it was in use by `mysqld`.
- Multiple `mysqld` servers are using the table and one server performed a `REPAIR TABLE` or `CHECK TABLE` on the table while it was in use by another server. In this setup, it is safe to use `CHECK TABLE`, although you might get the warning from other servers. However, `REPAIR TA-`

`BLE` should be avoided because when one server replaces the data file with a new one, this is not known to the other servers.

In general, it is a bad idea to share a data directory among multiple servers. See [Section 5.12](#), “Running Multiple MySQL Servers on the Same Machine”, for additional discussion.

## 14.2. The `InnoDB` Storage Engine

### 14.2.1. `InnoDB` Overview

`InnoDB` provides MySQL with a transaction-safe (`ACID` compliant) storage engine that has commit, rollback, and crash recovery capabilities. `InnoDB` does locking on the row level and also provides an Oracle-style consistent non-locking read in `SELECT` statements. These features increase multi-user concurrency and performance. There is no need for lock escalation in `InnoDB` because row-level locks fit in very little space. `InnoDB` also supports `FOREIGN KEY` constraints. You can freely mix `InnoDB` tables with tables from other MySQL storage engines, even within the same statement.

`InnoDB` has been designed for maximum performance when processing large data volumes. Its CPU efficiency is probably not matched by any other disk-based relational database engine.

Fully integrated with MySQL Server, the `InnoDB` storage engine maintains its own buffer pool for caching data and indexes in main memory. `InnoDB` stores its tables and indexes in a tablespace, which may consist of several files (or raw disk partitions). This is different from, for example, `MyISAM` tables where each table is stored using separate files. `InnoDB` tables can be of any size even on operating systems where file size is limited to 2GB.

`InnoDB` is included in binary distributions by default as of MySQL 4.0. For information about `InnoDB` support in MySQL 3.23, see [Section 14.2.3](#), “`InnoDB` in MySQL 3.23”. Starting from MySQL 4.1.5, the improved Windows installer makes `InnoDB` the MySQL default storage engine on Windows.

`InnoDB` is used in production at numerous large database sites requiring high performance. The famous Internet news site Slashdot.org runs on `InnoDB`. Myrix, Inc. stores over 1TB of data in `InnoDB`, and another site handles an average load of 800 inserts/updates per second in `InnoDB`.

`InnoDB` is published under the same GNU GPL License Version 2 (of June 1991) as MySQL. For more information on MySQL licensing, see <http://www.mysql.com/company/legal/licensing/>.

A forum dedicated to the `InnoDB` storage engine is available at <http://forums.mysql.com/list.php?22>.

### 14.2.2. `InnoDB` Contact Information

Contact information for Innobase Oy, producer of the `InnoDB` engine:

```
Web site: http://www.innodb.com/
Email: <sales@innodb.com>
Phone: +358-9-6969 3250 (office)
       +358-40-5617367 (mobile)
```

```
Innbase Oy Inc.
World Trade Center Helsinki
Aleksanterinkatu 17
P.O.Box 800
00101 Helsinki
Finland
```

### 14.2.3. `InnoDB` in MySQL 3.23

Beginning with MySQL 4.0, `InnoDB` is enabled by default, so the following information applies only to MySQL 3.23.

`InnoDB` tables are included in the MySQL source distribution starting from 3.23.34a and are activated in the MySQL-Max binaries of the 3.23 series. For Windows, the MySQL-Max binaries are included in the standard distribution.

If you have downloaded a binary version of MySQL that includes support for `InnoDB`, simply follow the instructions of the MySQL manual for installing a binary version of MySQL. If you have MySQL 3.23 installed, the simplest way to install MySQL-Max is to replace the executable `mysqld` server with the corresponding executable from the MySQL-Max distribution. MySQL and MySQL-Max differ only in the server executable. See [Section 2.8, “Installing MySQL on Other Unix-Like Systems”](#), and [Section 5.3, “The `mysqld-max` Extended MySQL Server”](#).

To compile the MySQL source code with `InnoDB` support, download MySQL 3.23.34a or newer from <http://www.mysql.com/> and configure MySQL with the `--with-innodb` option. See [Section 2.9, “MySQL Installation Using a Source Distribution”](#).

To use `InnoDB` tables with MySQL 3.23, you must specify configuration parameters in the `[mysqld]` section of the `my.cnf` option file. On Windows, you can use `my.ini` instead. If you do not configure `InnoDB` in the option file, `InnoDB` does not start. (From MySQL 4.0 on, `InnoDB` uses default parameters if you do not specify any. However, to get best performance, it is still recommended that you use parameters appropriate for your system, as discussed in [Section 14.2.4, “InnoDB Configuration”](#).)

In MySQL 3.23, you must specify at the minimum an `innodb_data_file_path` value to configure the `InnoDB` data files. For example, to configure `InnoDB` to use a single 500MB data file, place the following setting in the `[mysqld]` section of your option file:

```
[mysqld]
innodb_data_file_path=ibdata1:500M
```

`InnoDB` creates the `ibdata1` file in the MySQL data directory by default. To specify the location explicitly, specify an `innodb_data_home_dir` setting. See [Section 14.2.4, “InnoDB Configuration”](#).

## 14.2.4. InnoDB Configuration

To enable `InnoDB` tables in MySQL 3.23, see [Section 14.2.3, “InnoDB in MySQL 3.23”](#).

From MySQL 4.0 on, the `InnoDB` storage engine is enabled by default. If you don't want to use `InnoDB` tables, you can add the `skip-innodb` option to your MySQL option file.

**Note:** `InnoDB` provides MySQL with a transaction-safe (`ACID` compliant) storage engine that has commit, rollback, and crash recovery capabilities. **However, it cannot do so** if the underlying operating system or hardware does not work as advertised. Many operating systems or disk subsystems may delay or reorder write operations to improve performance. On some operating systems, the very system call that should wait until all unwritten data for a file has been flushed — `fsync()` — might actually return before the data has been flushed to stable storage. Because of this, an operating system crash or a power outage may destroy recently committed data, or in the worst case, even corrupt the database because of write operations having been reordered. If data integrity is important to you, you should perform some “pull-the-plug” tests before using anything in production. On Mac OS X 10.3 and up, `InnoDB` uses a special `fcntl()` file flush method. Under Linux, it is advisable to **disable the write-back cache**.

On ATAPI hard disks, a command such `hdparm -W0 /dev/hda` may work to disable the write-back cache. **Beware that some drives or disk controllers may be unable to disable the write-back cache.**

Two important disk-based resources managed by the `InnoDB` storage engine are its tablespace data files and its log files.

**Note:** If you specify no [InnoDB](#) configuration options, MySQL 4.0 and above create an auto-extending 10MB data file named `ibdata1` and two 5MB log files named `ib_logfile0` and `ib_logfile1` in the MySQL data directory. (In MySQL 4.0.0 and 4.0.1, the data file is 64MB and not auto-extending.) In MySQL 3.23, [InnoDB](#) does not start if you provide no configuration options. To get good performance, you should explicitly provide [InnoDB](#) parameters as discussed in the following examples. Naturally, you should edit the settings to suit your hardware and requirements.

The examples shown here are representative. See [Section 14.2.5, “InnoDB Startup Options and System Variables”](#) for additional information about [InnoDB](#)-related configuration parameters.

To set up the [InnoDB](#) tablespace files, use the `innodb_data_file_path` option in the `[mysqld]` section of the `my.cnf` option file. On Windows, you can use `my.ini` instead. The value of `innodb_data_file_path` should be a list of one or more data file specifications. If you name more than one data file, separate them by semicolon (`;`) characters:

```
innodb_data_file_path=datafile_spec1[:datafile_spec2]...
```

For example, a setting that explicitly creates a tablespace having the same characteristics as the MySQL 4.0 default is as follows:

```
[mysqld]
innodb_data_file_path=ibdata1:10M:autoextend
```

This setting configures a single 10MB data file named `ibdata1` that is auto-extending. No location for the file is given, so by default, [InnoDB](#) creates it in the MySQL data directory.

Sizes are specified using `M` or `G` suffix letters to indicate units of MB or GB.

A tablespace containing a fixed-size 50MB data file named `ibdata1` and a 50MB auto-extending file named `ibdata2` in the data directory can be configured like this:

```
[mysqld]
innodb_data_file_path=ibdata1:50M;ibdata2:50M:autoextend
```

The full syntax for a data file specification includes the filename, its size, and several optional attributes:

```
file_name:file_size[:autoextend[:max:max_file_size]]
```

The `autoextend` attribute and those following can be used only for the last data file in the `innodb_data_file_path` line. `autoextend` is available starting from MySQL 3.23.50 and 4.0.2.

If you specify the `autoextend` option for the last data file, [InnoDB](#) extends the data file if it runs out of free space in the tablespace. The increment is 8MB at a time by default. It can be modified by changing the `innodb_autoextend_increment` system variable.

If the disk becomes full, you might want to add another data file on another disk. Instructions for reconfiguring an existing tablespace are given in [Section 14.2.8, “Adding and Removing InnoDB Data and Log Files”](#).

[InnoDB](#) is not aware of the filesystem maximum file size, so be cautious on filesystems where the maximum file size is a small value such as 2GB. To specify a maximum size for an auto-extending data file, use the `max` attribute. The following configuration allows `ibdata1` to grow up to a limit of 500MB:

```
[mysqld]
innodb_data_file_path=ibdata1:10M:autoextend:max:500M
```

[InnoDB](#) creates tablespace files in the MySQL data directory by default. To specify a location explicitly, use the `innodb_data_home_dir` option. For example, to use two files named `ibdata1` and



`ibdata2` but create them in the `/ibdata` directory, configure `InnoDB` like this:

```
[mysqld]
innodb_data_home_dir = /ibdata
innodb_data_file_path=ibdata1:50M;ibdata2:50M:autoextend
```

**Note:** `InnoDB` does not create directories, so make sure that the `/ibdata` directory exists before you start the server. This is also true of any log file directories that you configure. Use the Unix or DOS `mkdir` command to create any necessary directories.

`InnoDB` forms the directory path for each data file by textually concatenating the value of `innodb_data_home_dir` to the data file name, adding a pathname separator (slash or backslash) between values if necessary. If the `innodb_data_home_dir` option is not mentioned in `my.cnf` at all, the default value is the “dot” directory `./`, which means the MySQL data directory. (The MySQL server changes its current working directory to its data directory when it begins executing.)

If you specify `innodb_data_home_dir` as an empty string, you can specify absolute paths for the data files listed in the `innodb_data_file_path` value. The following example is equivalent to the preceding one:

```
[mysqld]
innodb_data_home_dir =
innodb_data_file_path=/ibdata/ibdata1:50M;/ibdata/ibdata2:50M:autoextend
```

**A simple `my.cnf` example.** Suppose that you have a computer with 128MB RAM and one hard disk. The following example shows possible configuration parameters in `my.cnf` or `my.ini` for `InnoDB`. The example assumes the use of MySQL-Max 3.23.50 or later or MySQL 4.0.2 or later because it uses the `autoextend` attribute. The example suits most users, both on Unix and Windows, who do not want to distribute `InnoDB` data files and log files onto several disks. It creates an auto-extending data file `ibdata1` and two `InnoDB` log files `ib_logfile0` and `ib_logfile1` in the MySQL data directory. Also, the small archived `InnoDB` log file `ib_arch_log_0000000000` that `InnoDB` creates automatically ends up in the data directory.

```
[mysqld]
# You can write your other MySQL server options here
# ...
# Data files must be able to hold your data and indexes.
# Make sure that you have enough free disk space.
innodb_data_file_path = ibdata1:10M:autoextend
#
# Set buffer pool size to 50-80% of your computer's memory
set-variable = innodb_buffer_pool_size=70M
set-variable = innodb_additional_mem_pool_size=10M
#
# Set the log file size to about 25% of the buffer pool size
set-variable = innodb_log_file_size=20M
set-variable = innodb_log_buffer_size=8M
#
innodb_flush_log_at_trx_commit=1
```

Make sure that the MySQL server has the proper access rights to create files in the data directory. More generally, the server must have access rights in any directory where it needs to create data files or log files.

Note that data files must be less than 2GB in some filesystems. The combined size of the log files must be less than 4GB. The combined size of data files must be at least 10MB.

When you create an `InnoDB` tablespace for the first time, it is best that you start the MySQL server from the command prompt. `InnoDB` then prints the information about the database creation to the screen, so you can see what is happening. For example, on Windows, if `mysqld` is located in `C:\Program Files\MySQL\MySQL Server 4.1\bin`, you can start it like this:

```
C:\> "C:\Program Files\MySQL\MySQL Server 4.1\bin\mysqld" --console
```

If you do not send server output to the screen, check the server's error log to see what `InnoDB` prints during the startup process.

See [Section 14.2.6, “Creating the InnoDB Tablespace”](#), for an example of what the information displayed by `InnoDB` should look like.

You can place `InnoDB` options in the `[mysqld]` group of any option file that your server reads when it starts. The locations for option files are described in [Section 4.3.2, “Using Option Files”](#).

If you installed MySQL on Windows using the installation and configuration wizards, the option file will be the `my.ini` file located in your MySQL installation directory. See [Section 2.3.4.14, “The Location of the my.ini File”](#).

If your PC uses a boot loader where the `C:` drive is not the boot drive, your only option is to use the `my.ini` file in your Windows directory (typically `C:\WINDOWS` or `C:\WINNT`). You can use the `SET` command at the command prompt in a console window to print the value of `WINDIR`:

```
C:\> SET WINDIR
windir=C:\WINDOWS
```

If you want to make sure that `mysqld` reads options only from a specific file, you can use the `--defaults-file` option as the first option on the command line when starting the server:

```
mysqld --defaults-file=your_path_to_my_cnf
```

**An advanced `my.cnf` example.** Suppose that you have a Linux computer with 2GB RAM and three 60GB hard disks at directory paths `/`, `/dr2` and `/dr3`. The following example shows possible configuration parameters in `my.cnf` for `InnoDB`.

```
[mysqld]
# You can write your other MySQL server options here
# ...
innodb_data_home_dir =
#
# Data files must be able to hold your data and indexes
innodb_data_file_path = /ibdata/ibdata1:2000M;/dr2/ibdata/ibdata2:2000M:autoextend
#
# Set buffer pool size to 50-80% of your computer's memory,
# but make sure on Linux x86 total memory usage is < 2GB
set-variable = innodb_buffer_pool_size=1G
set-variable = innodb_additional_mem_pool_size=20M
innodb_log_group_home_dir = /dr3/iblogs
#
# innodb_log_arch_dir must be the same as innodb_log_group_home_dir
# (starting from 4.0.6, you can omit it)
innodb_log_arch_dir = /dr3/iblogs
set-variable = innodb_log_files_in_group=2
#
# Set the log file size to about 25% of the buffer pool size
set-variable = innodb_log_file_size=250M
set-variable = innodb_log_buffer_size=8M
#
innodb_flush_log_at_trx_commit=1
set-variable = innodb_lock_wait_timeout=50
#
# Uncomment the next lines if you want to use them
#set-variable = innodb_thread_concurrency=5
```

In some cases, database performance improves if all data is not placed on the same physical disk. Putting log files on a different disk from data is very often beneficial for performance. The example illustrates how to do this. It places the two data files on different disks and places the log files on the third disk. `InnoDB` fills the tablespace beginning with the first data file. You can also use raw disk partitions (raw devices) as `InnoDB` data files, which may speed up I/O. See [Section 14.2.4.2, “Using Raw Devices for the Shared Tablespace”](#).

**Warning:** On 32-bit GNU/Linux x86, you must be careful not to set memory usage too high. `glibc` may allow the process heap to grow over thread stacks, which crashes your server. It is a risk if the value of the following expression is close to or exceeds 2GB:

```
innodb_buffer_pool_size
+ key_buffer_size
+ max_connections*(sort_buffer_size+read_buffer_size+binlog_cache_size)
+ max_connections*2MB
```

Each thread uses a stack (often 2MB, but only 256KB in MySQL AB binaries) and in the worst case also uses `sort_buffer_size + read_buffer_size` additional memory.

In MySQL 4.1, by compiling MySQL yourself, you can use up to 64GB of physical memory in 32-bit Windows. See the description for `innodb_buffer_pool_awesome_mem_mb` in Section 14.2.5, “InnoDB Startup Options and System Variables”.

**How to tune other `mysqld` server parameters?** The following values are typical and suit most users:

```
[mysqld]
skip-external-locking
set-variable = max_connections=200
set-variable = read_buffer_size=1M
set-variable = sort_buffer_size=1M
#
# Set key_buffer to 5 - 50% of your RAM depending on how much
# you use MyISAM tables, but keep key_buffer_size + InnoDB
# buffer pool size < 80% of your RAM
set-variable = key_buffer_size=...
```

### 14.2.4.1. Using Per-Table Tablespaces

**Note:** There is a known bug in versions prior to 4.1.8 that manifests itself if you specify `innodb_file_per_table` in `my.cnf`. If you shut down `mysqld`, then records may disappear from the secondary indexes of a table. See [Bug#7496](http://bugs.mysql.com/7496) [http://bugs.mysql.com/7496] for more information and workarounds. This is fixed in 4.1.9, but another bug ([Bug#8021](http://bugs.mysql.com/8021) [http://bugs.mysql.com/8021]) bit the Windows version in 4.1.9, and in the Windows version of 4.1.9 you must put the line `innodb_flush_method=unbuffered` to your `my.cnf` or `my.ini` to get `mysqld` to work.

Starting from MySQL 4.1.1, you can store each InnoDB table and its indexes in its own file. This feature is called “multiple tablespaces” because in effect each table has its own tablespace.

Using multiple tablespaces can be beneficial to users who want to move specific tables to separate physical disks or who wish to restore backups of single tables quickly without interrupting the use of the remaining InnoDB tables.

If you need to downgrade to 4.0, you must make table dumps and re-create the whole InnoDB tablespace. If you have not created new InnoDB tables under MySQL 4.1.1 or later, and need to downgrade quickly, you can also do a direct downgrade to the MySQL 4.0.18 or later in the 4.0 series. Before doing the direct downgrade to 4.0.x, you have to end all client connections to the `mysqld` server that is to be downgraded, and let it run the purge and insert buffer merge operations to completion, so that `SHOW INNODB STATUS` shows the main thread in the state `waiting for server activity`. Then you can shut down `mysqld` and start 4.0.18 or later in the 4.0 series.

You can enable multiple tablespaces by adding a line to the `[mysqld]` section of `my.cnf`:

```
[mysqld]
innodb_file_per_table
```

After restarting the server, InnoDB stores each newly created table into its own file `tbl_name.ibd` in the database directory where the table belongs. This is similar to what the MyISAM storage engine does, but MyISAM divides the table into a data file `tbl_name.MYD` and the index file

`.MYI`. For `InnoDB`, the data and the indexes are stored together in the `.ibd` file. The `tbl_name.frm` file is still created as usual.

If you remove the `innodb_file_per_table` line from `my.cnf` and restart the server, `InnoDB` creates tables inside the shared tablespace files again.

`innodb_file_per_table` affects only table creation, not access to existing tables. If you start the server with this option, new tables are created using `.ibd` files, but you can still access tables that exist in the shared tablespace. If you remove the option and restart the server, new tables are created in the shared tablespace, but you can still access any tables that were created using multiple tablespaces.

**Note:** `InnoDB` always needs the shared tablespace because it puts its internal data dictionary and undo logs there. The `.ibd` files are not sufficient for `InnoDB` to operate.

**Note:** You cannot freely move `.ibd` files between database directories as you can with `MyISAM` table files. This is because the table definition that is stored in the `InnoDB` shared tablespace includes the database name, and because `InnoDB` must preserve the consistency of transaction IDs and log sequence numbers.

To move an `.ibd` file and the associated table from one database to another, use a `RENAME TABLE` statement:

```
RENAME TABLE db1.tbl_name TO db2.tbl_name;
```

If you have a “clean” backup of an `.ibd` file, you can restore it to the MySQL installation from which it originated as follows:

1. Issue this `ALTER TABLE` statement:

```
ALTER TABLE tbl_name DISCARD TABLESPACE;
```

**Caution:** This statement deletes the current `.ibd` file.

2. Put the backup `.ibd` file back in the proper database directory.
3. Issue this `ALTER TABLE` statement:

```
ALTER TABLE tbl_name IMPORT TABLESPACE;
```

In this context, a “clean” `.ibd` file backup means:

- There are no uncommitted modifications by transactions in the `.ibd` file.
- There are no unmerged insert buffer entries in the `.ibd` file.
- Purge has removed all delete-marked index records from the `.ibd` file.
- `mysqld` has flushed all modified pages of the `.ibd` file from the buffer pool to the file.

You can make a clean backup `.ibd` file using the following method:

1. Stop all activity from the `mysqld` server and commit all transactions.
2. Wait until `SHOW INNODB STATUS` shows that there are no active transactions in the database,

and the main thread status of `InnoDB` is `Waiting for server activity`. Then you can make a copy of the `.ibd` file.

Another method for making a clean copy of an `.ibd` file is to use the commercial `InnoDB Hot Backup` tool:

1. Use `InnoDB Hot Backup` to back up the `InnoDB` installation.
2. Start a second `mysqld` server on the backup and let it clean up the `.ibd` files in the backup.

### 14.2.4.2. Using Raw Devices for the Shared Tablespace

Starting from MySQL 3.23.41, you can use raw disk partitions as data files in the shared tablespace. By using a raw disk, you can perform non-buffered I/O on Windows and on some Unix systems without filesystem overhead, which might improve performance.

When you create a new data file, you must put the keyword `newraw` immediately after the data file size in `innodb_data_file_path`. The partition must be at least as large as the size that you specify. Note that 1MB in `InnoDB` is  $1024 \times 1024$  bytes, whereas 1MB in disk specifications usually means 1,000,000 bytes.

```
[mysqld]
innodb_data_home_dir=
innodb_data_file_path=/dev/hdd1:3Gnewraw;/dev/hdd2:2Gnewraw
```

The next time you start the server, `InnoDB` notices the `newraw` keyword and initializes the new partition. However, do not create or change any `InnoDB` tables yet. Otherwise, when you next restart the server, `InnoDB` reinitializes the partition and your changes are lost. (Starting from MySQL 3.23.44, as a safety measure `InnoDB` prevents users from modifying data when any partition with `newraw` is specified.)

After `InnoDB` has initialized the new partition, stop the server, change `newraw` in the data file specification to `raw`:

```
[mysqld]
innodb_data_home_dir=
innodb_data_file_path=/dev/hdd1:5Graw;/dev/hdd2:2Graw
```

Then restart the server and `InnoDB` allows changes to be made.

On Windows, starting from 4.1.1, you can allocate a disk partition as a data file like this:

```
[mysqld]
innodb_data_home_dir=
innodb_data_file_path=//./D:10Gnewraw
```

The `//./` corresponds to the Windows syntax of `\\.\` for accessing physical drives.

When you use raw disk partitions, be sure that they have permissions that allow read and write access by the account used for running the MySQL server.

## 14.2.5. InnoDB Startup Options and System Variables

This section describes the `InnoDB`-related command options and system variables. System variables that take a numeric value can be specified as `--var_name=value` on the command line or as `var_name=value` in option files. Many of the system variables can be changed at runtime (see [Sec-](#)

tion 5.2.3.2, “Dynamic System Variables”). (Before MySQL 4.0.2, system variable values should be specified using `--set-variable` syntax.) For more information on specifying options and system variables, see Section 4.3, “Specifying Program Options”.

InnoDB command options:

- `--innodb`  
Enables the InnoDB storage engine, if the server was compiled with InnoDB support. Use `--skip-innodb` to disable InnoDB.
- `--innodb_status_file`  
Causes InnoDB to create a file named `<datadir>/innodb_status.<pid>` in the MySQL data directory. InnoDB periodically writes the output of `SHOW ENGINE INNODB STATUS` to this file. This option is available as of MySQL 4.0.21.

InnoDB system variables:

- `innodb_additional_mem_pool_size`  
The size in bytes of a memory pool InnoDB uses to store data dictionary information and other internal data structures. The more tables you have in your application, the more memory you need to allocate here. If InnoDB runs out of memory in this pool, it starts to allocate memory from the operating system, and writes warning messages to the MySQL error log. The default value is 1MB.
- `innodb_autoextend_increment`  
The increment size (in MB) for extending the size of an auto-extending tablespace when it becomes full. The default value is 8. This variable is available starting from MySQL 4.0.24 and 4.1.5. As of MySQL 4.0.24 and 4.1.6, it can be changed at runtime as a global system variable.
- `innodb_buffer_pool_awe_mem_mb`  
The size of the buffer pool (in MB), if it is placed in the AWE memory. This is relevant only in 32-bit Windows. If your 32-bit Windows operating system supports more than 4GB memory, using so-called “Address Windowing Extensions,” you can allocate the InnoDB buffer pool into the AWE physical memory using this variable. The maximum possible value for this variable is 63000. If it is greater than 0, `innodb_buffer_pool_size` is the window in the 32-bit address space of `mysqld` where InnoDB maps that AWE memory. A good value for `innodb_buffer_pool_size` is 500MB. This variable is available as of MySQL 4.1.0.  
  
To take advantage of AWE memory, you will need to recompile MySQL yourself. The current project settings needed for doing this can be found in the `innobase/os/os0proj.c` source file.
- `innodb_buffer_pool_size`  
The size in bytes of the memory buffer InnoDB uses to cache data and indexes of its tables. The larger you set this value, the less disk I/O is needed to access data in tables. On a dedicated database server, you may set this to up to 80% of the machine physical memory size. However, do not set it too large because competition for physical memory might cause paging in the operating system.
- `innodb_data_file_path`  
The paths to individual data files and their sizes. The full directory path to each data file is acquired by concatenating `innodb_data_home_dir` to each path specified here. The file sizes are specified in MB or GB (1024MB) by appending `M` or `G` to the size value. The sum of the sizes of the

files must be at least 10MB. On some operating systems, files must be less than 2GB. If you do not specify `innodb_data_file_path`, the default behavior starting from 4.0 is to create a single 10MB auto-extending data file named `ibdata1`. Starting from 3.23.44, you can set the file size larger than 4GB on those operating systems that support big files. You can also use raw disk partitions as data files. See [Section 14.2.4.2, “Using Raw Devices for the Shared Tablespace”](#).

- `innodb_data_home_dir`

The common part of the directory path for all InnoDB data files. If you do not set this value, the default is the MySQL data directory. You can specify this also as an empty string, in which case you can use absolute file paths in `innodb_data_file_path`.

- `innodb_fast_shutdown`

If you set this variable to 0, InnoDB does a full purge and an insert buffer merge before a shutdown. These operations can take minutes, or even hours in extreme cases. If you set this variable to 1, InnoDB skips these operations at shutdown. The default value is 1 starting from 3.23.50.

- `innodb_file_io_threads`

The number of file I/O threads in InnoDB. Normally, this should be left at the default value of 4, but disk I/O on Windows may benefit from a larger number. On Unix, increasing the number has no effect; InnoDB always uses the default value. This variable is available as of MySQL 3.23.37.

- `innodb_file_per_table`

**NOTE:** A bug in versions  $\leq 4.1.8$  if you specify `innodb_file_per_table` in `my.cnf`! If you shut down `mysqld`, then records may disappear from the secondary indexes of a table. See [Bug#7496](http://bugs.mysql.com/7496) [http://bugs.mysql.com/7496] for more information and workarounds. This is fixed in 4.1.9, but another bug ([Bug#8021](http://bugs.mysql.com/8021) [http://bugs.mysql.com/8021]) bit the Windows version in 4.1.9, and in the Windows version of 4.1.9 you must put the line `innodb_flush_method=unbuffered` in your `my.cnf` or `my.ini` to get `mysqld` to work.

If this variable is enabled, InnoDB creates each new table using its own `.ibd` file for storing data and indexes, rather than in the shared tablespace. The default is to create tables in the shared tablespace. See [Section 14.2.4.1, “Using Per-Table Tablespaces”](#). This variable is available as of MySQL 4.1.1.

- `innodb_flush_log_at_trx_commit`

When `innodb_flush_log_at_trx_commit` is set to 0, the log buffer is written out to the log file once per second and the flush to disk operation is performed on the log file, but nothing is done at a transaction commit. When this value is 1 (the default), the log buffer is written out to the log file at each transaction commit and the flush to disk operation is performed on the log file. When set to 2, the log buffer is written out to the file at each commit, but the flush to disk operation is not performed on it. However, the flushing on the log file takes place once per second also when the value is 2. Note that the once-per-second flushing is not 100% guaranteed to happen every second, due to process scheduling issues.

The default value of this variable is 1, which is the value that is required for ACID compliance. You can achieve better performance by setting the value different from 1, but then you can lose at most one second worth of transactions in a crash. If you set the value to 0, then any `mysqld` process crash can erase the last second of transactions. If you set the value to 2, then only an operating system crash or a power outage can erase the last second of transactions. However, InnoDB's crash recovery is not affected and thus crash recovery does work regardless of the value. Note that many operating systems and some disk hardware fool the flush-to-disk operation. They may tell `mysqld` that the flush has taken place, even though it has not. Then the durability of transactions is not guaranteed even with the setting 1, and in the worst case a power outage can even corrupt the InnoDB database. Using a battery-backed disk cache in the SCSI disk controller or in the disk itself speeds up

file flushes, and makes the operation safer. You can also try using the Unix command `hdparm` to disable the caching of disk writes in hardware caches, or use some other command specific to the hardware vendor. The default value of this variable is 1 (prior to MySQL 4.0.13, the default is 0).

Note: For the greatest possible durability and consistency in a replication setup using `InnoDB` with transactions, you should use `innodb_flush_log_at_trx_commit=1`, `sync_binlog=1`, and `innodb_safe_binlog` in your master server `my.cnf` file.

- `innodb_flush_method`

If set to `fdatasync` (the default), `InnoDB` uses `fsync()` to flush both the data and log files. If set to `O_DSYNC`, `InnoDB` uses `O_SYNC` to open and flush the log files, but uses `fsync()` to flush the data files. If `O_DIRECT` is specified (available on some GNU/Linux versions starting from MySQL 4.0.14), `InnoDB` uses `O_DIRECT` to open the data files, and uses `fsync()` to flush both the data and log files. Note that starting from MySQL 3.23.41, `InnoDB` uses `fsync()` instead of `fdatasync()`, and it does not use `O_DSYNC` by default because there have been problems with it on many varieties of Unix. This variable is relevant only for Unix. On Windows, the flush method is always `async_unbuffered` and cannot be changed. This variable is available as of MySQL 3.23.40.

Different values of this variable can have a marked effect on `InnoDB` performance. For example, on some systems where `InnoDB` data and log files are located on a SAN, it has been found that setting `innodb_flush_method` to `O_DIRECT` can degrade performance of simple `SELECT` statements by a factor of three.

- `innodb_force_recovery`

The crash recovery mode. Warning: This variable should be set greater than 0 only in an emergency situation when you want to dump your tables from a corrupt database! Possible values are from 1 to 6. The meanings of these values are described in [Section 14.2.9.1, “Forcing InnoDB Recovery”](#). As a safety measure, `InnoDB` prevents any changes to its data when this variable is greater than 0. This variable is available starting from MySQL 3.23.44.

- `innodb_lock_wait_timeout`

The timeout in seconds an `InnoDB` transaction may wait for a lock before being rolled back. `InnoDB` automatically detects transaction deadlocks in its own lock table and rolls back the transaction. Beginning with MySQL 4.0.20 and 4.1.2, `InnoDB` notices locks set using the `LOCK TABLES` statement. Before that, if you use the `LOCK TABLES` statement, or other transaction-safe storage engines than `InnoDB` in the same transaction, a deadlock may arise that `InnoDB` cannot notice. In cases like this, the timeout is useful to resolve the situation. The default is 50 seconds.

- `innodb_locks_unsafe_for_binlog`

This variable controls next-key locking in `InnoDB` searches and index scans. By default, this variable is 0 (disabled), which means that next-key locking is enabled.

Normally, `InnoDB` uses an algorithm called *next-key locking*. `InnoDB` performs row-level locking in such a way that when it searches or scans a table index, it sets shared or exclusive locks on any index records it encounters. Thus, the row-level locks are actually index record locks. The locks that `InnoDB` sets on index records also affect the “gap” preceding that index record. If a user has a shared or exclusive lock on record *R* in an index, another user cannot insert a new index record immediately before *R* in the order of the index. Enabling this variable causes `InnoDB` not to use next-key locking in searches or index scans. Next-key locking is still used to ensure foreign key constraints and duplicate key checking. Note that enabling this variable may cause phantom problems: Suppose that you want to read and lock all children from the `child` table with an identifier value larger than 100, with the intention of updating some column in the selected rows later:

```
SELECT * FROM child WHERE id > 100 FOR UPDATE;
```



Suppose that there is an index on the `id` column. The query scans that index starting from the first record where `id` is larger than 100. If the locks set on the index records do not lock out inserts made in the gaps, another client can insert a new row into the table. If you execute the same `SELECT` within the same transaction, you see a new row in the result set returned by the query. This also means that if new items are added to the database, `InnoDB` does not guarantee serializability. Therefore, if this variable is enabled `InnoDB` guarantees at most isolation level `READ COMMITTED`. (Conflict serializability is still guaranteed.) This variable is available as of MySQL 4.1.4.

- `innodb_log_arch_dir`

The directory where fully written log files would be archived if we used log archiving. The value of this variable should currently be set the same as `innodb_log_group_home_dir`. Starting from MySQL 4.0.6, there is no need to set this variable.

- `innodb_log_archive`

Whether to log `InnoDB` archive files. This variable is unused. Recovery from a backup is done by MySQL using its own log files, so there is no need to archive `InnoDB` log files. The default for this variable is 0.

- `innodb_log_buffer_size`

The size in bytes of the buffer that `InnoDB` uses to write to the log files on disk. Sensible values range from 1MB to 8MB. The default is 1MB. A large log buffer allows large transactions to run without a need to write the log to disk before the transactions commit. Thus, if you have big transactions, making the log buffer larger saves disk I/O.

- `innodb_log_file_size`

The size in bytes of each log file in a log group. The combined size of log files must be less than 4GB on 32-bit computers. The default is 5MB. Sensible values range from 1MB to 1/*N*-th of the size of the buffer pool, where *N* is the number of log files in the group. The larger the value, the less checkpoint flush activity is needed in the buffer pool, saving disk I/O. But larger log files also mean that recovery is slower in case of a crash.

- `innodb_log_files_in_group`

The number of log files in the log group. `InnoDB` writes to the files in a circular fashion. The default (and recommended) is 2.

- `innodb_log_group_home_dir`

The directory path to the `InnoDB` log files. It must have the same value as `innodb_log_arch_dir`. If you do not specify any `InnoDB` log variables, the default is to create two 5MB files names `ib_logfile0` and `ib_logfile1` in the MySQL data directory.

- `innodb_max_dirty_pages_pct`

This is an integer in the range from 0 to 100. The default is 90. The main thread in `InnoDB` tries to write pages from the buffer pool so that the percentage of dirty (not yet written) pages will not exceed this value. Available starting from 4.0.13 and 4.1.1.

- `innodb_max_purge_lag`

This variable controls how to delay `INSERT`, `UPDATE` and `DELETE` operations when the purge operations are lagging (see Section 14.2.13, “Implementation of Multi-Versioning”). The default value

of this variable is 0, meaning that there are no delays. `innodb_max_purge_lag` is available as of MySQL 4.0.22 and 4.1.6.

The InnoDB transaction system maintains a list of transactions that have delete-marked index records by `UPDATE` or `DELETE` operations. Let the length of this list be `purge_lag`. When `purge_lag` exceeds `innodb_max_purge_lag`, each `INSERT`, `UPDATE` and `DELETE` operation is delayed by  $((\text{purge\_lag}/\text{innodb\_max\_purge\_lag}) \times 10) - 5$  milliseconds. The delay is computed in the beginning of a purge batch, every ten seconds. The operations are not delayed if purge cannot run because of an old consistent read view that could see the rows to be purged.

A typical setting for a problematic workload might be 1 million, assuming that our transactions are small, only 100 bytes in size, and we can allow 100MB of unpurged rows in our tables.

- `innodb_mirrored_log_groups`

The number of identical copies of log groups to keep for the database. Currently, this should be set to 1.

- `innodb_open_files`

This variable is relevant only if you use multiple tablespaces in InnoDB. It specifies the maximum number of `.ibd` files that InnoDB can keep open at one time. The minimum value is 10. The default is 300. This variable is available as of MySQL 4.1.1.

The file descriptors used for `.ibd` files are for InnoDB only. They are independent of those specified by the `--open-files-limit` server option, and do not affect the operation of the table cache.

- `innodb_safe_binlog`

Adds consistency guarantees between the content of InnoDB tables and the binary log. See [Section 5.11.4, “The Binary Log”](#).

- `innodb_table_locks`

Starting from MySQL 4.0.20, and 4.1.2, InnoDB honors `LOCK TABLES`; MySQL does not return from `LOCK TABLE ... WRITE` until all other threads have released all their locks to the table. In MySQL 4.0.19 and before, InnoDB ignored table locks, which allowed one to more easily simulate transactions with a combination of MyISAM and InnoDB tables. The default value is 1, which means that `LOCK TABLES` causes also InnoDB internally to take a table lock. In applications using `AUTOCOMMIT=1`, InnoDB's internal table locks can cause deadlocks. You can set `innodb_table_locks=0` in the server option file to remove that problem.

- `innodb_thread_concurrency`

InnoDB tries to keep the number of operating system threads concurrently inside InnoDB less than or equal to the limit given by this variable. The default value is 8. If you have low performance and `SHOW INNODB STATUS` reveals many threads waiting for semaphores, you may have thread thrashing and should try setting this variable lower or higher. If you have a computer with many processors and disks, you can try setting the value higher to better utilize the resources of your computer. A recommended value is the sum of the number of processors and disks your system has. A value of 500 or greater disables the concurrency checking. This variable is available starting from MySQL 3.23.44 and 4.0.1.

- `sync_binlog`

If the value of this variable is positive, the MySQL server synchronizes its binary log to disk (`fsync()`) after every `sync_binlog` writes to this binary log. Note that there is one write to the binary log per statement if in autocommit mode, and otherwise one write per transaction. The

default value is 0 which does no sync'ing to disk. A value of 1 is the safest choice, because in case of crash you lose at most one statement/transaction from the binary log; but it is also the slowest choice (unless the disk has a battery-backed cache, which makes sync'ing very fast). This variable was added in MySQL 4.1.3.

## 14.2.6. Creating the InnoDB Tablespace

Suppose that you have installed MySQL and have edited your option file so that it contains the necessary InnoDB configuration parameters. Before starting MySQL, you should verify that the directories you have specified for InnoDB data files and log files exist and that the MySQL server has access rights to those directories. InnoDB does not create directories, only files. Check also that you have enough disk space for the data and log files.

It is best to run the MySQL server `mysqld` from the command prompt when you first start the server with InnoDB enabled, not from the `mysqld_safe` wrapper or as a Windows service. When you run from a command prompt you see what `mysqld` prints and what is happening. On Unix, just invoke `mysqld`. On Windows, use the `--console` option.

When you start the MySQL server after initially configuring InnoDB in your option file, InnoDB creates your data files and log files, and prints something like this:

```
InnoDB: The first specified datafile /home/heikki/data/ibdata1
did not exist:
InnoDB: a new database to be created!
InnoDB: Setting file /home/heikki/data/ibdata1 size to 134217728
InnoDB: Database physically writes the file full: wait...
InnoDB: datafile /home/heikki/data/ibdata2 did not exist:
new to be created
InnoDB: Setting file /home/heikki/data/ibdata2 size to 262144000
InnoDB: Database physically writes the file full: wait...
InnoDB: Log file /home/heikki/data/logs/ib_logfile0 did not exist:
new to be created
InnoDB: Setting log file /home/heikki/data/logs/ib_logfile0 size
to 5242880
InnoDB: Log file /home/heikki/data/logs/ib_logfile1 did not exist:
new to be created
InnoDB: Setting log file /home/heikki/data/logs/ib_logfile1 size
to 5242880
InnoDB: Doublewrite buffer not found: creating new
InnoDB: Doublewrite buffer created
InnoDB: Creating foreign key constraint system tables
InnoDB: Foreign key constraint system tables created
InnoDB: Started
mysqld: ready for connections
```

At this point InnoDB has initialized its tablespace and log files. You can connect to the MySQL server with the usual MySQL client programs like `mysql`. When you shut down the MySQL server with `mysqladmin shutdown`, the output is like this:

```
010321 18:33:34 mysqld: Normal shutdown
010321 18:33:34 mysqld: Shutdown Complete
InnoDB: Starting shutdown...
InnoDB: Shutdown completed
```

You can look at the data file and log directories and you see the files created there. The log directory also contains a small file named `ib_arch_log_0000000000`. That file resulted from the database creation, after which InnoDB switched off log archiving. When MySQL is started again, the data files and log files have been created already, so the output is much briefer:

```
InnoDB: Started
mysqld: ready for connections
```

Starting from MySQL 4.1.1, you can add the option `innodb_file_per_table` to `my.cnf` to make

**InnoDB** store each table to its own `.ibd` file in the same MySQL database directory where the `.frm` file is created. See [Section 14.2.4.1, “Using Per-Table Tablespaces”](#).

### 14.2.6.1. Dealing with **InnoDB** Initialization Problems

If **InnoDB** prints an operating system error during a file operation, usually the problem has one of the following causes:

- You did not create the **InnoDB** data file directory or the **InnoDB** log directory.
- `mysqld` does not have access rights to create files in those directories.
- `mysqld` cannot read the proper `my.cnf` or `my.ini` option file, and consequently does not see the options that you specified.
- The disk is full or a disk quota is exceeded.
- You have created a subdirectory whose name is equal to a data file that you specified, so the name cannot be used as a filename.
- There is a syntax error in the `innodb_data_home_dir` or `innodb_data_file_path` value.

If something goes wrong when **InnoDB** attempts to initialize its tablespace or its log files, you should delete all files created by **InnoDB**. This means all `ibdata` files and all `ib_logfile` files. In case you have already created some **InnoDB** tables, delete the corresponding `.frm` files for these tables (and any `.ibd` files if you are using multiple tablespaces) from the MySQL database directories as well. Then you can try the **InnoDB** database creation again. It is best to start the MySQL server from a command prompt so that you see what is happening.

### 14.2.7. Creating and Using **InnoDB** Tables

To create an **InnoDB** table, specify an `ENGINE = InnoDB` option in the `CREATE TABLE` statement:

```
CREATE TABLE customers (a INT, b CHAR (20), INDEX (a)) ENGINE=InnoDB;
```

The older term `TYPE` is supported as a synonym for `ENGINE` for backward compatibility, but `ENGINE` is the preferred term and `TYPE` is deprecated.

The statement creates a table and an index on column `a` in the **InnoDB** tablespace that consists of the data files that you specified in `my.cnf`. In addition, MySQL creates a file `customers.frm` in the `test` directory under the MySQL database directory. Internally, **InnoDB** adds an entry for the table to its own data dictionary. The entry includes the database name. For example, if `test` is the database in which the `customers` table is created, the entry is for `'test/customers'`. This means you can create a table of the same name `customers` in some other database, and the table names do not collide inside **InnoDB**.

You can query the amount of free space in the **InnoDB** tablespace by issuing a `SHOW TABLE STATUS` statement for any **InnoDB** table. The amount of free space in the tablespace appears in the `Comment` section in the output of `SHOW TABLE STATUS`. For example:

```
SHOW TABLE STATUS FROM test LIKE 'customers'
```

Note that the statistics `SHOW` displays for **InnoDB** tables are only approximate. They are used in SQL optimization. Table and index reserved sizes in bytes are accurate, though.

#### 14.2.7.1. How to Use Transactions in **InnoDB** with Different APIs

By default, each client that connects to the MySQL server begins with autocommit mode enabled, which automatically commits every SQL statement as you execute it. To use multiple-statement transactions, you can switch autocommit off with the SQL statement `SET AUTOCOMMIT = 0` and use `COMMIT` and `ROLLBACK` to commit or roll back your transaction. If you want to leave autocommit on, you can enclose your transactions within `START TRANSACTION` and either `COMMIT` or `ROLLBACK`. Before MySQL 4.0.11, you have to use the keyword `BEGIN` instead of `START TRANSACTION`. The following example shows two transactions. The first is committed and the second is rolled back.

```
shell> mysql test

mysql> CREATE TABLE CUSTOMER (A INT, B CHAR (20), INDEX (A))
-> TYPE=InnoDB;
Query OK, 0 rows affected (0.00 sec)
mysql> BEGIN;
Query OK, 0 rows affected (0.00 sec)
mysql> INSERT INTO CUSTOMER VALUES (10, 'Heikki');
Query OK, 1 row affected (0.00 sec)
mysql> COMMIT;
Query OK, 0 rows affected (0.00 sec)
mysql> SET AUTOCOMMIT=0;
Query OK, 0 rows affected (0.00 sec)
mysql> INSERT INTO CUSTOMER VALUES (15, 'John');
Query OK, 1 row affected (0.00 sec)
mysql> ROLLBACK;
Query OK, 0 rows affected (0.00 sec)
mysql> SELECT * FROM CUSTOMER;
+-----+-----+
| A     | B           |
+-----+-----+
| 10   | Heikki     |
+-----+-----+
1 row in set (0.00 sec)
mysql>
```

In APIs such as PHP, Perl DBI, JDBC, ODBC, or the standard C call interface of MySQL, you can send transaction control statements such as `COMMIT` to the MySQL server as strings just like any other SQL statements such as `SELECT` or `INSERT`. Some APIs also offer separate special transaction commit and rollback functions or methods.

### 14.2.7.2. Converting MyISAM Tables to InnoDB

Important: Do not convert MySQL system tables in the `mysql` database (such as `user` or `host`) to the `InnoDB` type. This is an unsupported operation. The system tables must always be of the `MyISAM` type.

If you want all your (non-system) tables to be created as `InnoDB` tables, you can, starting from the MySQL 3.23.43, add the line `default-table-type=innodb` to the `[mysqld]` section of your server option file.

`InnoDB` does not have a special optimization for separate index creation the way the `MyISAM` storage engine does. Therefore, it does not pay to export and import the table and create indexes afterward. The fastest way to alter a table to `InnoDB` is to do the inserts directly to an `InnoDB` table. That is, use `ALTER TABLE ... TYPE=INNODB`, or create an empty `InnoDB` table with identical definitions and insert the rows with `INSERT INTO ... SELECT * FROM ...`

If you have `UNIQUE` constraints on secondary keys, starting from MySQL 3.23.52, you can speed up a table import by turning off the uniqueness checks temporarily during the import operation:

```
SET UNIQUE_CHECKS=0;
... import operation ...
SET UNIQUE_CHECKS=1;
```

For big tables, this saves a lot of disk I/O because `InnoDB` can then use its insert buffer to write secondary index records in a batch. Be certain that the data contains no duplicate keys. `UNIQUE_CHECKS` allows but does not require storage engines to ignore duplicate keys.

To get better control over the insertion process, it might be good to insert big tables in pieces:

```
INSERT INTO newtable SELECT * FROM oldtable
WHERE yourkey > something AND yourkey <= somethingelse;
```

After all records have been inserted, you can rename the tables.

During the conversion of big tables, you should increase the size of the [InnoDB](#) buffer pool to reduce disk I/O. Do not use more than 80% of the physical memory, though. You can also increase the sizes of the [InnoDB](#) log files.

Make sure that you do not fill up the tablespace: [InnoDB](#) tables require a lot more disk space than [MyISAM](#) tables. If an `ALTER TABLE` operation runs out of space, it starts a rollback, and that can take hours if it is disk-bound. For inserts, [InnoDB](#) uses the insert buffer to merge secondary index records to indexes in batches. That saves a lot of disk I/O. For rollback, no such mechanism is used, and the rollback can take 30 times longer than the insertion.

In the case of a runaway rollback, if you do not have valuable data in your database, it may be advisable to kill the database process rather than wait for millions of disk I/O operations to complete. For the complete procedure, see [Section 14.2.9.1, “Forcing InnoDB Recovery”](#).

### 14.2.7.3. How `AUTO_INCREMENT` Columns Work in [InnoDB](#)

If you specify an `AUTO_INCREMENT` column for an [InnoDB](#) table, the table handle in the [InnoDB](#) data dictionary contains a special counter called the auto-increment counter that is used in assigning new values for the column. This counter is stored only in main memory, not on disk.

[InnoDB](#) uses the following algorithm to initialize the auto-increment counter for a table `T` that contains an `AUTO_INCREMENT` column named `ai_col`: After a server startup, for the first insert into a table `T`, [InnoDB](#) executes the equivalent of this statement:

```
SELECT MAX(ai_col) FROM T FOR UPDATE;
```

[InnoDB](#) increments by one the value retrieved by the statement and assigns it to the column and to the auto-increment counter for the table. If the table is empty, [InnoDB](#) uses the value `1`. If a user invokes a `SHOW TABLE STATUS` statement that displays output for the table `T` and the auto-increment counter has not been initialized, [InnoDB](#) initializes but does not increment the value and stores it for use by later inserts. Note that this initialization uses a normal exclusive-locking read on the table and the lock lasts to the end of the transaction.

[InnoDB](#) follows the same procedure for initializing the auto-increment counter for a freshly created table.

After the auto-increment counter has been initialized, if a user does not explicitly specify a value for an `AUTO_INCREMENT` column, [InnoDB](#) increments the counter by one and assigns the new value to the column. If the user inserts a row that explicitly specifies the column value, and the value is bigger than the current counter value, the counter is set to the specified column value.

You may see gaps in the sequence of values assigned to the `AUTO_INCREMENT` column if you roll back transactions that have generated numbers using the counter.

If a user specifies `NULL` or `0` for the `AUTO_INCREMENT` column in an `INSERT`, [InnoDB](#) treats the row as if the value had not been specified and generates a new value for it.

The behavior of the auto-increment mechanism is not defined if a user assigns a negative value to the column or if the value becomes bigger than the maximum integer that can be stored in the specified integer type.

When accessing the auto-increment counter, **InnoDB** uses a special table-level **AUTO-INC** lock that it keeps to the end of the current SQL statement, not to the end of the transaction. The special lock release strategy was introduced to improve concurrency for inserts into a table containing an **AUTO\_INCREMENT** column. Nevertheless, two transactions cannot have the **AUTO-INC** lock on the same table simultaneously, which can have a performance impact if the **AUTO-INC** lock is held for a long time. That might be the case for a statement such as `INSERT INTO t1 ... SELECT ... FROM t2` that inserts all rows from one table into another.

**InnoDB** uses the in-memory auto-increment counter as long as the server runs. When the server is stopped and restarted, **InnoDB** reinitializes the counter for each table for the first **INSERT** to the table, as described earlier.

Beginning with MySQL 4.1.12, **InnoDB** supports the **AUTO\_INCREMENT = n** table option in **ALTER TABLE** statements, to set the initial counter value or alter the current counter value. The same is true as of MySQL 4.1.14 for **CREATE TABLE**. The effect of this option is canceled by a server restart, for reasons discussed earlier in this section.

#### 14.2.7.4. FOREIGN KEY Constraints

Starting from MySQL 3.23.44, **InnoDB** features foreign key constraints.

The syntax of a foreign key constraint definition in **InnoDB** looks like this:

```
[CONSTRAINT symbol] FOREIGN KEY [id] (index_col_name, ...)
REFERENCES tbl_name (index_col_name, ...)
[ON DELETE {RESTRICT | CASCADE | SET NULL | NO ACTION}]
[ON UPDATE {RESTRICT | CASCADE | SET NULL | NO ACTION}]
```

Foreign keys definitions are subject to the following conditions:

- Both tables must be **InnoDB** tables and they must not be **TEMPORARY** tables.
- In the referencing table, there must be an index where the foreign key columns are listed as the *first* columns in the same order. Starting with MySQL 4.1.2, such an index is created on the referencing table automatically if it does not exist.
- In the referenced table, there must be an index where the referenced columns are listed as the *first* columns in the same order.
- Index prefixes on foreign key columns are not supported. One consequence of this is that **BLOB** and **TEXT** columns cannot be included in a foreign key, because indexes on those columns must always include a prefix length.
- If the **CONSTRAINT *symbol*** clause is given, the *symbol* value must be unique in the database. If the clause is not given, **InnoDB** creates the name automatically.

**InnoDB** rejects any **INSERT** or **UPDATE** operation that attempts to create a foreign key value in a child table if there is no a matching candidate key value in the parent table. The action **InnoDB** takes for any **UPDATE** or **DELETE** operation that attempts to update or delete a candidate key value in the parent table that has some matching rows in the child table is dependent on the *referential action* specified using **ON UPDATE** and **ON DELETE** subclauses of the **FOREIGN KEY** clause. When the user attempts to delete or update a row from a parent table, and there are one or more matching rows in the child table, **InnoDB** supports five options regarding the action to be taken:

- **CASCADE**: Delete or update the row from the parent table and automatically delete or update the matching rows in the child table. **ON DELETE CASCADE** is supported starting from MySQL

3.23.50 and `ON UPDATE CASCADE` is supported starting from 4.0.8. Between two tables, you should not define several `ON UPDATE CASCADE` clauses that act on the same column in the parent table or in the child table.

- `SET NULL`: Delete or update the row from the parent table and set the foreign key column or columns in the child table to `NULL`. This is valid only if the foreign key columns do not have the `NOT NULL` qualifier specified. `ON DELETE SET NULL` is available starting from MySQL 3.23.50 and `ON UPDATE SET NULL` is available starting from 4.0.8.
- `NO ACTION`: In standard sQL, `NO ACTION` means *no action* in the sense that an attempt to delete or update a primary key value will not be allowed to proceed if there is a related foreign key value in the referenced table. Starting from 4.0.18 `InnoDB` rejects the delete or update operation for the parent table.
- `RESTRICT`: Rejects the delete or update operation for the parent table. `NO ACTION` and `RESTRICT` are the same as omitting the `ON DELETE` or `ON UPDATE` clause. (Some database systems have deferred checks, and `NO ACTION` is a deferred check. In MySQL, foreign key constraints are checked immediately, so `NO ACTION` and `RESTRICT` are the same.)
- `SET DEFAULT`: This action is recognized by the parser, but `InnoDB` rejects table definitions containing `ON DELETE SET DEFAULT` or `ON UPDATE SET DEFAULT` clauses.

Note that `InnoDB` supports foreign key references within a table. In these cases, “child table records” really refers to dependent records within the same table.

`InnoDB` requires indexes on foreign keys and referenced keys so that foreign key checks can be fast and not require a table scan. Starting with MySQL 4.1.2, the index on the foreign key is created automatically. In older versions, the indexes must be created explicitly or the creation of foreign key constraints fails.

Corresponding columns in the foreign key and the referenced key must have similar internal data types inside `InnoDB` so that they can be compared without a type conversion. *The size and sign of integer types must be the same.* The length of string types need not be the same. If you specify a `SET NULL` action, *make sure that you have not declared the columns in the child table as `NOT NULL`.*

If MySQL reports an error number 1005 from a `CREATE TABLE` statement, and the error message refers to `errno 150`, table creation failed because a foreign key constraint was not correctly formed. Similarly, if an `ALTER TABLE` fails and it refers to `errno 150`, that means a foreign key definition would be incorrectly formed for the altered table. Starting from MySQL 4.0.13, you can use `SHOW INNODB STATUS` to display a detailed explanation of the latest `InnoDB` foreign key error in the server.

Starting from MySQL 3.23.50, `InnoDB` does not check foreign key constraints on those foreign key or referenced key values that contain a `NULL` column.

**Note:** Currently, triggers are not activated by cascaded foreign key actions.

You cannot create a table with a column name that matches the name of an internal `InnoDB` column (including `DB_ROW_ID`, `DB_TRX_ID`, `DB_ROLL_PTR` and `DB_MIX_ID`). In versions of MySQL before 4.1.19 this would cause a crash, since 4.1.19 the server will report error 1005 and refers to `errno -1` in the error message.

**Deviation from SQL standards:** If there are several rows in the parent table that have the same referenced key value, `InnoDB` acts in foreign key checks as if the other parent rows with the same key value do not exist. For example, if you have defined a `RESTRICT` type constraint, and there is a child row with several parent rows, `InnoDB` does not allow the deletion of any of those parent rows.

`InnoDB` performs cascading operations through a depth-first algorithm, based on records in the indexes corresponding to the foreign key constraints.



**Deviation from SQL standards:** A `FOREIGN KEY` constraint that references a non-`UNIQUE` key is not standard SQL. It is an `InnoDB` extension to standard SQL.

**Deviation from SQL standards:** If `ON UPDATE CASCADE` or `ON UPDATE SET NULL` recurses to update the *same table* it has previously updated during the cascade, it acts like `RESTRICT`. This means that you cannot use self-referential `ON UPDATE CASCADE` or `ON UPDATE SET NULL` operations. This is to prevent infinite loops resulting from cascaded updates. A self-referential `ON DELETE SET NULL`, on the other hand, is possible from 4.0.13. A self-referential `ON DELETE CASCADE` has been possible since `ON DELETE` was implemented. Since 4.0.21, cascading operations may not be nested more than 15 levels.

**Deviation from SQL standards:** Like MySQL in general, in an SQL statement that inserts, deletes, or updates many rows, `InnoDB` checks `UNIQUE` and `FOREIGN KEY` constraints row-by-row. According to the SQL standard, the default behavior should be deferred checking. That is, constraints are only checked after the **whole SQL statement** has been processed. Until `InnoDB` implements deferred constraint checking, some things will be impossible, such as deleting a record that refers to itself via a foreign key.

Here is a simple example that relates `parent` and `child` tables through a single-column foreign key:

```
CREATE TABLE parent (id INT NOT NULL,
                     PRIMARY KEY (id)
) TYPE=INNODB;
CREATE TABLE child (id INT, parent_id INT,
                    INDEX par_ind (parent_id),
                    FOREIGN KEY (parent_id) REFERENCES parent(id)
                    ON DELETE CASCADE
) TYPE=INNODB;
```

A more complex example in which a `product_order` table has foreign keys for two other tables. One foreign key references a two-column index in the `product` table. The other references a single-column index in the `customer` table:

```
CREATE TABLE product (category INT NOT NULL, id INT NOT NULL,
                      price DECIMAL,
                      PRIMARY KEY(category, id)) TYPE=INNODB;
CREATE TABLE customer (id INT NOT NULL,
                       PRIMARY KEY (id)) TYPE=INNODB;
CREATE TABLE product_order (no INT NOT NULL AUTO_INCREMENT,
                             product_category INT NOT NULL,
                             product_id INT NOT NULL,
                             customer_id INT NOT NULL,
                             PRIMARY KEY(no),
                             INDEX (product_category, product_id),
                             FOREIGN KEY (product_category, product_id)
                             REFERENCES product(category, id)
                             ON UPDATE CASCADE ON DELETE RESTRICT,
                             INDEX (customer_id),
                             FOREIGN KEY (customer_id)
                             REFERENCES customer(id)) TYPE=INNODB;
```

Starting from MySQL 3.23.50, `InnoDB` allows you to add a new foreign key constraint to a table by using `ALTER TABLE`:

```
ALTER TABLE tbl_name
ADD [CONSTRAINT symbol] FOREIGN KEY [id] (index_col_name, ...)
REFERENCES tbl_name (index_col_name, ...)
[ON DELETE {RESTRICT | CASCADE | SET NULL | NO ACTION}]
[ON UPDATE {RESTRICT | CASCADE | SET NULL | NO ACTION}]
```

**Remember to create the required indexes first.** You can also add a self-referential foreign key constraint to a table using `ALTER TABLE`.

Starting from MySQL 4.0.13, `InnoDB` supports the use of `ALTER TABLE` to drop foreign keys:

```
ALTER TABLE tbl_name DROP FOREIGN KEY fk_symbol;
```

If the **FOREIGN KEY** clause included a **CONSTRAINT** name when you created the foreign key, you can refer to that name to drop the foreign key. (A constraint name can be given as of MySQL 4.0.18.) Otherwise, the *fk\_symbol* value is internally generated by InnoDB when the foreign key is created. To find out the symbol when you want to drop a foreign key, use the **SHOW CREATE TABLE** statement. For example:

```
mysql> SHOW CREATE TABLE ibtest11c\G
***** 1. row *****
      Table: ibtest11c
Create Table: CREATE TABLE `ibtest11c` (
  `A` int(11) NOT NULL auto_increment,
  `D` int(11) NOT NULL default '0',
  `B` varchar(200) NOT NULL default '',
  `C` varchar(175) default NULL,
  PRIMARY KEY (`A`,`D`,`B`),
  KEY `B` (`B`,`C`),
  KEY `C` (`C`),
  CONSTRAINT `0_38775` FOREIGN KEY (`A`, `D`)
REFERENCES `ibtest11a` (`A`, `D`)
ON DELETE CASCADE ON UPDATE CASCADE,
  CONSTRAINT `0_38776` FOREIGN KEY (`B`, `C`)
REFERENCES `ibtest11a` (`B`, `C`)
ON DELETE CASCADE ON UPDATE CASCADE
) TYPE=InnoDB CHARSET=latin1
1 row in set (0.01 sec)

mysql> ALTER TABLE ibtest11c DROP FOREIGN KEY `0_38775`;
```

You cannot add a foreign key and drop a foreign key in separate clauses of a single **ALTER TABLE** statement. Separate statements are required.

Starting from MySQL 3.23.50, the InnoDB parser allows table and column identifiers in a **FOREIGN KEY ... REFERENCES ...** clause to be quoted within backticks. (Alternatively, double quotes can be used if the **ANSI\_QUOTES** SQL mode is enabled.) The InnoDB parser also takes into account the setting of the **lower\_case\_table\_names** system variable.

Before MySQL 3.23.50, **ALTER TABLE** or **CREATE INDEX** should not be used in connection with tables that have foreign key constraints or that are referenced in foreign key constraints: Any **ALTER TABLE** removes all foreign key constraints defined for the table. You should not use **ALTER TABLE** with the referenced table, either. Instead, use **DROP TABLE** and **CREATE TABLE** to modify the schema. When MySQL does an **ALTER TABLE** it may internally use **RENAME TABLE**, and that confuses the foreign key constraints that refer to the table. In MySQL, a **CREATE INDEX** statement is processed as an **ALTER TABLE**, so the same considerations apply.

Starting from MySQL 3.23.50, InnoDB returns the foreign key definitions of a table as part of the output of the **SHOW CREATE TABLE** statement:

```
SHOW CREATE TABLE tbl_name;
```

**mysqldump** also produces correct definitions of tables to the dump file, and does not forget about the foreign keys.

You can also display the foreign key constraints for a table like this:

```
SHOW TABLE STATUS FROM db_name LIKE 'tbl_name';
```

The foreign key constraints are listed in the **Comment** column of the output.

When performing foreign key checks, InnoDB sets shared row-level locks on child or parent records it has to look at. InnoDB checks foreign key constraints immediately; the check is not deferred to transac-

tion commit.

To make it easier to reload dump files for tables that have foreign key relationships, `mysqldump` automatically includes a statement in the dump output to set `FOREIGN_KEY_CHECKS` to 0 as of MySQL 4.1.1. This avoids problems with tables having to be reloaded in a particular order when the dump is reloaded. For earlier versions, you can disable the variable manually within `mysql` when loading the dump file like this:

```
mysql> SET FOREIGN_KEY_CHECKS = 0;
mysql> SOURCE dump_file_name;
mysql> SET FOREIGN_KEY_CHECKS = 1;
```

This allows you to import the tables in any order if the dump file contains tables that are not correctly ordered for foreign keys. It also speeds up the import operation. Setting `FOREIGN_KEY_CHECKS` to 0 can also be useful for ignoring foreign key constraints during `LOAD DATA` or `ALTER TABLE` operations. However, even if `FOREIGN_KEY_CHECKS=0`, InnoDB does not allow the creation of a foreign key constraint where a column references a non-matching column type. `FOREIGN_KEY_CHECKS` is available starting from MySQL 3.23.52 and 4.0.3.

InnoDB does not allow you to drop a table that is referenced by a `FOREIGN KEY` constraint, unless you do `SET FOREIGN_KEY_CHECKS=0`. When you drop a table, the constraints that were defined in its create statement are also dropped.

If you re-create a table that was dropped, it must have a definition that conforms to the foreign key constraints referencing it. It must have the right column names and types, and it must have indexes on the referenced keys, as stated earlier. If these are not satisfied, MySQL returns error number 1005 and refers to errno 150 in the error message.

### 14.2.7.5. InnoDB and MySQL Replication

MySQL replication works for InnoDB tables as it does for MyISAM tables. It is also possible to use replication in a way where the storage engine on the slave is not the same as the original storage engine on the master. For example, you can replicate modifications to an InnoDB table on the master to a MyISAM table on the slave.

To set up a new slave for a master, you have to make a copy of the InnoDB tablespace and the log files, as well as the `.frm` files of the InnoDB tables, and move the copies to the slave. If the `innodb_file_per_table` variable is enabled, you must also copy the `.ibd` files as well. For the proper procedure to do this, see Section 14.2.9, “Backing Up and Recovering an InnoDB Database”.

If you can shut down the master or an existing slave, you can take a cold backup of the InnoDB tablespace and log files and use that to set up a slave. To make a new slave without taking down any server you can also use the non-free (commercial) InnoDB Hot Backup tool [http://www.innodb.com/order.html].

There are minor limitations in InnoDB replication:

- `LOAD TABLE FROM MASTER` does not work for InnoDB type tables. There are workarounds: 1) dump the table on the master and import the dump file into the slave, or 2) use `ALTER TABLE tbl_name TYPE=MyISAM` on the master before setting up replication with `LOAD TABLE tbl_name FROM MASTER`, and then use `ALTER TABLE` to alter the master table back to the InnoDB type afterward. However, this should not be done for tables that have foreign key definitions because the definitions will be lost.
- Before MySQL 4.0.6, `SLAVE STOP` did not respect the boundary of a multiple-statement transaction. An incomplete transaction would be rolled back, and the next `SLAVE START` would only execute the remaining part of the half transaction. That would cause replication to fail.

- Before MySQL 4.0.6, a slave crash in the middle of a multiple-statement transaction would cause the same problem as `SLAVE STOP`.
- Before MySQL 4.0.11, replication of the `SET FOREIGN_KEY_CHECKS=0` statement does not work properly.

Most of these limitations can be eliminated by using more recent server versions for which the limitations do not apply.

Transactions that fail on the master do not affect replication at all. MySQL replication is based on the binary log where MySQL writes SQL statements that modify data. A transaction that fails (for example, because of a foreign key violation, or because it is rolled back) is not written to the binary log, so it is not sent to slaves. See [Section 13.4.1](#), “`START TRANSACTION`, `COMMIT`, and `ROLLBACK Syntax`”.

## 14.2.8. Adding and Removing InnoDB Data and Log Files

This section describes what you can do when your InnoDB tablespace runs out of room or when you want to change the size of the log files.

From MySQL 3.23.50 and 4.0.2, the easiest way to increase the size of the InnoDB tablespace is to configure it from the beginning to be auto-extending. Specify the `autoextend` attribute for the last data file in the tablespace definition. Then InnoDB increases the size of that file automatically in 8MB increments when it runs out of space. Starting with MySQL 4.0.24 and 4.1.5, the increment size can be changed by setting the value of the `innodb_autoextend_increment` system variable, which is measured in MB.

Alternatively, you can increase the size of your tablespace by adding another data file. To do this, you have to shut down the MySQL server, change the tablespace configuration to add a new data file to the end of `innodb_data_file_path`, and start the server again.

If your last data file was defined with the keyword `autoextend`, the procedure for reconfiguring the tablespace must take into account the size to which the last data file has grown. Obtain the size of the data file, round it down to the closest multiple of  $1024 \times 1024$  bytes (= 1MB), and specify the rounded size explicitly in `innodb_data_file_path`. Then you can add another data file. Remember that only the last data file in the `innodb_data_file_path` can be specified as auto-extending.

As an example, assume that the tablespace has just one auto-extending data file `ibdata1`:

```
innodb_data_home_dir =
innodb_data_file_path = /ibdata/ibdata1:10M:autoextend
```

Suppose that this data file, over time, has grown to 988MB. Here is the configuration line after modifying the original data file to not be auto-extending and adding another auto-extending data file:

```
innodb_data_home_dir =
innodb_data_file_path = /ibdata/ibdata1:988M;/disk2/ibdata2:50M:autoextend
```

When you add a new file to the tablespace configuration, make sure that it does not exist. InnoDB will create and initialize the file when you restart the server.

Currently, you cannot remove a data file from the tablespace. To decrease the size of your tablespace, use this procedure:

1. Use `mysqldump` to dump all your InnoDB tables.
2. Stop the server.

3. Remove all the existing tablespace files.
4. Configure a new tablespace.
5. Restart the server.
6. Import the dump files.

If you want to change the number or the size of your `InnoDB` log files, you must stop the MySQL server and make sure that it shuts down without errors (to ensure that there is no information for outstanding transactions in the logs). Then copy the old log files into a safe place just in case something went wrong in the shutdown and you need them to recover the tablespace. Delete the old log files from the log file directory, edit `my.cnf` to change the log file configuration, and start the MySQL server again. `mysqld` sees that no log files exist at startup and tells you that it is creating new ones.

## 14.2.9. Backing Up and Recovering an `InnoDB` Database

The key to safe database management is making regular backups.

`InnoDB Hot Backup` is an online backup tool you can use to backup your `InnoDB` database while it is running. `InnoDB Hot Backup` does not require you to shut down your database and it does not set any locks or disturb your normal database processing. `InnoDB Hot Backup` is a non-free (commercial) add-on tool with an annual license fee of €390 per computer on which the MySQL server is run. See the `InnoDB Hot Backup` home page [<http://www.innodb.com/order.html>] for detailed information and screenshots.

If you are able to shut down your MySQL server, you can make a binary backup that consists of all files used by `InnoDB` to manage its tables. Use the following procedure:

1. Shut down your MySQL server and make sure that it shuts down without errors.
2. Copy all your data files (`ibdata` files and `.ibd` files) into a safe place.
3. Copy all your `ib_logfile` files to a safe place.
4. Copy your `my.cnf` configuration file or files to a safe place.
5. Copy all the `.frm` files for your `InnoDB` tables to a safe place.

Replication works with `InnoDB` tables, so you can use MySQL replication capabilities to keep a copy of your database at database sites requiring high availability.

In addition to making binary backups as just described, you should also regularly make dumps of your tables with `mysqldump`. The reason for this is that a binary file might be corrupted without you noticing it. Dumped tables are stored into text files that are human-readable, so spotting table corruption becomes easier. Also, because the format is simpler, the chance for serious data corruption is smaller. `mysqldump` also has a `--single-transaction` option that you can use to make a consistent snapshot without locking out other clients.

To be able to recover your `InnoDB` database to the present from the binary backup just described, you have to run your MySQL server with binary logging turned on. Then you can apply the binary log to the backup database to achieve point-in-time recovery:

```
mysqlbinlog yourhostname-bin.123 | mysql
```

To recover from a crash of your MySQL server, the only requirement is to restart it. [InnoDB](#) automatically checks the logs and performs a roll-forward of the database to the present. [InnoDB](#) automatically rolls back uncommitted transactions that were present at the time of the crash. During recovery, [mysqld](#) displays output something like this:

```
InnoDB: Database was not shut down normally.
InnoDB: Starting recovery from log files...
InnoDB: Starting log scan based on checkpoint at
InnoDB: log sequence number 0 13674004
InnoDB: Doing recovery: scanned up to log sequence number 0 13739520
InnoDB: Doing recovery: scanned up to log sequence number 0 13805056
InnoDB: Doing recovery: scanned up to log sequence number 0 13870592
InnoDB: Doing recovery: scanned up to log sequence number 0 13936128
...
InnoDB: Doing recovery: scanned up to log sequence number 0 20555264
InnoDB: Doing recovery: scanned up to log sequence number 0 20620800
InnoDB: Doing recovery: scanned up to log sequence number 0 20664692
InnoDB: 1 uncommitted transaction(s) which must be rolled back
InnoDB: Starting rollback of uncommitted transactions
InnoDB: Rolling back trx no 16745
InnoDB: Rolling back of trx no 16745 completed
InnoDB: Rollback of uncommitted transactions completed
InnoDB: Starting an apply batch of log records to the database...
InnoDB: Apply batch completed
InnoDB: Started
mysqld: ready for connections
```

If your database gets corrupted or your disk fails, you have to do the recovery from a backup. In the case of corruption, you should first find a backup that is not corrupted. After restoring the base backup, do the recovery from the binary log files using [mysqlbinlog](#) and [mysql](#) to restore the changes performed after the backup was made.

In some cases of database corruption it is enough just to dump, drop, and re-create one or a few corrupt tables. You can use the [CHECK TABLE](#) SQL statement to check whether a table is corrupt, although [CHECK TABLE](#) naturally cannot detect every possible kind of corruption. You can use [innodb\\_tablespace\\_monitor](#) to check the integrity of the file space management inside the tablespace files.

In some cases, apparent database page corruption is actually due to the operating system corrupting its own file cache, and the data on disk may be okay. It is best first to try restarting your computer. Doing so may eliminate errors that appeared to be database page corruption.

### 14.2.9.1. Forcing [InnoDB](#) Recovery

If there is database page corruption, you may want to dump your tables from the database with [SELECT INTO outfile](#). Usually, most of the data obtained in this way is intact. Even so, the corruption may cause [SELECT \\* FROM tbl\\_name](#) statements or [InnoDB](#) background operations to crash or assert, or even to cause [InnoDB](#) roll-forward recovery to crash. Starting from MySQL 3.23.44, there is an [InnoDB](#) variable that you can use to force the [InnoDB](#) storage engine to start up, and you can also prevent background operations from running, so that you are able to dump your tables. For example, you can add the following line to the [\[mysqld\]](#) section of your option file before restarting the server:

```
[mysqld]
innodb_force_recovery = 4
```

Before MySQL 4.0, use this syntax instead:

```
[mysqld]
set-variable = innodb_force_recovery=4
```

The allowable non-zero values for [innodb\\_force\\_recovery](#) follow. A larger number includes all precautions of smaller numbers. If you are able to dump your tables with an option value of at most 4, then you are relatively safe that only some data on corrupt individual pages is lost. A value of 6 is more

drastic because database pages are left in an obsolete state, which in turn may introduce more corruption into B-trees and other database structures.

- 1 (`SRV_FORCE_IGNORE_CORRUPT`)

Let the server run even if it detects a corrupt page. Try to make `SELECT * FROM tbl_name` jump over corrupt index records and pages, which helps in dumping tables.

- 2 (`SRV_FORCE_NO_BACKGROUND`)

Prevent the main thread from running. If a crash would occur during the purge operation, this recovery value prevents it.

- 3 (`SRV_FORCE_NO_TRX_UNDO`)

Do not run transaction rollbacks after recovery.

- 4 (`SRV_FORCE_NO_IBUF_MERGE`)

Prevent also insert buffer merge operations. If they would cause a crash, do not do them. Do not calculate table statistics.

- 5 (`SRV_FORCE_NO_UNDO_LOG_SCAN`)

Do not look at undo logs when starting the database: `InnoDB` treats even incomplete transactions as committed.

- 6 (`SRV_FORCE_NO_LOG_REDO`)

Do not do the log roll-forward in connection with recovery.

Starting from MySQL 3.23.53 and 4.0.4, you can `SELECT` from tables to dump them, or `DROP` or `CREATE` a table even if forced recovery is used. If you know that a certain table is causing a crash in rollback, you can drop it. You can use this also to stop a runaway rollback caused by a failing mass import or `ALTER TABLE`. You can kill the `mysqld` process and set `innodb_force_recovery` to 3 to bring the database up without the rollback, then `DROP` the table that is causing the runaway rollback.

*The database must not otherwise be used with any non-zero value of `innodb_force_recovery`. As a safety measure, `InnoDB` prevents users from performing `INSERT`, `UPDATE`, or `DELETE` operations when `innodb_force_recovery` is greater than 0.*

## 14.2.9.2. Checkpoints

`InnoDB` implements a checkpoint mechanism known as “fuzzy” checkpointing. `InnoDB` flushes modified database pages from the buffer pool in small batches. There is no need to flush the buffer pool in one single batch, which would in practice stop processing of user SQL statements during the checkpointing process.

During crash recovery, `InnoDB` looks for a checkpoint label written to the log files. It knows that all modifications to the database before the label are present in the disk image of the database. Then `InnoDB` scans the log files forward from the checkpoint, applying the logged modifications to the database.

`InnoDB` writes to its log files on a rotating basis. All committed modifications that make the database pages in the buffer pool different from the images on disk must be available in the log files in case `InnoDB` has to do a recovery. This means that when `InnoDB` starts to reuse a log file, it has to make sure that the database page images on disk contain the modifications logged in the log file that `InnoDB` is going to reuse. In other words, `InnoDB` must create a checkpoint and this often involves flushing of

modified database pages to disk.

The preceding description explains why making your log files very large may save disk I/O in check-pointing. It often makes sense to set the total size of the log files as big as the buffer pool or even bigger. The drawback of using large log files is that crash recovery can take longer because there is more logged information to apply to the database.

## 14.2.10. Moving an InnoDB Database to Another Machine

On Windows, InnoDB always stores database and table names internally in lowercase. To move databases in a binary format from Unix to Windows or from Windows to Unix, you should have all table and database names in lowercase. A convenient way to accomplish this is to add the following line to the `[mysqld]` section of your `my.cnf` or `my.ini` file before creating any databases or tables:

```
[mysqld]
set-variable = lower_case_table_names=1
```

Like MyISAM data files, InnoDB data and log files are binary-compatible on all platforms having the same floating-point number format. You can move an InnoDB database simply by copying all the relevant files listed in Section 14.2.9, “Backing Up and Recovering an InnoDB Database”. If the floating-point formats differ but you have not used `FLOAT` or `DOUBLE` data types in your tables, then the procedure is the same: simply copy the relevant files. If the formats differ and your tables contain floating-point data, you must use `mysqldump` to dump your tables on one machine and then import the dump files on the other machine.

One way to increase performance is to switch off autocommit mode when importing data, assuming that the tablespace has enough space for the big rollback segment that the import transactions generate. Do the commit only after importing a whole table or a segment of a table.

## 14.2.11. InnoDB Transaction Model and Locking

In the InnoDB transaction model, the goal is to combine the best properties of a multi-versioning database with traditional two-phase locking. InnoDB does locking on the row level and runs queries as non-locking consistent reads by default, in the style of Oracle. The lock table in InnoDB is stored so space-efficiently that lock escalation is not needed: Typically several users are allowed to lock every row in the database, or any random subset of the rows, without InnoDB running out of memory.

### 14.2.11.1. InnoDB Lock Modes

InnoDB implements standard row-level locking where there are two types of locks:

- A shared (*S*) lock allows a transaction to read a row (tuple).
- An exclusive (*X*) lock allows a transaction to update or delete a row.

If transaction **T1** holds a shared (*S*) lock on tuple **t**, then

- A request from some distinct transaction **T2** for an *S* lock on **t** can be granted immediately. As a result, both **T1** and **T2** hold an *S* lock on **t**.
- A request from some distinct transaction **T2** for an *X* lock on **t** cannot be granted immediately.

If a transaction **T1** holds an exclusive (*X*) lock on tuple **t**, then a request from some distinct transaction **T2** for a lock of either type on **t** cannot be granted immediately. Instead, transaction **T2** has to wait for



transaction **T1** to release its lock on tuple **t**.

Additionally, **InnoDB** supports *multiple granularity locking* which allows coexistence of record locks and locks on entire tables. To make locking at multiple granularity levels practical, additional types of locks called *intention locks* are used. Intention locks are table locks in **InnoDB**. The idea behind intention locks is for a transaction to indicate which type of lock (shared or exclusive) it will require later for a row in that table. There are two types of intention locks used in **InnoDB** (assume that transaction **T** has requested a lock of the indicated type on table **R**):

- Intention shared (**IS**): Transaction **T** intends to set **S** locks on individual rows in table **R**.
- Intention exclusive (**IX**): Transaction **T** intends to set **X** locks on those rows.

The intention locking protocol is as follows:

- Before a given transaction can acquire an **S** lock on a given row, it must first acquire an **IS** or stronger lock on the table containing that row.
- Before a given transaction can acquire an **X** lock on a given row, it must first acquire an **IX** lock on the table containing that row.

These rules can be conveniently summarized by means of a *lock type compatibility matrix*:

	<b>X</b>	<b>IX</b>	<b>S</b>	<b>IS</b>
<b>X</b>	Conflict	Conflict	Conflict	Conflict
<b>IX</b>	Conflict	Compatible	Conflict	Compatible
<b>S</b>	Conflict	Conflict	Compatible	Compatible
<b>IS</b>	Conflict	Compatible	Compatible	Compatible

A lock is granted to a requesting transaction if it is compatible with existing locks. A lock is not granted to a requesting transaction if it conflicts with existing locks. A transaction waits until the conflicting existing lock is released. If a lock request conflicts with an existing lock and cannot be granted because it would cause deadlock, an error occurs.

Thus, intention locks do not block anything except full table requests (for example, **LOCK TABLES ... WRITE**). The main purpose of **IX** and **IS** locks is to show that someone is locking a row, or going to lock a row in the table.

The following example illustrates how an error can occur when a lock request would cause a deadlock. The example involves two clients, A and B.

First, client A creates a table containing one row, and then begins a transaction. Within the transaction, A obtains an **S** lock on the row by selecting it in share mode:

```
mysql> CREATE TABLE t (i INT) ENGINE = InnoDB;
Query OK, 0 rows affected (1.07 sec)

mysql> INSERT INTO t (i) VALUES(1);
Query OK, 1 row affected (0.09 sec)

mysql> START TRANSACTION;
Query OK, 0 rows affected (0.00 sec)

mysql> SELECT * FROM t WHERE i = 1 LOCK IN SHARE MODE;
+-----+
| i     |
+-----+
```

```
| 1 |
+----+
1 row in set (0.10 sec)
```

Next, client B begins a transaction and attempts to delete the row from the table:

```
mysql> START TRANSACTION;
Query OK, 0 rows affected (0.00 sec)

mysql> DELETE FROM t WHERE i = 1;
```

The delete operation requires an *X* lock. The lock cannot be granted because it is incompatible with the *S* lock that client A holds, so the request goes on the queue of lock requests for the row and client B blocks.

Finally, client A also attempts to delete the row from the table:

```
mysql> DELETE FROM t WHERE i = 1;
ERROR 1213 (40001): Deadlock found when trying to get lock;
try restarting transaction
```

Deadlock occurs here because client A needs an *X* lock to delete the row. However, that lock request cannot be granted because client B is already has a request for an *X* lock and is waiting for client A to release its *S* lock. Nor can the *S* lock held by A be upgraded to an *X* lock because of the prior request by B for an *X* lock. As a result, *InnoDB* generates an error for client A and releases its locks. At that point, the lock request for client B can be granted and B deletes the row from the table.

### 14.2.11.2. *InnoDB* and AUTOCOMMIT

In *InnoDB*, all user activity occurs inside a transaction. If the autocommit mode is enabled, each SQL statement forms a single transaction on its own. By default, MySQL starts new connections with autocommit enabled.

If the autocommit mode is switched off with `SET AUTOCOMMIT = 0`, then we can consider that a user always has a transaction open. An SQL `COMMIT` or `ROLLBACK` statement ends the current transaction and a new one starts. A `COMMIT` means that the changes made in the current transaction are made permanent and become visible to other users. A `ROLLBACK` statement, on the other hand, cancels all modifications made by the current transaction. Both statements release all *InnoDB* locks that were set during the current transaction.

If the connection has autocommit enabled, the user can still perform a multiple-statement transaction by starting it with an explicit `START TRANSACTION` or `BEGIN` statement and ending it with `COMMIT` or `ROLLBACK`.

### 14.2.11.3. *InnoDB* and TRANSACTION ISOLATION LEVEL

In terms of the SQL:1992 transaction isolation levels, the *InnoDB* default is `REPEATABLE READ`. Starting from MySQL 4.0.5, *InnoDB* offers all four different transaction isolation levels described by the SQL standard. You can set the default isolation level for all connections by using the `--transaction-isolation` option on the command line or in an option file. For example, you can set the option in the `[mysqld]` section of an option file like this:

```
[mysqld]
transaction-isolation = {READ-UNCOMMITTED | READ-COMMITTED
                        | REPEATABLE-READ | SERIALIZABLE}
```

A user can change the isolation level for a single session or for all new incoming connections with the `SET TRANSACTION` statement. Its syntax is as follows:

```
SET [SESSION | GLOBAL] TRANSACTION ISOLATION LEVEL
```

```
{READ UNCOMMITTED | READ COMMITTED
 | REPEATABLE READ | SERIALIZABLE}
```

Note that there are hyphens in the level names for the `--transaction-isolation` option, but not for the `SET TRANSACTION` statement.

The default behavior is to set the isolation level for the next (not started) transaction. If you use the `GLOBAL` keyword, the statement sets the default transaction level globally for all new connections created from that point on (but not for existing connections). You need the `SUPER` privilege to do this. Using the `SESSION` keyword sets the default transaction level for all future transactions performed on the current connection.

Any client is free to change the session isolation level (even in the middle of a transaction), or the isolation level for the next transaction.

Before MySQL 3.23.50, `SET TRANSACTION` had no effect on `InnoDB` tables. Before 4.0.5, only `REPEATABLE READ` and `SERIALIZABLE` were available.

You can determine the global and session transaction isolation levels by checking the value of the `tx_isolation` system variable with these statements:

```
SELECT @@global.tx_isolation;
SELECT @@tx_isolation;
```

In row-level locking, `InnoDB` uses next-key locking. That means that besides index records, `InnoDB` can also lock the “gap” preceding an index record to block insertions by other users immediately before the index record. A next-key lock refers to a lock that locks an index record and the gap before it. A gap lock refers to a lock that only locks a gap before some index record.

A detailed description of each isolation level in `InnoDB` follows:

- `READ UNCOMMITTED`

`SELECT` statements are performed in a non-locking fashion, but a possible earlier version of a record might be used. Thus, using this isolation level, such reads are not consistent. This is also called a “dirty read.” Otherwise, this isolation level works like `READ COMMITTED`.

- `READ COMMITTED`

A somewhat Oracle-like isolation level. All `SELECT ... FOR UPDATE` and `SELECT ... LOCK IN SHARE MODE` statements lock only the index records, not the gaps before them, and thus allow the free insertion of new records next to locked records. `UPDATE` and `DELETE` statements using a unique index with a unique search condition lock only the index record found, not the gap before it. In range-type `UPDATE` and `DELETE` statements, `InnoDB` must set next-key or gap locks and block insertions by other users to the gaps covered by the range. This is necessary because “phantom rows” must be blocked for MySQL replication and recovery to work.

Consistent reads behave as in Oracle: Each consistent read, even within the same transaction, sets and reads its own fresh snapshot. See [Section 14.2.11.4, “Consistent Non-Locking Read”](#).

- `REPEATABLE READ`

This is the default isolation level of `InnoDB`. `SELECT ... FOR UPDATE`, `SELECT ... LOCK IN SHARE MODE`, `UPDATE`, and `DELETE` statements that use a unique index with a unique search condition lock only the index record found, not the gap before it. With other search conditions, these operations employ next-key locking, locking the index range scanned with next-key or gap locks, and block new insertions by other users.

In consistent reads, there is an important difference from the `READ COMMITTED` isolation level: All consistent reads within the same transaction read the same snapshot established by the first read. This convention means that if you issue several plain `SELECT` statements within the same transaction, these `SELECT` statements are consistent also with respect to each other. See [Section 14.2.11.4, “Consistent Non-Locking Read”](#).

- `SERIALIZABLE`

This level is like `REPEATABLE READ`, but `InnoDB` implicitly commits all plain `SELECT` statements to `SELECT ... LOCK IN SHARE MODE`.

#### 14.2.11.4. Consistent Non-Locking Read

A consistent read means that `InnoDB` uses multi-versioning to present to a query a snapshot of the database at a point in time. The query sees the changes made by those transactions that committed before that point of time, and no changes made by later or uncommitted transactions. The exception to this rule is that the query sees the changes made by earlier statements within the same transaction. Note that the exception to the rule causes the following anomaly: if you update some rows in a table, a `SELECT` will see the latest version of the updated rows, while it sees the old version of other rows. If other users simultaneously update the same table, the anomaly means that you may see the table in a state that never existed in the database.

If you are running with the default `REPEATABLE READ` isolation level, all consistent reads within the same transaction read the snapshot established by the first such read in that transaction. You can get a fresher snapshot for your queries by committing the current transaction and after that issuing new queries.

Consistent read is the default mode in which `InnoDB` processes `SELECT` statements in `READ COMMITTED` and `REPEATABLE READ` isolation levels. A consistent read does not set any locks on the tables it accesses, and therefore other users are free to modify those tables at the same time a consistent read is being performed on the table.

Note that consistent read does not work over `DROP TABLE` and over `ALTER TABLE`. Consistent read does not work over `DROP TABLE` because MySQL can't use a table that has been dropped and `InnoDB` destroys the table. Consistent read does not work over `ALTER TABLE` because `ALTER TABLE` works by making a temporary copy of the original table and deleting the original table when the temporary copy is built. When you reissue a consistent read within a transaction, rows in the new table are not visible because those rows did not exist when the transaction's snapshot was taken.

#### 14.2.11.5. `SELECT ... FOR UPDATE` and `SELECT ... LOCK IN SHARE MODE` Locking Reads

In some circumstances, a consistent read is not convenient. For example, you might want to add a new row into your table `child`, and make sure that the child has a parent in table `parent`. The following example shows how to implement referential integrity in your application code.

Suppose that you use a consistent read to read the table `parent` and indeed see the parent of the child in the table. Can you safely add the child row to table `child`? No, because it may happen that meanwhile some other user deletes the parent row from the table `parent` without you being aware of it.

The solution is to perform the `SELECT` in a locking mode using `LOCK IN SHARE MODE`:

```
SELECT * FROM parent WHERE NAME = 'Jones' LOCK IN SHARE MODE;
```

Performing a read in share mode means that we read the latest available data, and set a shared mode lock

on the rows we read. A shared mode lock prevents others from updating or deleting the row we have read. Also, if the latest data belongs to a yet uncommitted transaction of another client connection, we wait until that transaction commits. After we see that the preceding query returns the parent 'Jones', we can safely add the child record to the `child` table and commit our transaction.

Let us look at another example: We have an integer counter field in a table `child_codes` that we use to assign a unique identifier to each child added to table `child`. Obviously, using a consistent read or a shared mode read to read the present value of the counter is not a good idea because two users of the database may then see the same value for the counter, and a duplicate-key error occurs if two users attempt to add children with the same identifier to the table.

Here, `LOCK IN SHARE MODE` is not a good solution because if two users read the counter at the same time, at least one of them ends up in deadlock when attempting to update the counter.

In this case, there are two good ways to implement the reading and incrementing of the counter: (1) update the counter first by incrementing it by 1 and only after that read it, or (2) read the counter first with a lock mode `FOR UPDATE`, and increment after that. The latter approach can be implemented as follows:

```
SELECT counter_field FROM child_codes FOR UPDATE;
UPDATE child_codes SET counter_field = counter_field + 1;
```

A `SELECT ... FOR UPDATE` reads the latest available data, setting exclusive locks on each row it reads. Thus, it sets the same locks a searched SQL `UPDATE` would set on the rows.

The preceding description is merely an example of how `SELECT ... FOR UPDATE` works. In MySQL, the specific task of generating a unique identifier actually can be accomplished using only a single access to the table:

```
UPDATE child_codes SET counter_field = LAST_INSERT_ID(counter_field + 1);
SELECT LAST_INSERT_ID();
```

The `SELECT` statement merely retrieves the identifier information (specific to the current connection). It does not access any table.

Locks set by `IN SHARE MODE` and `FOR UPDATE` reads are released when the transaction is committed or rolled back.

### 14.2.11.6. Next-Key Locking: Avoiding the Phantom Problem

In row-level locking, `InnoDB` uses an algorithm called *next-key locking*. `InnoDB` performs the row-level locking in such a way that when it searches or scans an index of a table, it sets shared or exclusive locks on the index records it encounters. Thus, the row-level locks are actually index record locks.

The locks `InnoDB` sets on index records also affect the “gap” before that index record. If a user has a shared or exclusive lock on record `R` in an index, another user cannot insert a new index record immediately before `R` in the index order. This locking of gaps is done to prevent the so-called “phantom problem.” Suppose that you want to read and lock all children from the `child` table having an identifier value greater than 100, with the intention of updating some column in the selected rows later:

```
SELECT * FROM child WHERE id > 100 FOR UPDATE;
```

Suppose that there is an index on the `id` column. The query scans that index starting from the first record where `id` is bigger than 100. If the locks set on the index records would not lock out inserts made in the gaps, a new row might meanwhile be inserted to the table. If you execute the same `SELECT` within the same transaction, you would see a new row in the result set returned by the query. This is contrary to the isolation principle of transactions: A transaction should be able to run so that the data it has read

does not change during the transaction. If we regard a set of rows as a data item, the new “phantom” child would violate this isolation principle.

When `InnoDB` scans an index, it can also lock the gap after the last record in the index. Just that happens in the previous example: The locks set by `InnoDB` prevent any insert to the table where `id` would be bigger than 100.

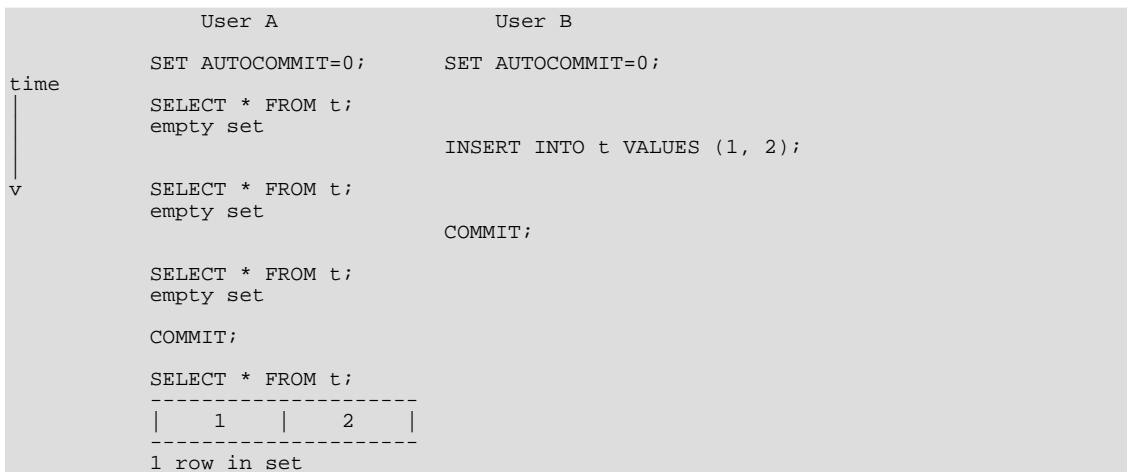
You can use next-key locking to implement a uniqueness check in your application: If you read your data in share mode and do not see a duplicate for a row you are going to insert, then you can safely insert your row and know that the next-key lock set on the successor of your row during the read prevents anyone meanwhile inserting a duplicate for your row. Thus, the next-key locking allows you to “lock” the non-existence of something in your table.

### 14.2.11.7. An Example of Consistent Read in `InnoDB`

Suppose that you are running in the default `REPEATABLE READ` isolation level. When you issue a consistent read (that is, an ordinary `SELECT` statement), `InnoDB` gives your transaction a timepoint according to which your query sees the database. If another transaction deletes a row and commits after your timepoint was assigned, you do not see the row as having been deleted. Inserts and updates are treated similarly.

You can advance your timepoint by committing your transaction and then doing another `SELECT`.

This is called *multi-versioned concurrency control*.



In this example, user A sees the row inserted by B only when B has committed the insert and A has committed as well, so that the timepoint is advanced past the commit of B.

If you want to see the “freshest” state of the database, you should use either the `READ COMMITTED` isolation level or a locking read:

```
SELECT * FROM t LOCK IN SHARE MODE;
```

### 14.2.11.8. Locks Set by Different SQL Statements in `InnoDB`

A locking read, an `UPDATE`, or a `DELETE` generally set record locks on every index record that is scanned in the processing of the SQL statement. It does not matter if there are `WHERE` conditions in the statement that would exclude the row. `InnoDB` does not remember the exact `WHERE` condition, but only knows which index ranges were scanned. The record locks are normally next-key locks that also block inserts to the “gap” immediately before the record.

If the locks to be set are exclusive, `InnoDB` always retrieves also the clustered index record and sets a lock on it.

If you do not have indexes suitable for your statement and MySQL has to scan the whole table to process the statement, every row of the table becomes locked, which in turn blocks all inserts by other users to the table. It is important to create good indexes so that your queries do not unnecessarily need to scan many rows.

`InnoDB` sets specific types of locks as follows:

- `SELECT ... FROM` is a consistent read, reading a snapshot of the database and setting no locks unless the transaction isolation level is set to `SERIALIZABLE`. For `SERIALIZABLE` level, this sets shared next-key locks on the index records it encounters.
- `SELECT ... FROM ... LOCK IN SHARE MODE` sets shared next-key locks on all index records the read encounters.
- `SELECT ... FROM ... FOR UPDATE` sets exclusive next-key locks on all index records the read encounters.
- `INSERT INTO ... VALUES (...)` sets an exclusive lock on the inserted row. Note that this lock is not a next-key lock and does not prevent other users from inserting to the gap before the inserted row. If a duplicate-key error occurs, a shared lock on the duplicate index record is set.
- While initializing a previously specified `AUTO_INCREMENT` column on a table, `InnoDB` sets an exclusive lock on the end of the index associated with the `AUTO_INCREMENT` column. In accessing the auto-increment counter, `InnoDB` uses a specific table lock mode `AUTO-INC` where the lock lasts only to the end of the current SQL statement, not to the end of the entire transaction. Note that other clients cannot insert into the table while the `AUTO-INC` table lock is held; see [Section 14.2.11.2, “InnoDB and AUTOCOMMIT”](#).

Before MySQL 3.23.50, `SHOW TABLE STATUS` applied to a table with an `AUTO_INCREMENT` column sets an exclusive row-level lock to the high end of the `AUTO_INCREMENT` index. This means also that `SHOW TABLE STATUS` could cause a deadlock of transactions, something that may surprise users. Starting from MySQL 3.23.50, `InnoDB` fetches the value of a previously initialized `AUTO_INCREMENT` column without setting any locks.

- `INSERT INTO T SELECT ... FROM S WHERE ...` sets an exclusive (non-next-key) lock on each row inserted into `T`. It does the search on `S` as a consistent read, but sets shared next-key locks on `S` if MySQL binary logging is turned on. `InnoDB` has to set locks in the latter case: In roll-forward recovery from a backup, every SQL statement has to be executed in exactly the same way it was done originally.
- `CREATE TABLE ... SELECT ...` performs the `SELECT` as a consistent read or with shared locks, as in the previous item.
- `REPLACE` is done like an insert if there is no collision on a unique key. Otherwise, an exclusive next-key lock is placed on the row that has to be updated.
- `UPDATE ... WHERE ...` sets an exclusive next-key lock on every record the search encounters.
- `DELETE FROM ... WHERE ...` sets an exclusive next-key lock on every record the search encounters.
- If a `FOREIGN KEY` constraint is defined on a table, any insert, update, or delete that requires the constraint condition to be checked sets shared record-level locks on the records that it looks at to check the constraint. `InnoDB` also sets these locks in the case where the constraint fails.

- `LOCK TABLES` sets table locks, but it is the higher MySQL layer above the `InnoDB` layer that sets these locks. Beginning with MySQL 4.0.20 and 4.1.2, `InnoDB` is aware of table locks if `innodb_table_locks=1` (the default) and `AUTOCOMMIT=0`, and the MySQL layer above `InnoDB` knows about row-level locks. Otherwise, `InnoDB`'s automatic deadlock detection cannot detect deadlocks where such table locks are involved. Also, because the higher MySQL layer does not know about row-level locks, it is possible to get a table lock on a table where another user currently has row-level locks. However, this does not endanger transaction integrity, as discussed in [Section 14.2.11.10, “Deadlock Detection and Rollback”](#). See also [Section 14.2.17, “Restrictions on InnoDB Tables”](#).

### 14.2.11.9. Implicit Transaction Commit and Rollback

By default, MySQL begins each client connection with autocommit mode enabled. When autocommit is enabled, MySQL does a commit after each SQL statement if that statement did not return an error. If an SQL statement returns an error, the commit or rollback behavior depends on the error. See [Section 14.2.16, “InnoDB Error Handling”](#).

If you have the autocommit mode off and close a connection without explicitly committing the final transaction, MySQL rolls back that transaction.

Each of the following statements (and any synonyms for them) implicitly end a transaction, as if you had done a `COMMIT` before executing the statement:

- `ALTER TABLE`, `BEGIN`, `CREATE DATABASE`, `CREATE INDEX`, `CREATE TABLE`, `DROP DATABASE`, `DROP INDEX`, `DROP TABLE`, `LOAD MASTER DATA`, `LOCK TABLES`, `RENAME TABLE`, `SET AUTOCOMMIT=1`, `START TRANSACTION`, `TRUNCATE TABLE`, `UNLOCK TABLES`.
- `UNLOCK TABLES` commits a transaction only if any tables currently are locked.
- Prior to MySQL 4.0.13, `CREATE TABLE` commits a transaction if the binary update log is enabled. The `CREATE TABLE`, `CREATE DATABASE DROP DATABASE`, and `TRUNCATE TABLE` statements cause an implicit commit beginning with MySQL 4.1.13.
- The `CREATE TABLE` statement in `InnoDB` is processed as a single transaction. This means that a `ROLLBACK` from the user does not undo `CREATE TABLE` statements the user made during that transaction.

Transactions cannot be nested. This is a consequence of the implicit `COMMIT` performed for any current transaction when you issue a `START TRANSACTION` statement or one of its synonyms.

### 14.2.11.10. Deadlock Detection and Rollback

`InnoDB` automatically detects a deadlock of transactions and rolls back a transaction or transactions to break the deadlock. Starting from MySQL 4.0.5, `InnoDB` tries to pick small transactions to roll back, the size of a transaction being determined by the number of rows inserted, updated, or deleted. Prior to 4.0.5, `InnoDB` always rolled back the transaction whose lock request was the last one to build a deadlock, that is, a cycle in the “waits-for” graph of transactions.

Beginning with MySQL 4.0.20 and 4.1.2, `InnoDB` is aware of table locks if `innodb_table_locks=1` (the default) and `AUTOCOMMIT=0`, and the MySQL layer above `InnoDB` knows about row-level locks. Before that, `InnoDB` cannot detect deadlocks where a table lock set by a MySQL `LOCK TABLES` statement is involved, or if a lock set by another storage engine than `InnoDB` is involved. You have to resolve these situations by setting the value of the `innodb_lock_wait_timeout` system variable.



When `InnoDB` performs a complete rollback of a transaction, all locks set by the transaction are released. However, if just a single SQL statement is rolled back as a result of an error, some of the locks set by the statement may be preserved. This happens because `InnoDB` stores row locks in a format such that it cannot know afterward which lock was set by which statement.

### 14.2.11.11. How to Cope with Deadlocks

Deadlocks are a classic problem in transactional databases, but they are not dangerous unless they are so frequent that you cannot run certain transactions at all. Normally, you must write your applications so that they are always prepared to re-issue a transaction if it gets rolled back because of a deadlock.

`InnoDB` uses automatic row-level locking. You can get deadlocks even in the case of transactions that just insert or delete a single row. That is because these operations are not really “atomic”; they automatically set locks on the (possibly several) index records of the row inserted or deleted.

You can cope with deadlocks and reduce the likelihood of their occurrence with the following techniques:

- Use `SHOW INNODB STATUS` to determine the cause of the latest deadlock. That can help you to tune your application to avoid deadlocks. This strategy can be used as of MySQL 3.23.52 and 4.0.3, depending on your MySQL series. From 4.1.2 on, use `SHOW ENGINE INNODB STATUS`.
- Always be prepared to re-issue a transaction if it fails due to deadlock. Deadlocks are not dangerous. Just try again.
- Commit your transactions often. Small transactions are less prone to collision.
- If you are using locking reads (`SELECT ... FOR UPDATE` or `... LOCK IN SHARE MODE`), try using a lower isolation level such as `READ COMMITTED`.
- Access your tables and rows in a fixed order. Then transactions form well-defined queues and do not deadlock.
- Add well-chosen indexes to your tables. Then your queries need to scan fewer index records and consequently set fewer locks. Use `EXPLAIN SELECT` to determine which indexes the MySQL server regards as the most appropriate for your queries.
- Use less locking. If you can afford to allow a `SELECT` to return data from an old snapshot, do not add the clause `FOR UPDATE` or `LOCK IN SHARE MODE` to it. Using the `READ COMMITTED` isolation level is good here, because each consistent read within the same transaction reads from its own fresh snapshot.
- If nothing else helps, serialize your transactions with table-level locks. The correct way to use `LOCK TABLES` with transactional tables, such as `InnoDB` tables, is to set `AUTOCOMMIT = 0` and not to call `UNLOCK TABLES` until after you commit the transaction explicitly. For example, if you need to write to table `t1` and read from table `t2`, you can do this:

```
SET AUTOCOMMIT=0;
LOCK TABLES t1 WRITE, t2 READ, ...;
... do something with tables t1 and t2 here ...
COMMIT;
UNLOCK TABLES;
```

Table-level locks make your transactions queue nicely, and deadlocks are avoided.

- Another way to serialize transactions is to create an auxiliary “semaphore” table that contains just a single row. Have each transaction update that row before accessing other tables. In that way, all transactions happen in a serial fashion. Note that the `InnoDB` instant deadlock detection algorithm

also works in this case, because the serializing lock is a row-level lock. With MySQL table-level locks, the timeout method must be used to resolve deadlocks.

- In applications that use `AUTOCOMMIT=1` and MySQL's `LOCK TABLES` command, InnoDB's internal table locks that were present from 4.0.20 to 4.0.23 can cause deadlocks. Starting from 4.0.22, you can set `innodb_table_locks=0` in `my.cnf` to fall back to the old behavior and remove the problem. 4.0.24 does not set InnoDB table locks if `AUTOCOMMIT=1`.

## 14.2.12. InnoDB Performance Tuning Tips

- In InnoDB, having a long `PRIMARY KEY` wastes a lot of disk space because its value must be stored with every secondary index record. (See [Section 14.2.14, “InnoDB Table and Index Structures”](#).) Create an `AUTO_INCREMENT` column as the primary key if your primary key is long.
- If the Unix `top` tool or the Windows Task Manager shows that the CPU usage percentage with your workload is less than 70%, your workload is probably disk-bound. Maybe you are making too many transaction commits, or the buffer pool is too small. Making the buffer pool bigger can help, but do not set it equal to more than 80% of physical memory.
- Wrap several modifications into one transaction. InnoDB must flush the log to disk at each transaction commit if that transaction made modifications to the database. The rotation speed of a disk is typically at most 167 revolutions/second, which constrains the number of commits to the same 167<sup>th</sup> of a second if the disk does not “fool” the operating system.
- If you can afford the loss of some of the latest committed transactions if a crash occurs, you can set the `innodb_flush_log_at_trx_commit` parameter to 0. InnoDB tries to flush the log once per second anyway, although the flush is not guaranteed.
- Make your log files big, even as big as the buffer pool. When InnoDB has written the log files full, it has to write the modified contents of the buffer pool to disk in a checkpoint. Small log files cause many unnecessary disk writes. The drawback of big log files is that the recovery time is longer.
- Make the log buffer quite large as well (on the order of 8MB).
- Use the `VARCHAR` data type instead of `CHAR` if you are storing variable-length strings or if the column may contain many `NULL` values. A `CHAR(N)` column always takes `N` characters to store data, even if the string is shorter or its value is `NULL`. Smaller tables fit better in the buffer pool and reduce disk I/O.
- (Relevant from 3.23.39 up.) In some versions of GNU/Linux and Unix, flushing files to disk with the Unix `fsync()` call (which InnoDB uses by default) and other similar methods is surprisingly slow. If you are dissatisfied with database write performance, you might try setting the `innodb_flush_method` parameter to `O_DSYNC`. Although `O_DSYNC` seems to be slower on most systems, yours might not be one of them.
- (Verified using MySQL 4.1, assumed for other MySQL versions, given that this is a platform architecture issue.) When using the InnoDB storage engine on Solaris 10 for x86\_64 architecture (AMD Opteron), it is important to mount any filesystems used for storing InnoDB-related files using the `forcedirectio` option. (The default on Solaris 10/x86\_64 is *not* to use this option.) Failure to use `forcedirectio` causes a serious degradation of InnoDB's speed and performance on this platform.

When using the InnoDB storage engine with a large `innodb_buffer_pool_size` value on any release of Solaris 2.6 and up and any platform (sparc/x86/x64/amd64), a significant performance gain can be achieved by placing InnoDB data files and log files on raw devices or on a separate direct I/O UFS filesystem (using mount option `forcedirectio`; see `mount_ufs(1M)`). Users of

the Veritas filesystem VxFS should use the mount option `convosync=direct`.

Other MySQL data files, such as those for `MyISAM` tables, should not be placed on a direct I/O filesystem. Executables or libraries *must not* be placed on a direct I/O filesystem.

- When importing data into `InnoDB`, make sure that MySQL does not have autocommit mode enabled because that requires a log flush to disk for every insert. To disable autocommit during your import operation, surround it with `SET AUTOCOMMIT` and `COMMIT` statements:

```
SET AUTOCOMMIT=0;
... SQL import statements ...
COMMIT;
```

If you use the `mysqldump` option `--opt`, you get dump files that are fast to import into an `InnoDB` table, even without wrapping them with the `SET AUTOCOMMIT` and `COMMIT` statements.

- Beware of big rollbacks of mass inserts: `InnoDB` uses the insert buffer to save disk I/O in inserts, but no such mechanism is used in a corresponding rollback. A disk-bound rollback can take 30 times as long to perform as the corresponding insert. Killing the database process does not help because the rollback starts again on server startup. The only way to get rid of a runaway rollback is to increase the buffer pool so that the rollback becomes CPU-bound and runs fast, or to use a special procedure. See [Section 14.2.9.1, “Forcing InnoDB Recovery”](#).
- Beware also of other big disk-bound operations. Use `DROP TABLE` and `CREATE TABLE` to empty a table, not `DELETE FROM tbl_name`.
- Use the multiple-row `INSERT` syntax to reduce communication overhead between the client and the server if you need to insert many rows:

```
INSERT INTO yourtable VALUES (1,2), (5,5), ...;
```

This tip is valid for inserts into any table, not just `InnoDB` tables.

- If you have `UNIQUE` constraints on secondary keys, starting from MySQL 3.23.52 and 4.0.3, you can speed up table imports by temporarily turning off the uniqueness checks during the import session:

```
SET UNIQUE_CHECKS=0;
... import operation ...
SET UNIQUE_CHECKS=1;
```

For big tables, this saves a lot of disk I/O because `InnoDB` can use its insert buffer to write secondary index records in a batch. Be certain that the data contains no duplicate keys. `UNIQUE_CHECKS` allows but does not require storage engines to ignore duplicate keys.

- If you have `FOREIGN KEY` constraints in your tables, starting from MySQL 3.23.52 and 4.0.3, you can speed up table imports by turning the foreign key checks off for a while in the import session:

```
SET FOREIGN_KEY_CHECKS=0;
... import operation ...
SET FOREIGN_KEY_CHECKS=1;
```

For big tables, this can save a lot of disk I/O.

- If you often have recurring queries for tables that are not updated frequently, use the query cache available as of MySQL 4.0:

```
[mysqld]
query_cache_type = ON
query_cache_size = 10M
```

In MySQL 4.0, the query cache works only with autocommit enabled. This restriction is removed in MySQL 4.1.1 and up.

- Unlike `MyISAM`, `InnoDB` does not store an index cardinality value in its tables. Instead, `InnoDB` computes a cardinality for a table the first time it accesses it after startup. With a large number of tables, this might take significant time. It is the initial table open operation that is important, so to “warm up” a table for later use, you might want to use it immediately after start up by issuing a statement such as `SELECT 1 FROM tbl_name LIMIT 1`.

### 14.2.12.1. `SHOW ENGINE INNODB STATUS` and the `InnoDB` Monitors

Starting from MySQL 3.23.42, `InnoDB` includes `InnoDB` Monitors that print information about the `InnoDB` internal state. Starting from MySQL 3.23.52 and 4.0.3, you can use the `SHOW INNODB STATUS` SQL statement at any time to fetch the output of the standard `InnoDB` Monitor to your SQL client. (In MySQL 4.1.2 and up, the statement is `SHOW ENGINE INNODB STATUS`.) The information is useful in performance tuning. If you are using the `mysql` interactive SQL client, the output is more readable if you replace the usual semicolon statement terminator by `\G`: For a discussion about the `InnoDB` lock modes, see [Section 14.2.11.1, “InnoDB Lock Modes”](#).

```
mysql> SHOW INNODB STATUS\G
```

Another way to use `InnoDB` Monitors is to let them periodically write data to the standard output of the `mysqld` server. In this case, no output is sent to clients. When switched on, `InnoDB` Monitors print data about every 15 seconds. Server output usually is directed to the `.err` log in the MySQL data directory. This data is useful in performance tuning. On Windows, you must start the server from a command prompt in a console window with the `--console` option if you want to direct the output to the window rather than to the error log.

Monitor output includes the following types of information:

- Table and record locks held by each active transaction
- Lock waits of a transactions
- Semaphore waits of threads
- Pending file I/O requests
- Buffer pool statistics
- Purge and insert buffer merge activity of the main `InnoDB` thread

To cause the standard `InnoDB` Monitor to write to the standard output of `mysqld`, use the following SQL statement:

```
CREATE TABLE innodb_monitor (a INT) TYPE=InnoDB;
```

The monitor can be stopped by issuing the following statement:

```
DROP TABLE innodb_monitor;
```

The `CREATE TABLE` syntax is just a way to pass a command to the `InnoDB` engine through MySQL's

SQL parser: The only things that matter are the table name `innodb_monitor` and that it be an `InnoDB` table. The structure of the table is not relevant at all for the `InnoDB` Monitor. If you shut down the server, the monitor does not restart automatically when you restart the server. You must drop the monitor table and issue a new `CREATE TABLE` statement to start the monitor. (This syntax may change in a future release.)

In a similar way, you can start `innodb_lock_monitor`, which is otherwise the same as `innodb_monitor` but also prints a lot of lock information. A separate `innodb_tablespace_monitor` prints a list of created file segments existing in the tablespace and also validates the tablespace allocation data structures. Starting from 3.23.44, there is `innodb_table_monitor` with which you can print the contents of the `InnoDB` internal data dictionary.

A sample of `InnoDB` Monitor output:

```
mysql> SHOW INNODB STATUS\G
***** 1. row *****
Status:
=====
030709 13:00:59 INNODB MONITOR OUTPUT
=====
Per second averages calculated from the last 18 seconds
-----
SEMAPHORES
-----
OS WAIT ARRAY INFO: reservation count 413452, signal count 378357
--Thread 32782 has waited at btr0sea.c line 1477 for 0.00 seconds the
semaphore: X-lock on RW-latch at 41a28668 created in file btr0sea.c line 135
a writer (thread id 32782) has reserved it in mode wait exclusive
number of readers 1, waiters flag 1
Last time read locked in file btr0sea.c line 731
Last time write locked in file btr0sea.c line 1347
Mutex spin waits 0, rounds 0, OS waits 0
RW-shared spins 108462, OS waits 37964; RW-excl spins 681824, OS waits
375485
-----
LATEST FOREIGN KEY ERROR
-----
030709 13:00:59 Transaction:
TRANSACTION 0 290328284, ACTIVE 0 sec, process no 3195, OS thread id 34831
inserting
15 lock struct(s), heap size 2496, undo log entries 9
MySQL thread id 25, query id 4668733 localhost heikki update
insert into ibtest11a (D, B, C) values (5, 'khDk', 'khDk')
Foreign key constraint fails for table test/ibtest11a:
'
  CONSTRAINT `0_219242` FOREIGN KEY (`A`, `D`) REFERENCES `ibtest11b` (`A`,
  `D`) ON DELETE CASCADE ON UPDATE CASCADE
Trying to add in child table, in index PRIMARY tuple:
  0: len 4; hex 80000101; asc ....;; 1: len 4; hex 80000005; asc ....;; 2:
  len 4; hex 6b68446b; asc khDk;; 3: len 6; hex 0000114e0edc; asc ...N.;; 4:
  len 7; hex 00000000c3e0a7; asc .....;; 5: len 4; hex 6b68446b; asc khDk;;
But in parent table test/ibtest11b, in index PRIMARY,
the closest match we can find is record:
RECORD: info bits 0 0: len 4; hex 8000015b; asc ...[;; 1: len 4; hex
80000005; asc ....;; 2: len 3; hex 6b6864; asc khD;; 3: len 6; hex
0000111ef3eb; asc .....;; 4: len 7; hex 800001001e0084; asc .....;; 5:
len 3; hex 6b6864; asc khD;;
-----
LATEST DETECTED DEADLOCK
-----
030709 12:59:58
*** (1) TRANSACTION:
TRANSACTION 0 290252780, ACTIVE 1 sec, process no 3185, OS thread id 30733
inserting
LOCK WAIT 3 lock struct(s), heap size 320, undo log entries 146
MySQL thread id 21, query id 4553379 localhost heikki update
INSERT INTO alex1 VALUES(86, 86, 794,'aA35818','bb','c79166','d4766t',
'e187358f','g84586','h794',date_format('2001-04-03 12:54:22','%Y-%m-%d
%H:%i'),7
*** (1) WAITING FOR THIS LOCK TO BE GRANTED:
RECORD LOCKS space id 0 page no 48310 n bits 568 table test/alex1 index
sybole trx id 0 290252780 lock mode S waiting
Record lock, heap no 324 RECORD: info bits 0 0: len 7; hex 61613335383138;
asc aa35818;; 1:
*** (2) TRANSACTION:
```



```

25.44 reads/s, 18494 avg bytes/read, 17.55 writes/s, 2.33 fsyncs/s
-----
INSERT BUFFER AND ADAPTIVE HASH INDEX
-----
Ibuf for space 0: size 1, free list len 19, seg size 21,
85004 inserts, 85004 merged recs, 26669 merges
Hash table size 207619, used cells 14461, node heap has 16 buffer(s)
1877.67 hash searches/s, 5121.10 non-hash searches/s
---
LOG
---
Log sequence number 18 1212842764
Log flushed up to 18 1212665295
Last checkpoint at 18 1135877290
0 pending log writes, 0 pending chkp writes
4341 log i/o's done, 1.22 log i/o's/second
-----
BUFFER POOL AND MEMORY
-----
Total memory allocated 84966343; in additional pool allocated 1402624
Buffer pool size 3200
Free buffers 110
Database pages 3074
Modified db pages 2674
Pending reads 0
Pending writes: LRU 0, flush list 0, single page 0
Pages read 171380, created 51968, written 194688
28.72 reads/s, 20.72 creates/s, 47.55 writes/s
Buffer pool hit rate 999 / 1000
-----
ROW OPERATIONS
-----
0 queries inside InnoDB, 0 queries in queue
Main thread process no. 3004, id 7176, state: purging
Number of rows inserted 3738558, updated 127415, deleted 33707, read 755779
1586.13 inserts/s, 50.89 updates/s, 28.44 deletes/s, 107.88 reads/s
-----
END OF INNODB MONITOR OUTPUT
=====

```

Some notes on the output:

- If the [TRANSACTIONS](#) section reports lock waits, your applications may have lock contention. The output can also help to trace the reasons for transaction deadlocks.
- The [SEMAPHORES](#) section reports threads waiting for a semaphore and statistics on how many times threads have needed a spin or a wait on a mutex or a rw-lock semaphore. A large number of threads waiting for semaphores may be a result of disk I/O, or contention problems inside [InnoDB](#). Contention can be due to heavy parallelism of queries or problems in operating system thread scheduling. Setting [innodb\\_thread\\_concurrency](#) smaller than the default value can help in such situations.
- The [BUFFER POOL AND MEMORY](#) section gives you statistics on pages read and written. You can calculate from these numbers how many data file I/O operations your queries currently are doing.
- The [ROW OPERATIONS](#) section shows what the main thread is doing.

Beginning with MySQL 4.0.19, [InnoDB](#) sends diagnostic output to stderr or files instead of stdout or fixed-size memory buffers, to avoid potential buffer overflow errors. As a side effect, the output of [SHOW INNODB STATUS](#) is written to a status file in the MySQL data directory every fifteen seconds. The name of the file is [innodb\\_status.pid](#), where [pid](#) is the server process ID. [InnoDB](#) removes the file for a normal shutdown. If abnormal shutdowns have occurred, instances of these status files may be present and must be removed manually. Before removing them, you might want to examine them to see whether they contain useful information about the cause of abnormal shutdowns. Beginning with MySQL 4.0.21, the [innodb\\_status.pid](#) file is created only if the configuration option [innodb\\_status\\_file=1](#) is set.

## 14.2.13. Implementation of Multi-Versioning

Because **InnoDB** is a multi-versioned storage engine, it must keep information about old versions of rows in the tablespace. This information is stored in a data structure called a *rollback segment* (after an analogous data structure in Oracle).

Internally, **InnoDB** adds two fields to each row stored in the database. A 6-byte field indicates the transaction identifier for the last transaction that inserted or updated the row. Also, a deletion is treated internally as an update where a special bit in the row is set to mark it as deleted. Each row also contains a 7-byte field called the roll pointer. The roll pointer points to an undo log record written to the rollback segment. If the row was updated, the undo log record contains the information necessary to rebuild the content of the row before it was updated.

**InnoDB** uses the information in the rollback segment to perform the undo operations needed in a transaction rollback. It also uses the information to build earlier versions of a row for a consistent read.

Undo logs in the rollback segment are divided into insert and update undo logs. Insert undo logs are needed only in transaction rollback and can be discarded as soon as the transaction commits. Update undo logs are used also in consistent reads, but they can be discarded only after there is no transaction present for which **InnoDB** has assigned a snapshot that in a consistent read could need the information in the update undo log to build an earlier version of a database row.

You must remember to commit your transactions regularly, including those transactions that issue only consistent reads. Otherwise, **InnoDB** cannot discard data from the update undo logs, and the rollback segment may grow too big, filling up your tablespace.

The physical size of an undo log record in the rollback segment is typically smaller than the corresponding inserted or updated row. You can use this information to calculate the space need for your rollback segment.

In the **InnoDB** multi-versioning scheme, a row is not physically removed from the database immediately when you delete it with an SQL statement. Only when **InnoDB** can discard the update undo log record written for the deletion can it also physically remove the corresponding row and its index records from the database. This removal operation is called a purge, and it is quite fast, usually taking the same order of time as the SQL statement that did the deletion.

In a scenario where the user inserts and deletes rows in smallish batches at about the same rate in the table, it is possible that the purge thread starts to lag behind, and the table grows bigger and bigger, making everything disk-bound and very slow. Even if the table would carry just 10MB of useful data, it may grow to occupy 10GB with all the dead rows. In such a case, it would be good to throttle new row operations, and allocate more resources for the purge thread. Starting with MySQL 4.0.22 and 4.1.6, the `innodb_max_purge_lag` system variable exists for exactly this purpose. See [Section 14.2.5, “InnoDB Startup Options and System Variables”](#), for more information.

## 14.2.14. InnoDB Table and Index Structures

MySQL stores its data dictionary information for tables in `.frm` files in database directories. This is true for all MySQL storage engines. But every **InnoDB** table also has its own entry in the **InnoDB** internal data dictionary inside the tablespace. When MySQL drops a table or a database, it has to delete both an `.frm` file or files, and the corresponding entries inside the **InnoDB** data dictionary. This is the reason why you cannot move **InnoDB** tables between databases simply by moving the `.frm` files. It is also the reason why `DROP DATABASE` did not work for **InnoDB** type tables before MySQL 3.23.44.

Every **InnoDB** table has a special index called the *clustered index* where the data for the rows is stored. If you define a **PRIMARY KEY** on your table, the index of the primary key is the clustered index.

If you do not define a **PRIMARY KEY** for your table, MySQL picks the first **UNIQUE** index that has only **NOT NULL** columns as the primary key and **InnoDB** uses it as the clustered index. If there is no



such index in the table, [InnoDB](#) internally generates a clustered index where the rows are ordered by the row ID that [InnoDB](#) assigns to the rows in such a table. The row ID is a 6-byte field that increases monotonically as new rows are inserted. Thus, the rows ordered by the row ID are physically in insertion order.

Accessing a row through the clustered index is fast because the row data is on the same page where the index search leads. If a table is large, the clustered index architecture often saves a disk I/O when compared to the traditional solution. (In many database systems, data storage uses a different page from the index record.)

In [InnoDB](#), the records in non-clustered indexes (also called secondary indexes) contain the primary key value for the row. [InnoDB](#) uses this primary key value to search for the row from the clustered index. Note that if the primary key is long, the secondary indexes use more space.

[InnoDB](#) compares [CHAR](#) and [VARCHAR](#) strings of different lengths such that the remaining length in the shorter string is treated as if padded with spaces.

### 14.2.14.1. Physical Structure of an Index

All [InnoDB](#) indexes are B-trees where the index records are stored in the leaf pages of the tree. The default size of an index page is 16KB. When new records are inserted, [InnoDB](#) tries to leave 1/16 of the page free for future insertions and updates of the index records.

If index records are inserted in a sequential order (ascending or descending), the resulting index pages are about 15/16 full. If records are inserted in a random order, the pages are from 1/2 to 15/16 full. If the fill factor of an index page drops below 1/2, [InnoDB](#) tries to contract the index tree to free the page.

### 14.2.14.2. Insert Buffering

It is a common situation in database applications that the primary key is a unique identifier and new rows are inserted in the ascending order of the primary key. Thus, the insertions to the clustered index do not require random reads from a disk.

On the other hand, secondary indexes are usually non-unique, and insertions into secondary indexes happen in a relatively random order. This would cause a lot of random disk I/O operations without a special mechanism used in [InnoDB](#).

If an index record should be inserted to a non-unique secondary index, [InnoDB](#) checks whether the secondary index page is in the buffer pool. If that is the case, [InnoDB](#) does the insertion directly to the index page. If the index page is not found in the buffer pool, [InnoDB](#) inserts the record to a special insert buffer structure. The insert buffer is kept so small that it fits entirely in the buffer pool, and insertions can be done very fast.

Periodically, the insert buffer is merged into the secondary index trees in the database. Often it is possible to merge several insertions to the same page of the index tree, saving disk I/O operations. It has been measured that the insert buffer can speed up insertions into a table up to 15 times.

The insert buffer merging may continue to happen *after* the inserting transaction has been committed. In fact, it may continue to happen after a server shutdown and restart (see [Section 14.2.9.1, “Forcing InnoDB Recovery”](#)).

The insert buffer merging may take many hours, when many secondary indexes must be updated, and many rows have been inserted. During this time, disk I/O will be increased, which can cause significant slowdown on disk-bound queries. Another significant background I/O operation is the purge thread (see [Section 14.2.13, “Implementation of Multi-Versioning”](#)).

### 14.2.14.3. Adaptive Hash Indexes

If a table fits almost entirely in main memory, the fastest way to perform queries on it is to use hash indexes. [InnoDB](#) has a mechanism that monitors index searches made to the indexes defined for a table. If [InnoDB](#) notices that queries could benefit from building a hash index, it does so automatically.

Note that the hash index is always built based on an existing B-tree index on the table. [InnoDB](#) can build a hash index on a prefix of any length of the key defined for the B-tree, depending on the pattern of searches that [InnoDB](#) observes for the B-tree index. A hash index can be partial: It is not required that the whole B-tree index is cached in the buffer pool. [InnoDB](#) builds hash indexes on demand for those pages of the index that are often accessed.

In a sense, [InnoDB](#) tailors itself through the adaptive hash index mechanism to ample main memory, coming closer to the architecture of main-memory databases.

#### 14.2.14.4. Physical Row Structure

Records in [InnoDB](#) tables have the following characteristics:

- Each index record contains a six-byte header. The header is used to link together consecutive records, and also in row-level locking.
- Records in the clustered index contain fields for all user-defined columns. In addition, there is a six-byte field for the transaction ID and a seven-byte field for the roll pointer.
- If no primary key was defined for a table, each clustered index record also contains a six-byte row ID field.
- Each secondary index record contains also all the fields defined for the clustered index key.
- A record contains also a pointer to each field of the record. If the total length of the fields in a record is less than 128 bytes, the pointer is one byte; otherwise, two bytes. The array of these pointers is called the record directory. The area where these pointers point is called the data part of the record.
- Internally, [InnoDB](#) stores fixed-length character columns such as `CHAR(10)` in a fixed-length format. [InnoDB](#) truncates trailing spaces from `VARCHAR` columns. Note that MySQL may internally convert `CHAR` columns to `VARCHAR`. See [Section 13.1.5.1, “Silent Column Specification Changes”](#).
- An SQL `NULL` value reserves 1 or 2 bytes in the record directory. Besides that, an SQL `NULL` value reserves zero bytes in the data part of the record if stored in a variable length column. In a fixed-length column, it reserves the fixed length of the column in the data part of the record. The motivation behind reserving the fixed space for `NULL` values is that it enables an update of the column from `NULL` to a non-`NULL` value to be done in place without causing fragmentation of the index page.

### 14.2.15. [InnoDB](#) File Space Management and Disk I/O

#### 14.2.15.1. [InnoDB](#) Disk I/O

[InnoDB](#) uses simulated asynchronous disk I/O: [InnoDB](#) creates a number of threads to take care of I/O operations, such as read-ahead.

There are two read-ahead heuristics in [InnoDB](#):

- In sequential read-ahead, if [InnoDB](#) notices that the access pattern to a segment in the tablespace is sequential, it posts in advance a batch of reads of database pages to the I/O system.

- In random read-ahead, if `InnoDB` notices that some area in a tablespace seems to be in the process of being fully read into the buffer pool, it posts the remaining reads to the I/O system.

`InnoDB` uses a novel file flush technique called *doublewrite*. It adds safety to recovery following an operating system crash or a power outage, and improves performance on most varieties of Unix by reducing the need for `fsync()` operations.

Doublewrite means that before writing pages to a data file, `InnoDB` first writes them to a contiguous tablespace area called the doublewrite buffer. Only after the write and the flush to the doublewrite buffer has completed does `InnoDB` write the pages to their proper positions in the data file. If the operating system crashes in the middle of a page write, `InnoDB` can later find a good copy of the page from the doublewrite buffer during recovery.

### 14.2.15.2. File Space Management

The data files that you define in the configuration file form the tablespace of `InnoDB`. The files are simply concatenated to form the tablespace. There is no striping in use. Currently, you cannot define where within the tablespace your tables are allocated. However, in a newly created tablespace, `InnoDB` allocates space starting from the first data file.

The tablespace consists of database pages with a default size of 16KB. The pages are grouped into extents of 64 consecutive pages. The “files” inside a tablespace are called *segments* in `InnoDB`. The term “rollback segment” is somewhat confusing because it actually contains many tablespace segments.

Two segments are allocated for each index in `InnoDB`. One is for non-leaf nodes of the B-tree, the other is for the leaf nodes. The idea here is to achieve better sequentiality for the leaf nodes, which contain the data.

When a segment grows inside the tablespace, `InnoDB` allocates the first 32 pages to it individually. After that `InnoDB` starts to allocate whole extents to the segment. `InnoDB` can add to a large segment up to 4 extents at a time to ensure good sequentiality of data.

Some pages in the tablespace contain bitmaps of other pages, and therefore a few extents in an `InnoDB` tablespace cannot be allocated to segments as a whole, but only as individual pages.

When you ask for available free space in the tablespace by issuing a `SHOW TABLE STATUS` statement, `InnoDB` reports the extents that are definitely free in the tablespace. `InnoDB` always reserves some extents for cleanup and other internal purposes; these reserved extents are not included in the free space.

When you delete data from a table, `InnoDB` contracts the corresponding B-tree indexes. Whether the freed space becomes available for other users depends on whether the pattern of deletes frees individual pages or extents to the tablespace. Dropping a table or deleting all rows from it is guaranteed to release the space to other users, but remember that deleted rows are physically removed only in an (automatic) purge operation after they are no longer needed for transaction rollbacks or consistent reads. (See [Section 14.2.13, “Implementation of Multi-Versioning”](#).)

### 14.2.15.3. Defragmenting a Table

If there are random insertions into or deletions from the indexes of a table, the indexes may become fragmented. Fragmentation means that the physical ordering of the index pages on the disk is not close to the index ordering of the records on the pages, or that there are many unused pages in the 64-page blocks that were allocated to the index.

A symptom of fragmentation is that a table takes more space than it “should” take. How much that is exactly, is difficult to determine. All `InnoDB` data and indexes are stored in B-trees, and their fill factor may vary from 50% to 100%. Another symptom of fragmentation is that a table scan such as this takes more time than it “should” take:

```
SELECT COUNT(*) FROM t WHERE a_non_indexed_column <> 12345;
```

(In the preceding query, we are “fooling” the SQL optimizer into scanning the clustered index, rather than a secondary index.) Most disks can read 10 to 50MB/s, which can be used to estimate how fast a table scan should run.

It can speed up index scans if you periodically perform a “null” `ALTER TABLE` operation:

```
ALTER TABLE tbl_name TYPE=InnoDB;
```

That causes MySQL to rebuild the table. Another way to perform a defragmentation operation is to use `mysqldump` to dump the table to a text file, drop the table, and reload it from the dump file.

If the insertions to an index are always ascending and records are deleted only from the end, the `InnoDB` filespace management algorithm guarantees that fragmentation in the index does not occur.

## 14.2.16. InnoDB Error Handling

Error handling in `InnoDB` is not always the same as specified in the SQL standard. According to the standard, any error during an SQL statement should cause the rollback of that statement. `InnoDB` sometimes rolls back only part of the statement, or the whole transaction. The following items describe how `InnoDB` performs error handling:

- If you run out of file space in the tablespace, a MySQL `Table is full` error occurs and `InnoDB` rolls back the SQL statement.
- A transaction deadlock or a timeout in a lock wait causes `InnoDB` to roll back the whole transaction.

When a transaction rollback occurs due to a deadlock or lock wait timeout, it cancels the effect of the statements within the transaction. But if the start-transaction statement was `START TRANSACTION` or `BEGIN` statement, rollback does not cancel that statement. Further SQL statements become part of the transaction until the occurrence of `COMMIT`, `ROLLBACK`, or some SQL statement that causes an implicit commit.

- A duplicate-key error rolls back the SQL statement, if you have not specified the `IGNORE` option in your statement.
- A `row too long error` rolls back the SQL statement.
- Other errors are mostly detected by the MySQL layer of code (above the `InnoDB` storage engine level), and they roll back the corresponding SQL statement. Locks are not released in a rollback of a single SQL statement.

During such implicit rollbacks, as well as during the explicit `ROLLBACK SQL` command, `SHOW PROCESSLIST` displays “Rolling back” in the `State` column for the connection (starting from MySQL 4.1.8).

### 14.2.16.1. InnoDB Error Codes

The following is a non-exhaustive list of common `InnoDB`-specific errors that you may encounter, with information about why each occurs and how to resolve the problem.

- 1005 (`ER_CANT_CREATE_TABLE`)

Cannot create table. If the error message refers to `errno` 150, table creation failed because a foreign key constraint was not correctly formed. If the error message refers to `errno` -1, table creation probably failed because the table included a column name that matched the name of an internal InnoDB table.

- 1016 (`ER_CANT_OPEN_FILE`)

Cannot find the InnoDB table from the InnoDB data files, although the `.frm` file for the table exists. See [Section 14.2.18.1, “Troubleshooting InnoDB Data Dictionary Operations”](#).

- 1114 (`ER_RECORD_FILE_FULL`)

InnoDB has run out of free space in the tablespace. You should reconfigure the tablespace to add a new data file.

- 1205 (`ER_LOCK_WAIT_TIMEOUT`)

Lock wait timeout expired. Transaction was rolled back.

- 1213 (`ER_LOCK_DEADLOCK`)

Transaction deadlock. You should rerun the transaction.

- 1216 (`ER_NO_REFERENCED_ROW`)

You are trying to add a row but there is no parent row, and a foreign key constraint fails. You should add the parent row first.

- 1217 (`ER_ROW_IS_REFERENCED`)

You are trying to delete a parent row that has children, and a foreign key constraint fails. You should delete the children first.

## 14.2.16.2. Operating System Error Codes

To print the meaning of an operating system error number, use the `pererror` program that comes with the MySQL distribution.

The following table provides a list of some common Linux system error codes. For a more complete list, see [Linux source code](http://www.iglu.org.il/lxr/source/include/asm-i386/errno.h) [<http://www.iglu.org.il/lxr/source/include/asm-i386/errno.h>].

- 1 (`EPERM`)

Operation not permitted

- 2 (`ENOENT`)

No such file or directory

- 3 (`ESRCH`)

No such process

- 4 (`EINTR`)

Interrupted system call

- 5 (EIO)  
I/O error
- 6 (ENXIO)  
No such device or address
- 7 (E2BIG)  
Arg list too long
- 8 (ENOEXEC)  
Exec format error
- 9 (EBADF)  
Bad file number
- 10 (ECHILD)  
No child processes
- 11 (EAGAIN)  
Try again
- 12 (ENOMEM)  
Out of memory
- 13 (EACCES)  
Permission denied
- 14 (EFAULT)  
Bad address
- 15 (ENOTBLK)  
Block device required
- 16 (EBUSY)  
Device or resource busy
- 17 (EEXIST)  
File exists
- 18 (EXDEV)  
Cross-device link
- 19 (ENODEV)  
No such device

- 20 (ENOTDIR)  
Not a directory
- 21 (EISDIR)  
Is a directory
- 22 (EINVAL)  
Invalid argument
- 23 (ENFILE)  
File table overflow
- 24 (EMFILE)  
Too many open files
- 25 (ENOTTY)  
Inappropriate ioctl for device
- 26 (ETXTBSY)  
Text file busy
- 27 (EFBIG)  
File too large
- 28 (ENOSPC)  
No space left on device
- 29 (ESPIPE)  
Illegal seek
- 30 (EROFS)  
Read-only file system
- 31 (EMLINK)  
Too many links

The following table provides a list of some common Windows system error codes. For a complete list see the [Microsoft Web site](http://msdn.microsoft.com/library/default.asp?url=/library/en-us/debug/base/system_error_codes.asp) [http://msdn.microsoft.com/library/default.asp?url=/library/en-us/debug/base/system\_error\_codes.asp].

- 1 (ERROR\_INVALID\_FUNCTION)  
Incorrect function.
- 2 (ERROR\_FILE\_NOT\_FOUND)

The system cannot find the file specified.

- 3 ( `ERROR_PATH_NOT_FOUND` )

The system cannot find the path specified.

- 4 ( `ERROR_TOO_MANY_OPEN_FILES` )

The system cannot open the file.

- 5 ( `ERROR_ACCESS_DENIED` )

Access is denied.

- 6 ( `ERROR_INVALID_HANDLE` )

The handle is invalid.

- 7 ( `ERROR_ARENA_TRASHED` )

The storage control blocks were destroyed.

- 8 ( `ERROR_NOT_ENOUGH_MEMORY` )

Not enough storage is available to process this command.

- 9 ( `ERROR_INVALID_BLOCK` )

The storage control block address is invalid.

- 10 ( `ERROR_BAD_ENVIRONMENT` )

The environment is incorrect.

- 11 ( `ERROR_BAD_FORMAT` )

An attempt was made to load a program with an incorrect format.

- 12 ( `ERROR_INVALID_ACCESS` )

The access code is invalid.

- 13 ( `ERROR_INVALID_DATA` )

The data is invalid.

- 14 ( `ERROR_OUTOFMEMORY` )

Not enough storage is available to complete this operation.

- 15 ( `ERROR_INVALID_DRIVE` )

The system cannot find the drive specified.

- 16 ( `ERROR_CURRENT_DIRECTORY` )

The directory cannot be removed.

- 17 ( `ERROR_NOT_SAME_DEVICE` )



The system cannot move the file to a different disk drive.

- 18 ( `ERROR_NO_MORE_FILES` )

There are no more files.

- 19 ( `ERROR_WRITE_PROTECT` )

The media is write protected.

- 20 ( `ERROR_BAD_UNIT` )

The system cannot find the device specified.

- 21 ( `ERROR_NOT_READY` )

The device is not ready.

- 22 ( `ERROR_BAD_COMMAND` )

The device does not recognize the command.

- 23 ( `ERROR_CRC` )

Data error (cyclic redundancy check).

- 24 ( `ERROR_BAD_LENGTH` )

The program issued a command but the command length is incorrect.

- 25 ( `ERROR_SEEK` )

The drive cannot locate a specific area or track on the disk.

- 26 ( `ERROR_NOT_DOS_DISK` )

The specified disk or diskette cannot be accessed.

- 27 ( `ERROR_SECTOR_NOT_FOUND` )

The drive cannot find the sector requested.

- 28 ( `ERROR_OUT_OF_PAPER` )

The printer is out of paper.

- 29 ( `ERROR_WRITE_FAULT` )

The system cannot write to the specified device.

- 30 ( `ERROR_READ_FAULT` )

The system cannot read from the specified device.

- 31 ( `ERROR_GEN_FAILURE` )

A device attached to the system is not functioning.

- 32 ( `ERROR_SHARING_VIOLATION` )

The process cannot access the file because it is being used by another process.

- 33 ( [ERROR\\_LOCK\\_VIOLATION](#) )

The process cannot access the file because another process has locked a portion of the file.

- 34 ( [ERROR\\_WRONG\\_DISK](#) )

The wrong diskette is in the drive. Insert %2 (Volume Serial Number: %3) into drive %1.

- 36 ( [ERROR\\_SHARING\\_BUFFER\\_EXCEEDED](#) )

Too many files opened for sharing.

- 38 ( [ERROR\\_HANDLE\\_EOF](#) )

Reached the end of the file.

- 39 ( [ERROR\\_HANDLE\\_DISK\\_FULL](#) )

The disk is full.

- 87 ( [ERROR\\_INVALID\\_PARAMETER](#) )

The parameter is incorrect. (If this error occurs on MySQL 4.1.9 on Windows and you have set `innodb_file_per_table` in a server option file, this is [Bug#8021](#) [<http://bugs.mysql.com/8021>], and a workaround is to add the line `innodb_flush_method=unbuffered` to the file as well.)

- 112 ( [ERROR\\_DISK\\_FULL](#) )

The disk is full.

- 123 ( [ERROR\\_INVALID\\_NAME](#) )

The filename, directory name, or volume label syntax is incorrect.

- 1450 ( [ERROR\\_NO\\_SYSTEM\\_RESOURCES](#) )

Insufficient system resources exist to complete the requested service.

## 14.2.17. Restrictions on InnoDB Tables

- **Warning:** Do *not* convert MySQL system tables in the `mysql` database from `MyISAM` to `InnoDB` tables! This is an unsupported operation. If you do this, MySQL does not restart until you restore the old system tables from a backup or re-generate them with the `mysql_install_db` script.
- A table cannot contain more than 1000 columns.
- The internal maximum key length is 3500 bytes, but MySQL itself restricts this to 1024 bytes.
- The maximum row length, except for `VARCHAR`, `BLOB` and `TEXT` columns, is slightly less than half of a database page. That is, the maximum row length is about 8000 bytes. `LONGBLOB` and `LONGTEXT` columns must be less than 4GB, and the total row length, including also `BLOB` and `TEXT` columns, must be less than 4GB. `InnoDB` stores the first 768 bytes of a `VARCHAR`, `BLOB`, or `TEXT` column in the row, and the rest into separate pages.

- On some older operating systems, files must be less than 2GB. This is not a limitation of `InnoDB` itself, but if you require a large tablespace, you will need to configure it using several smaller data files rather than one or a file large data files.
- The combined size of the `InnoDB` log files must be less than 4GB.
- The minimum tablespace size is 10MB. The maximum tablespace size is four billion database pages (64TB). This is also the maximum size for a table.
- `InnoDB` tables do not support `FULLTEXT` indexes.
- `InnoDB` tables do not support spatial data types.
- `ANALYZE TABLE` determines index cardinality (as displayed in the `Cardinality` column of `SHOW INDEX` output) by doing ten random dives to each of the index trees and updating index cardinality estimates accordingly. Note that because these are only estimates, repeated runs of `ANALYZE TABLE` may produce different numbers. This makes `ANALYZE TABLE` fast on `InnoDB` tables but not 100% accurate as it doesn't take all rows into account.

MySQL uses index cardinality estimates only in join optimization. If some join is not optimized in the right way, you can try using `ANALYZE TABLE`. In the few cases that `ANALYZE TABLE` doesn't produce values good enough for your particular tables, you can use `FORCE INDEX` with your queries to force the use of a particular index, or set the `max_seeks_for_key` system variable to ensure that MySQL prefers index lookups over table scans. See [Section 5.2.2, “System Variables”](#), and [Section A.6, “Optimizer-Related Issues”](#).

- `SHOW TABLE STATUS` does not give accurate statistics on `InnoDB` tables, except for the physical size reserved by the table. The row count is only a rough estimate used in SQL optimization.
- `InnoDB` does not keep an internal count of rows in a table. (In practice, this would be somewhat complicated due to multi-versioning.) To process a `SELECT COUNT(*) FROM t` statement, `InnoDB` must scan an index of the table, which takes some time if the index is not entirely in the buffer pool. To get a fast count, you have to use a counter table you create yourself and let your application update it according to the inserts and deletes it does. If your table does not change often, using the MySQL query cache is a good solution. `SHOW TABLE STATUS` also can be used if an approximate row count is sufficient. See [Section 14.2.12, “InnoDB Performance Tuning Tips”](#).
- On Windows, `InnoDB` always stores database and table names internally in lowercase. To move databases in binary format from Unix to Windows or from Windows to Unix, you should always use explicitly lowercase names when creating databases and tables.
- For an `AUTO_INCREMENT` column, you must always define an index for the table, and that index must contain just the `AUTO_INCREMENT` column. In `MyISAM` tables, the `AUTO_INCREMENT` column may be part of a multi-column index.
- Before MySQL 4.1.12, `InnoDB` does not support the `AUTO_INCREMENT` table option for setting the initial sequence value in an `ALTER TABLE` statement. Before MySQL 4.1.14, the same is true for `CREATE TABLE`. To set the value with `InnoDB`, insert a dummy row with a value one less and delete that dummy row, or insert the first row with an explicit value specified.
- While initializing a previously specified `AUTO_INCREMENT` column on a table, `InnoDB` sets an exclusive lock on the end of the index associated with the `AUTO_INCREMENT` column. In accessing the auto-increment counter, `InnoDB` uses a specific table lock mode `AUTO-INC` where the lock lasts only to the end of the current SQL statement, not to the end of the entire transaction. Note that other clients cannot insert into the table while the `AUTO-INC` table lock is held; see [Section 14.2.11.2, “InnoDB and AUTOCOMMIT”](#).
- When you restart the MySQL server, `InnoDB` may reuse an old value that was generated for an `AUTO_INCREMENT` column but never stored (that is, a value that was generated during an old

transaction that was rolled back).

- When an `AUTO_INCREMENT` column runs out of values, `InnoDB` wraps a `BIGINT` to `-9223372036854775808` and `BIGINT UNSIGNED` to `1`. However, `BIGINT` values have 64 bits, so do note that if you were to insert one million rows per second, it would still take nearly three hundred thousand years before `BIGINT` reached its upper bound. With all other integer type columns, a duplicate-key error results. This is similar to how `MyISAM` works, because it is mostly general MySQL behavior and not about any storage engine in particular.
- `DELETE FROM tbl_name` does not regenerate the table but instead deletes all rows, one by one.
- Under some conditions, `TRUNCATE tbl_name` for an `InnoDB` table is mapped to `DELETE FROM+ tbl_name` and doesn't reset the `AUTO_INCREMENT` counter. See [Section 13.2.9, “TRUNCATE Syntax”](#).
- Before MySQL 4.0.14 or 4.1.0, if you tried to create a unique index on a prefix of a column you got an error:

```
CREATE TABLE T (A CHAR(20), B INT, UNIQUE (A(5))) TYPE = InnoDB;
```

If you created a non-unique index on a prefix of a column, `InnoDB` created an index over the whole column. These restrictions were removed in MySQL 4.0.14.

- Before MySQL 4.0.20 or 4.1.2, the MySQL `LOCK TABLES` operation does not know about `InnoDB` row-level locks set by completed SQL statements. This means that you can get a table lock on a table even if there still exist transactions by other users who have row-level locks on the same table. Thus, your operations on the table may have to wait if they collide with these locks of other users. Also a deadlock is possible. However, this does not endanger transaction integrity, because the row-level locks set by `InnoDB` always take care of the integrity. Also, a table lock prevents other transactions from acquiring more row-level locks (in a conflicting lock mode) on the table.
- Beginning with MySQL 4.0.20 and 4.1.2, the MySQL `LOCK TABLES` operation acquires two locks on each table if `innodb_table_locks=1` (the default). In addition to a table lock on the MySQL layer, it also acquires an `InnoDB` table lock. Older versions of MySQL do not acquire `InnoDB` table locks. Beginning with MySQL 4.0.22 and 4.1.7, the old behavior can be selected by setting `innodb_table_locks=0`. If no `InnoDB` table lock is acquired, `LOCK TABLES` completes even if some records of the tables are being locked by other transactions.
- All `InnoDB` locks held by a transaction are released when the transaction is committed or aborted. Thus, it does not make much sense to invoke `LOCK TABLES` on `InnoDB` tables in `AUTOCOMMIT=1` mode, because the acquired `InnoDB` table locks would be released immediately.
- Sometimes it would be useful to lock further tables in the course of a transaction. Unfortunately, `LOCK TABLES` in MySQL performs an implicit `COMMIT` and `UNLOCK TABLES`. An `InnoDB` variant of `LOCK TABLES` has been planned that can be executed in the middle of a transaction.
- Before MySQL 3.23.52, replication always ran with autocommit enabled. Therefore consistent reads in the slave would also see partially processed transactions, and thus the read would not be really consistent in the slave. This restriction was removed in MySQL 3.23.52.
- The `LOAD TABLE FROM MASTER` statement for setting up replication slave servers does not work for `InnoDB` tables. A workaround is to alter the table to `MyISAM` on the master, do then the load, and after that alter the master table back to `InnoDB`. Do not do this if the tables use `InnoDB`-specific features such as foreign keys.
- The default database page size in `InnoDB` is 16KB. By recompiling the code, you can set it to values ranging from 8KB to 64KB. You must update the values of `UNIV_PAGE_SIZE` and `UNIV_PAGE_SIZE_SHIFT` in the `univ.i` source file.

## 14.2.18. InnoDB Troubleshooting

The following general guidelines apply to troubleshooting InnoDB problems:

- When an operation fails or you suspect a bug, you should look at the MySQL server error log, which is the file in the data directory that has a suffix of `.err`.
- When troubleshooting, it is usually best to run the MySQL server from the command prompt, rather than through the `mysqld_safe` wrapper or as a Windows service. You can then see what `mysqld` prints to the console, and so have a better grasp of what is going on. On Windows, you must start the server with the `--console` option to direct the output to the console window.
- Use the InnoDB Monitors to obtain information about a problem (see Section 14.2.12.1, “SHOW ENGINE INNODB STATUS and the InnoDB Monitors”). If the problem is performance-related, or your server appears to be hung, you should use `innodb_monitor` to print information about the internal state of InnoDB. If the problem is with locks, use `innodb_lock_monitor`. If the problem is in creation of tables or other data dictionary operations, use `innodb_table_monitor` to print the contents of the InnoDB internal data dictionary.
- If you suspect that a table is corrupt, run `CHECK TABLE` on that table.

### 14.2.18.1. Troubleshooting InnoDB Data Dictionary Operations

A specific issue with tables is that the MySQL server keeps data dictionary information in `.frm` files it stores in the database directories, whereas InnoDB also stores the information into its own data dictionary inside the tablespace files. If you move `.frm` files around, or use `DROP DATABASE` in MySQL versions before 3.23.44, or the server crashes in the middle of a data dictionary operation, the locations of the `.frm` files may end up out of synchrony with the locations recorded in the InnoDB internal data dictionary.

A symptom of an out-of-sync data dictionary is that a `CREATE TABLE` statement fails. If this occurs, you should look in the server's error log. If the log says that the table already exists inside the InnoDB internal data dictionary, you have an orphaned table inside the InnoDB tablespace files that has no corresponding `.frm` file. The error message looks like this:

```
InnoDB: Error: table test/parent already exists in InnoDB internal
InnoDB: data dictionary. Have you deleted the .frm file
InnoDB: and not used DROP TABLE? Have you used DROP DATABASE
InnoDB: for InnoDB tables in MySQL version <= 3.23.43?
InnoDB: See the Restrictions section of the InnoDB manual.
InnoDB: You can drop the orphaned table inside InnoDB by
InnoDB: creating an InnoDB table with the same name in another
InnoDB: database and moving the .frm file to the current database.
InnoDB: Then MySQL thinks the table exists, and DROP TABLE will
InnoDB: succeed.
```

You can drop the orphaned table by following the instructions given in the error message. If you are still unable to use `DROP TABLE` successfully, the problem may be due to name completion in the `mysql` client. To work around this problem, start the `mysql` client with the `--skip-auto-rehash` option and try `DROP TABLE` again. (With name completion on, `mysql` tries to construct a list of table names, which fails when a problem such as just described exists.)

Another symptom of an out-of-sync data dictionary is that MySQL prints an error that it cannot open a `.InnoDB` file:

```
ERROR 1016: Can't open file: 'child2.InnoDB'. (errno: 1)
```

In the error log you can find a message like this:

```
InnoDB: Cannot find table test/child2 from the internal data dictionary
InnoDB: of InnoDB though the .frm file for the table exists. Maybe you
InnoDB: have deleted and recreated InnoDB data files but have forgotten
InnoDB: to delete the corresponding .frm files of InnoDB tables?
```

This means that there is an orphaned `.frm` file without a corresponding table inside `InnoDB`. You can drop the orphaned `.frm` file by deleting it manually.

If MySQL crashes in the middle of an `ALTER TABLE` operation, you may end up with an orphaned temporary table inside the `InnoDB` tablespace. With `innodb_table_monitor` you see a table whose name is `#sql-...`. Starting from MySQL 4.0.6, you can perform SQL statements also on tables whose name contains the character `#` if you enclose the name within backticks. Thus, you can drop such an orphaned table like any other orphaned table using the method described earlier. Note that to copy or rename a file in the Unix shell, you need to put the file name in double quotes if the file name contains `#`.

Older MySQL versions did not allow accessing any table with a name containing `#`. The solution in older MySQL versions is to use a special `InnoDB` mechanism available starting from MySQL 3.23.48. When you have an orphaned table `#sql-id` inside the tablespace, you can cause `InnoDB` to rename it to `rsql-id_recover_innodb_tmp_table` with the following statement:

```
CREATE TABLE `rsql-id_recover_innodb_tmp_table` (...) TYPE=InnoDB;
```

## 14.3. The `MERGE` Storage Engine

The `MERGE` storage engine was introduced in MySQL 3.23.25. It is also known as the `MRG_MyISAM` engine.

A `MERGE` table is a collection of identical `MyISAM` tables that can be used as one. “Identical” means that all tables have identical column and index information. You cannot merge `MyISAM` tables in which the columns are listed in a different order, do not have exactly the same columns, or have the indexes in different order. However, any or all of the `MyISAM` tables can be compressed with `myisampack`. See [Section 8.5, “myisampack — Generate Compressed, Read-Only MyISAM Tables”](#). Differences in table options such as `AVG_ROW_LENGTH`, `MAX_ROWS`, or `PACK_KEYS` do not matter.

When you create a `MERGE` table, MySQL creates two files on disk. The files have names that begin with the table name and have an extension to indicate the file type. An `.frm` file stores the table format, and an `.MRG` file contains the names of the tables that should be used as one. (Originally, all used tables had to be in the same database as the `MERGE` table itself. This restriction has been lifted as of MySQL 4.1.1.)

You can use `SELECT`, `DELETE`, `UPDATE`, and (as of MySQL 4.0) `INSERT` on `MERGE` tables. You must have `SELECT`, `UPDATE`, and `DELETE` privileges on the `MyISAM` tables that you map to a `MERGE` table.

**Note:** The use of `MERGE` tables entails the following security issue: If a user has access to `MyISAM` table `t`, that user can create a `MERGE` table `m` that accesses `t`. However, if the user's privileges on `t` are subsequently revoked, the user can continue to access `t` by doing so through `m`. If this behavior is undesirable, you can start the server with the new `--skip-merge` option to disable the `MERGE` storage engine. This option is available as of MySQL 4.1.21.

If you `DROP` the `MERGE` table, you are dropping only the `MERGE` specification. The underlying tables are not affected.

To create a `MERGE` table, you must specify a `UNION=(list-of-tables)` clause that indicates which `MyISAM` tables you want to use as one. You can optionally specify an `INSERT_METHOD` option if you want inserts for the `MERGE` table to take place in the first or last table of the `UNION` list. Use a value of `FIRST` or `LAST` to cause inserts to be made in the first or last table, respectively. If you do not specify an `INSERT_METHOD` option or if you specify it with a value of `NO`, attempts to insert rows into

the `MERGE` table result in an error.

The following example shows how to create a `MERGE` table:

```
mysql> CREATE TABLE t1 (
->   a INT NOT NULL AUTO_INCREMENT PRIMARY KEY,
->   message CHAR(20)) ENGINE=MyISAM;
mysql> CREATE TABLE t2 (
->   a INT NOT NULL AUTO_INCREMENT PRIMARY KEY,
->   message CHAR(20)) ENGINE=MyISAM;
mysql> INSERT INTO t1 (message) VALUES ('Testing'),('table'),('t1');
mysql> INSERT INTO t2 (message) VALUES ('Testing'),('table'),('t2');
mysql> CREATE TABLE total (
->   a INT NOT NULL AUTO_INCREMENT,
->   message CHAR(20), INDEX(a))
->   ENGINE=MERGE UNION=(t1,t2) INSERT_METHOD=LAST;
```

The older term `TYPE` is supported as a synonym for `ENGINE` for backward compatibility, but `ENGINE` is the preferred term from MySQL 4.0.18 on and `TYPE` is deprecated.

Note that the `a` column is indexed as a `PRIMARY KEY` in the underlying `MyISAM` tables, but not in the `MERGE` table. There it is indexed but not as a `PRIMARY KEY` because a `MERGE` table cannot enforce uniqueness over the set of underlying tables.

After creating the `MERGE` table, you can issue queries that operate on the group of tables as a whole:

```
mysql> SELECT * FROM total;
+----+-----+
| a | message |
+----+-----+
| 1 | Testing |
| 2 | table   |
| 3 | t1      |
| 1 | Testing |
| 2 | table   |
| 3 | t2      |
+----+-----+
```

To remap a `MERGE` table to a different collection of `MyISAM` tables, you can use one of the following methods:

- `DROP` the `MERGE` table and re-create it.
- Use `ALTER TABLE tbl_name UNION=(...)` to change the list of underlying tables.

`MERGE` tables can help you solve the following problems:

- Easily manage a set of log tables. For example, you can put data from different months into separate tables, compress some of them with `myisampack`, and then create a `MERGE` table to use them as one.
- Obtain more speed. You can split a big read-only table based on some criteria, and then put individual tables on different disks. A `MERGE` table on this could be much faster than using the big table.
- Perform more efficient searches. If you know exactly what you are looking for, you can search in just one of the split tables for some queries and use a `MERGE` table for others. You can even have many different `MERGE` tables that use overlapping sets of tables.
- Perform more efficient repairs. It is easier to repair individual tables that are mapped to a `MERGE` table than to repair a single large table.
- Instantly map many tables as one. A `MERGE` table need not maintain an index of its own because it

uses the indexes of the individual tables. As a result, **MERGE** table collections are *very* fast to create or remap. (Note that you must still specify the index definitions when you create a **MERGE** table, even though no indexes are created.)

- If you have a set of tables from which you create a large table on demand, you should instead create a **MERGE** table on them on demand. This is much faster and saves a lot of disk space.
- Exceed the file size limit for the operating system. Each **MyISAM** table is bound by this limit, but a collection of **MyISAM** tables is not.
- You can create an alias or synonym for a **MyISAM** table by defining a **MERGE** table that maps to that single table. There should be no really notable performance impact from doing this (only a couple of indirect calls and `memcpy()` calls for each read).

The disadvantages of **MERGE** tables are:

- You can use only identical **MyISAM** tables for a **MERGE** table.
- You cannot use a number of **MyISAM** features in **MERGE** tables. For example, you cannot create **FULLTEXT** indexes on **MERGE** tables. (You can, of course, create **FULLTEXT** indexes on the underlying **MyISAM** tables, but you cannot search the **MERGE** table with a full-text search.)
- If the **MERGE** table is non-temporary, all underlying **MyISAM** tables must be non-temporary, too. If the **MERGE** table is temporary, the **MyISAM** tables can be any mix of temporary and non-temporary.
- **MERGE** tables use more file descriptors. If 10 clients are using a **MERGE** table that maps to 10 tables, the server uses  $(10 \times 10) + 10$  file descriptors. (10 data file descriptors for each of the 10 clients, and 10 index file descriptors shared among the clients.)
- Key reads are slower. When you read a key, the **MERGE** storage engine needs to issue a read on all underlying tables to check which one most closely matches the given key. To read the next key, the **MERGE** storage engine needs to search the read buffers to find the next key. Only when one key buffer is used up does the storage engine need to read the next key block. This makes **MERGE** keys much slower on `eq_ref` searches, but not much slower on `ref` searches. See [Section 7.2.1, “Optimizing Queries with EXPLAIN”](#), for more information about `eq_ref` and `ref`.

#### Additional resources

- A forum dedicated to the **MERGE** storage engine is available at <http://forums.mysql.com/list.php?93>.

### 14.3.1. **MERGE** Table Problems

The following are known problems with **MERGE** tables:

- If you use `ALTER TABLE` to change a **MERGE** table to another storage engine, the mapping to the underlying tables is lost. Instead, the rows from the underlying **MyISAM** tables are copied into the altered table, which then uses the specified storage engine.
- Before MySQL 4.1.1, all underlying tables and the **MERGE** table itself had to be in the same database.
- `REPLACE` does not work.



- You cannot use `DROP TABLE`, `ALTER TABLE`, `DELETE` without a `WHERE` clause, `REPAIR TABLE`, `TRUNCATE TABLE`, `OPTIMIZE TABLE`, or `ANALYZE TABLE` on any of the tables that are mapped into an open `MERGE` table. If you do so, the `MERGE` table may still refer to the original table, which yields unexpected results. The easiest way to work around this deficiency is to ensure that no `MERGE` tables remain open by issuing a `FLUSH TABLES` statement prior to performing any of those operations.
- `DROP TABLE` on a table that is in use by a `MERGE` table does not work on Windows because the `MERGE` storage engine's table mapping is hidden from the upper layer of MySQL. Windows does not allow open files to be deleted, so you first must flush all `MERGE` tables (with `FLUSH TABLES`) or drop the `MERGE` table before dropping the table.
- A `MERGE` table cannot maintain uniqueness constraints over the entire table. When you perform an `INSERT`, the data goes into the first or last `MyISAM` table (depending on the value of the `INSERT_METHOD` option). MySQL ensures that unique key values remain unique within that `MyISAM` table, but not across all the tables in the collection.
- Before MySQL 3.23.49, `DELETE FROM merge_table` used without a `WHERE` clause only clears the mapping for the table. That is, it incorrectly empties the `.MRG` file rather than deleting records from the mapped tables.
- Using `RENAME TABLE` on an active `MERGE` table may corrupt the table. This is fixed in MySQL 4.1.x.
- When you create or alter `MERGE` table, there is no check to ensure that the underlying tables are existing `MyISAM` tables and have identical structures. When the `MERGE` table is used, MySQL checks that the row length for all mapped tables is equal, but this is not foolproof. If you create a `MERGE` table from dissimilar `MyISAM` tables, you are very likely to run into strange problems.

Similarly, if you create a `MERGE` table from non-`MyISAM` tables, or if you drop an underlying table or alter it to be a non-`MyISAM` table, no error for the `MERGE` table occurs until later when you attempt to use it.

- The order of indexes in the `MERGE` table and its underlying tables should be the same. If you use `ALTER TABLE` to add a `UNIQUE` index to a table used in a `MERGE` table, and then use `ALTER TABLE` to add a non-unique index on the `MERGE` table, the index ordering is different for the tables if there was already a non-unique index in the underlying table. (This happens because `ALTER TABLE` puts `UNIQUE` indexes before non-unique indexes to facilitate rapid detection of duplicate keys.) Consequently, queries on tables with such indexes may return unexpected results.
- If you encounter an error message similar to `ERROR 1017 (HY000): Can't find file: 'mm.MRG' (errno: 2)` it generally indicates that some of the base tables are not using the `MyISAM` storage engine. Confirm that all tables are `MyISAM`.

## 14.4. The `MEMORY (HEAP)` Storage Engine

The `MEMORY` storage engine creates tables with contents that are stored in memory. Before MySQL 4.1, `MEMORY` tables are called `HEAP` tables. As of 4.1, `MEMORY` is the preferred term, although `HEAP` remains supported for backward compatibility.

Each `MEMORY` table is associated with one disk file. The filename begins with the table name and has an extension of `.frm` to indicate that it stores the table definition.

To specify explicitly that you want to create a `MEMORY` table, indicate that with an `ENGINE` table option:

```
CREATE TABLE t (i INT) ENGINE = MEMORY;
```

The older term `TYPE` is supported as a synonym for `ENGINE` for backward compatibility, but `ENGINE` is the preferred term from MySQL 4.0.18 on and `TYPE` is deprecated.

As indicated by the name, `MEMORY` tables are stored in memory. They use hash indexes by default, which makes them very fast, and very useful for creating temporary tables. However, when the server shuts down, all rows stored in `MEMORY` tables are lost. The tables themselves continue to exist because their definitions are stored in `.frm` files on disk, but they are empty when the server restarts.

This example shows how you might create, use, and remove a `MEMORY` table:

```
mysql> CREATE TABLE test TYPE=MEMORY
-> SELECT ip,SUM(downloads) AS down
-> FROM log_table GROUP BY ip;
mysql> SELECT COUNT(ip),AVG(down) FROM test;
mysql> DROP TABLE test;
```

`MEMORY` tables have the following characteristics:

- Space for `MEMORY` tables is allocated in small blocks. Tables use 100% dynamic hashing for inserts. No overflow area or extra key space is needed. No extra space is needed for free lists. Deleted rows are put in a linked list and are reused when you insert new data into the table. `MEMORY` tables also have none of the problems commonly associated with deletes plus inserts in hashed tables.
- `MEMORY` tables allow up to 32 indexes per table and 16 columns per index. Previously, the maximum key length supported by this storage engine was 255 bytes; as of MySQL 4.1.13, `MEMORY` tables support a maximum key length of 500 bytes. (See [Section D.1.20](#), “Changes in release 4.1.3 (28 June 2004: Beta)”.)
- Before MySQL 4.1, the `MEMORY` storage engine implements only hash indexes. From MySQL 4.1 on, hash indexes are still the default, but you can specify explicitly that a `MEMORY` table index should be a `HASH` or `BTREE` by adding a `USING` clause as shown here:

```
CREATE TABLE lookup
(id INT, INDEX USING HASH (id))
ENGINE = MEMORY;
CREATE TABLE lookup
(id INT, INDEX USING BTREE (id))
ENGINE = MEMORY;
```

General characteristics of B-tree and hash indexes are described in [Section 7.4.5](#), “How MySQL Uses Indexes”.

- You can have non-unique keys in a `MEMORY` table. (This is an uncommon feature for implementations of hash indexes.)
- If you have a hash index on a `MEMORY` table that has a high degree of key duplication (many index entries containing the same value), updates to the table that affect key values and all deletes are significantly slower. The degree of this slowdown is proportional to the degree of duplication (or, inversely proportional to the index cardinality). You can use a `BTREE` index to avoid this problem.
- As of MySQL 4.0.2, columns that are indexed can contain `NULL` values.
- `MEMORY` tables use a fixed-length row storage format.
- `MEMORY` tables cannot contain `BLOB` or `TEXT` columns.
- `MEMORY` supports `AUTO_INCREMENT` columns as of MySQL 4.1.0.
- As of MySQL 4.1, you can use `INSERT DELAYED` with `MEMORY` tables. See [Section 13.2.4.2](#), “`INSERT DELAYED` Syntax”.

- `MEMORY` tables are shared among all clients (just like any other non-`TEMPORARY` table).
- `MEMORY` table contents are stored in memory, which is a property that `MEMORY` tables share with internal tables that the server creates on the fly while processing queries. However, the two types of tables differ in that `MEMORY` tables are not subject to storage conversion, whereas internal tables are:
  - If an internal table becomes too large, the server automatically converts it to an on-disk table. The size limit is determined by the value of the `tmp_table_size` system variable.
  - `MEMORY` tables are never converted to disk tables. To ensure that you don't accidentally do anything foolish, you can set the `max_heap_table_size` system variable to impose a maximum size on `MEMORY` tables. For individual tables, you can also specify a `MAX_ROWS` table option in the `CREATE TABLE` statement.
- The server needs sufficient memory to maintain all `MEMORY` tables that are in use at the same time.
- To free memory used by a `MEMORY` table when you no longer require its contents, you should execute `DELETE` or `TRUNCATE TABLE`, or remove the table altogether using `DROP TABLE`.
- If you want to populate a `MEMORY` table when the MySQL server starts, you can use the `--init-file` option. For example, you can put statements such as `INSERT INTO ... SELECT` or `LOAD DATA INFILE` into this file to load the table from a persistent data source. See [Section 5.2.1, “Command Options”](#), and [Section 13.2.5, “LOAD DATA INFILE Syntax”](#).
- If you are using replication, the master server's `MEMORY` tables become empty when it is shut down and restarted. However, a slave is not aware that these tables have become empty, so it returns out-of-date content if you select data from them. Beginning with MySQL 4.0.18, when a `MEMORY` table is used on the master for the first time since the master was started, a `DELETE FROM` statement is written to the master's binary log automatically, thus synchronizing the slave to the master again. Note that even with this strategy, the slave still has outdated data in the table during the interval between the master's restart and its first use of the table. However, if you use the `--init-file` option to populate the `MEMORY` table on the master at startup, it ensures that this time interval is zero.
- The memory needed for one row in a `MEMORY` table is calculated using the following expression:

```
SUM_OVER_ALL_BTREE_KEYS(max_length_of_key + sizeof(char*) × 4)
+ SUM_OVER_ALL_HASH_KEYS(sizeof(char*) × 2)
+ ALIGN(length_of_row+1, sizeof(char*))
```

`ALIGN()` represents a round-up factor to cause the row length to be an exact multiple of the `char` pointer size. `sizeof(char*)` is 4 on 32-bit machines and 8 on 64-bit machines.

#### Additional resources

- A forum dedicated to the `MEMORY` storage engine is available at <http://forums.mysql.com/list.php?92>.

## 14.5. The BDB (BerkeleyDB) Storage Engine

Sleepycat Software has provided MySQL with the Berkeley DB transactional storage engine. This storage engine typically is called `BDB` for short. `BDB` tables may have a greater chance of surviving crashes and are also capable of `COMMIT` and `ROLLBACK` operations on transactions.

Support for the `BDB` storage engine is included in MySQL source distributions starting from version

3.23.34a and is activated in MySQL-Max binary distributions. The MySQL source distribution comes with a BDB distribution that is patched to make it work with MySQL. You cannot use a non-patched version of BDB with MySQL.

We at MySQL AB work in close cooperation with Sleepycat to keep the quality of the MySQL/BDB interface high. (Even though Berkeley DB is in itself very tested and reliable, the MySQL interface is still considered gamma quality. We continue to improve and optimize it.)

When it comes to support for any problems involving BDB tables, we are committed to helping our users locate the problem and create reproducible test cases. Any such test case is forwarded to Sleepycat, which in turn helps us find and fix the problem. As this is a two-stage operation, any problems with BDB tables may take a little longer for us to fix than for other storage engines. However, we anticipate no significant difficulties with this procedure because the Berkeley DB code itself is used in many applications other than MySQL.

For general information about Berkeley DB, please visit the Sleepycat Web site, <http://www.sleepycat.com/>.

## 14.5.1. Operating Systems Supported by BDB

Currently, we know that the BDB storage engine works with the following operating systems:

- Linux 2.x Intel
- Sun Solaris (SPARC and x86)
- FreeBSD 4.x/5.x (x86, sparc64)
- IBM AIX 4.3.x
- SCO OpenServer
- SCO UnixWare 7.1.x
- Windows NT/2000/XP

The BDB storage engine does *not* work with the following operating systems:

- Linux 2.x Alpha
- Linux 2.x AMD64
- Linux 2.x IA-64
- Linux 2.x s390
- Mac OS X

**Note:** The preceding lists are not complete. We update them as we receive more information.

If you build MySQL from source with support for BDB tables, but the following error occurs when you start `mysqld`, it means that the BDB storage engine is not supported for your architecture:

```
bdb: architecture lacks fast mutexes: applications cannot be threaded
Can't init databases
```

In this case, you must rebuild MySQL without `BDB` support or start the server with the `--skip-bdb` option.

## 14.5.2. Installing `BDB`

If you have downloaded a binary version of MySQL that includes support for Berkeley DB, simply follow the usual binary distribution installation instructions. (MySQL-Max distributions include `BDB` support.)

If you build MySQL from source, you can enable `BDB` support by invoking `configure` with the `--with-berkeley-db` option in addition to any other options that you normally use. Download a distribution for MySQL 3.23.34 or newer, change location into its top-level directory, and run this command:

```
shell> ./configure --with-berkeley-db [other-options]
```

For more information, see [Section 5.3, “The `mysqld-max` Extended MySQL Server”](#), [Section 2.8, “Installing MySQL on Other Unix-Like Systems”](#), and [Section 2.9, “MySQL Installation Using a Source Distribution”](#).

## 14.5.3. `BDB` Startup Options

The following options to `mysqld` can be used to change the behavior of the `BDB` storage engine. For more information, see [Section 5.2.1, “Command Options”](#).

- `--bdb-home=path`

The base directory for `BDB` tables. This should be the same directory that you use for `--datadir`.

- `--bdb-lock-detect=method`

The `BDB` lock detection method. The option value should be `DEFAULT`, `OLDEST`, `RANDOM`, or `YOUNGEST`.

- `--bdb-logdir=file_name`

The `BDB` log file directory.

- `--bdb-no-recover`

Do not start Berkeley DB in recover mode.

- `--bdb-no-sync`

Don't synchronously flush the `BDB` logs. This option is deprecated as of MySQL 4.0.18; use `--skip-sync-bdb-logs` instead (see the description for `--sync-bdb-logs`).

- `--bdb-shared-data`

Start Berkeley DB in multi-process mode. (Do not use `DB_PRIVATE` when initializing Berkeley DB.)

- `--bdb-tmpdir=path`

The `BDB` temporary file directory.

- `--skip-bdb`

Disable the [BDB](#) storage engine.

- `--sync-bdb-logs`

Synchronously flush the [BDB](#) logs. This option is enabled by default. Use `--skip-sync-bdb-logs` to disable it. This option was added in MySQL 4.0.18.

If you use the `--skip-bdb` option, MySQL does not initialize the Berkeley DB library and this saves a lot of memory. However, if you use this option, you cannot use [BDB](#) tables. If you try to create a [BDB](#) table, MySQL uses the default storage engine instead.

Normally, you should start `mysqld` without the `--bdb-no-recover` option if you intend to use [BDB](#) tables. However, this may cause problems when you try to start `mysqld` if the [BDB](#) log files are corrupted. See [Section 2.10.2.3, “Starting and Troubleshooting the MySQL Server”](#).

With the `bdb_max_lock` variable, you can specify the maximum number of locks that can be active on a [BDB](#) table. The default is 10,000. You should increase this if errors such as the following occur when you perform long transactions or when `mysqld` has to examine many rows to execute a query:

```
bdb: Lock table is out of available locks
Got error 12 from ...
```

You may also want to change the `binlog_cache_size` and `max_binlog_cache_size` variables if you are using large multiple-statement transactions. See [Section 5.11.4, “The Binary Log”](#).

See also [Section 5.2.2, “System Variables”](#).

## 14.5.4. Characteristics of [BDB](#) Tables

Each [BDB](#) table is stored on disk in two files. The files have names that begin with the table name and have an extension to indicate the file type. An `.frm` file stores the table format, and a `.db` file contains the table data and indexes.

To specify explicitly that you want a [BDB](#) table, indicate that with an [ENGINE](#) table option:

```
CREATE TABLE t (i INT) ENGINE = BDB;
```

The older term [TYPE](#) is supported as a synonym for [ENGINE](#) for backward compatibility, but [ENGINE](#) is the preferred term from MySQL 4.0.18 on and [TYPE](#) is deprecated.

[BerkeleyDB](#) is a synonym for [BDB](#) in the [ENGINE](#) table option.

The [BDB](#) storage engine provides transactional tables. The way you use these tables depends on the autocommit mode:

- If you are running with autocommit enabled (which is the default), changes to [BDB](#) tables are committed immediately and cannot be rolled back.
- If you are running with autocommit disabled, changes do not become permanent until you execute a [COMMIT](#) statement. Instead of committing, you can execute [ROLLBACK](#) to forget the changes.

You can start a transaction with the [START TRANSACTION](#) or [BEGIN](#) statement to suspend autocommit, or with [SET AUTOCOMMIT=0](#) to disable autocommit explicitly.

For more information about transactions, see [Section 13.4.1, “START TRANSACTION, COMMIT, and ROLLBACK Syntax”](#).

The [BDB](#) storage engine has the following characteristics:

- [BDB](#) tables can have up to 31 indexes per table, 16 columns per index, and a maximum key size of 1024 bytes (500 bytes before MySQL 4.0).
- MySQL requires a primary key in each [BDB](#) table so that each row can be uniquely identified. If you don't create one explicitly by declaring a [PRIMARY KEY](#), MySQL creates and maintains a hidden primary key for you. The hidden key has a length of five bytes and is incremented for each insert attempt. This key does not appear in the output of [SHOW CREATE TABLE](#) or [DESCRIBE](#).
- The primary key is faster than any other index, because it is stored together with the row data. The other indexes are stored as the key data plus the primary key, so it's important to keep the primary key as short as possible to save disk space and get better speed.

This behavior is similar to that of [InnoDB](#), where shorter primary keys save space not only in the primary index but in secondary indexes as well.

- If all columns that you access in a [BDB](#) table are part of the same index or part of the primary key, MySQL can execute the query without having to access the actual row. In a [MyISAM](#) table, this can be done only if the columns are part of the same index.
- Sequential scanning is slower for [BDB](#) tables than for [MyISAM](#) tables because the data in [BDB](#) tables is stored in B-trees and not in a separate data file.
- Key values are not prefix- or suffix-compressed like key values in [MyISAM](#) tables. In other words, key information takes a little more space in [BDB](#) tables compared to [MyISAM](#) tables.
- There are often holes in the [BDB](#) table to allow you to insert new rows in the middle of the index tree. This makes [BDB](#) tables somewhat larger than [MyISAM](#) tables.
- [SELECT COUNT\(\\*\) FROM tbl\\_name](#) is slow for [BDB](#) tables, because no row count is maintained in the table.
- The optimizer needs to know the approximate number of rows in the table. MySQL solves this by counting inserts and maintaining this in a separate segment in each [BDB](#) table. If you don't issue a lot of [DELETE](#) or [ROLLBACK](#) statements, this number should be accurate enough for the MySQL optimizer. However, MySQL stores the number only on close, so it may be incorrect if the server terminates unexpectedly. It should not be fatal even if this number is not 100% correct. You can update the row count by using [ANALYZE TABLE](#) or [OPTIMIZE TABLE](#). See [Section 13.5.2.1, “ANALYZE TABLE Syntax”](#), and [Section 13.5.2.5, “OPTIMIZE TABLE Syntax”](#).
- Internal locking in [BDB](#) tables is done at the page level.
- [LOCK TABLES](#) works on [BDB](#) tables as with other tables. If you do not use [LOCK TABLES](#), MySQL issues an internal multiple-write lock on the table (a lock that does not block other writers) to ensure that the table is properly locked if another thread issues a table lock.
- To support transaction rollback, the [BDB](#) storage engine maintains log files. For maximum performance, you can use the [--bdb-logdir](#) option to place the [BDB](#) logs on a different disk than the one where your databases are located.
- MySQL performs a checkpoint each time a new [BDB](#) log file is started, and removes any [BDB](#) log files that are not needed for current transactions. You can also use [FLUSH LOGS](#) at any time to checkpoint the Berkeley DB tables.

For disaster recovery, you should use table backups plus MySQL's binary log. See [Section 5.9.1, “Database Backups”](#).

**Warning:** If you delete old log files that are still in use, **BDB** is not able to do recovery at all and you may lose data if something goes wrong.

- Applications must always be prepared to handle cases where any change of a **BDB** table may cause an automatic rollback and any read may fail with a deadlock error.
- If you get a full disk with a **BDB** table, you get an error (probably error 28) and the transaction should roll back. This contrasts with **MyISAM** and **ISAM** tables, for which `mysqld` waits for sufficient free disk space before continuing.

### 14.5.5. Restrictions on **BDB** Tables

The following list indicates restrictions that you must observe when using **BDB** tables:

- Each **BDB** table stores in its `.db` file the path to the file as it was created. This is done to enable detection of locks in a multi-user environment that supports symlinks. As a consequence of this, it is not possible to move **BDB** table files from one database directory to another.
- When making backups of **BDB** tables, you must either use `mysqldump` or else make a backup that includes the files for each **BDB** table (the `.frm` and `.db` files) as well as the **BDB** log files. The **BDB** storage engine stores unfinished transactions in its log files and requires them to be present when `mysqld` starts. The **BDB** logs are the files in the data directory with names of the form `log.NNNNNNNNNN` (ten digits).
- If a column that allows **NULL** values has a unique index, only a single **NULL** value is allowed. This differs from other storage engines, which allow multiple **NULL** values in unique indexes.

### 14.5.6. Errors That May Occur When Using **BDB** Tables

- If the following error occurs when you start `mysqld` after upgrading, it means that the current version of **BDB** doesn't support the old log file format:

```
bdb: Ignoring log file: ../log.NNNNNNNNNN:
unsupported log version #
```

In this case, you must delete all **BDB** logs from your data directory (the files that have names of the form `log.NNNNNNNNNN`) and restart `mysqld`. We also recommend that you then use `mysql-dump --opt` to dump your **BDB** tables, drop the tables, and restore them from the dump file.

- If autocommit mode is disabled and you drop a **BDB** table that is referenced in another transaction, you may get error messages of the following form in your MySQL error log:

```
001119 23:43:56 bdb: Missing log fileid entry
001119 23:43:56 bdb: txn_abort: Log undo failed for LSN:
1 3644744: Invalid
```

This is not fatal, but the fix is not trivial. Until the problem is fixed, we recommend that you not drop **BDB** tables except while autocommit mode is enabled.



## 14.6. The **EXAMPLE** Storage Engine

The **EXAMPLE** storage engine was added in MySQL 4.1.3. It is a “stub” engine that does nothing. Its purpose is to serve as an example in the MySQL source code that illustrates how to begin writing new storage engines. As such, it is primarily of interest to developers.

The **EXAMPLE** storage engine is included in MySQL-Max binary distributions. To enable this storage engine if you build MySQL from source, invoke `configure` with the `-with-example-storage-engine` option.

To examine the source for the **EXAMPLE** engine, look in the `sql/examples` directory of a MySQL source distribution.

When you create an **EXAMPLE** table, the server creates a table format file in the database directory. The file begins with the table name and has an `.frm` extension. No other files are created. No data can be stored into the table. Retrievals return an empty result.

```
mysql> CREATE TABLE test (i INT) ENGINE = EXAMPLE;
Query OK, 0 rows affected (0.78 sec)

mysql> INSERT INTO test VALUES(1),(2),(3);
ERROR 1031 (HY000): Table storage engine for 'test' doesn't have this option

mysql> SELECT * FROM test;
Empty set (0.31 sec)
```

The **EXAMPLE** storage engine does not support indexing.

## 14.7. The **ARCHIVE** Storage Engine

The **ARCHIVE** storage engine was added in MySQL 4.1.3. It is used for storing large amounts of data without indexes in a very small footprint.

The **ARCHIVE** storage engine is included in MySQL binary distributions. To enable this storage engine if you build MySQL from source, invoke `configure` with the `-with-archive-storage-engine` option.

To examine the source for the **ARCHIVE** engine, look in the `sql` directory of a MySQL source distribution.

You can check whether the **ARCHIVE** storage engine is available with this statement:

```
mysql> SHOW VARIABLES LIKE 'have_archive';
```

When you create an **ARCHIVE** table, the server creates a table format file in the database directory. The file begins with the table name and has an `.frm` extension. The storage engine creates other files, all having names beginning with the table name. The data and metadata files have extensions of `.ARZ` and `.ARM`, respectively. An `.ARN` file may appear during optimization operations.

The **ARCHIVE** engine supports `INSERT` and `SELECT`, but not `DELETE`, `REPLACE`, or `UPDATE`. It does support `ORDER BY` operations, `BLOB` columns, and basically all but spatial data types (see [Section 16.4.1, “MySQL Spatial Data Types”](#)). The **ARCHIVE** engine uses row-level locking.

**Storage:** Rows are compressed as they are inserted. The **ARCHIVE** engine uses `zlib` lossless data compression (see <http://www.zlib.net/>). You can use `OPTIMIZE TABLE` to analyze the table and pack it into a smaller format (for a reason to use `OPTIMIZE TABLE`, see later in this section). There are several types of insertions that are used:

- An `INSERT` statement just pushes rows into a compression buffer, and that buffer flushes as necessary. The insertion into the buffer is protected by a lock. A `SELECT` forces a flush to occur, unless the only insertions that have come in were `INSERT DELAYED` (those flush as necessary). See [Section 13.2.4.2, “INSERT DELAYED Syntax”](#).
- A bulk insert is visible only after it completes, unless other inserts occur at the same time, in which case it can be seen partially. A `SELECT` never causes a flush of a bulk insert unless a normal insert occurs while it is loading.

**Retrieval:** On retrieval, rows are uncompressed on demand; there is no row cache. A `SELECT` operation performs a complete table scan: When a `SELECT` occurs, it finds out how many rows are currently available and reads that number of rows. `SELECT` is performed as a consistent read. Note that lots of `SELECT` statements during insertion can deteriorate the compression, unless only bulk or delayed inserts are used. To achieve better compression, you can use `OPTIMIZE TABLE` or `REPAIR TABLE`. The number of rows in `ARCHIVE` tables reported by `SHOW TABLE STATUS` is always accurate. See [Section 13.5.2.5, “OPTIMIZE TABLE Syntax”](#), [Section 13.5.2.6, “REPAIR TABLE Syntax”](#), and [Section 13.5.4.18, “SHOW TABLE STATUS Syntax”](#).

#### Additional resources

- A forum dedicated to the `ARCHIVE` storage engine is available at <http://forums.mysql.com/list.php?112>.

## 14.8. The `CSV` Storage Engine

The `CSV` storage engine was added in MySQL 4.1.4. This engine stores data in text files using comma-separated values format.

To enable this storage engine, use the `--with-csv-storage-engine` option to `configure` when you build MySQL.

The `CSV` storage engine is included in MySQL-Max binary distributions. To enable this storage engine if you build MySQL from source, invoke `configure` with the `--with-csv-storage-engine` option.

To examine the source for the `CSV` engine, look in the `sql/examples` directory of a MySQL source distribution.

When you create a `CSV` table, the server creates a table format file in the database directory. The file begins with the table name and has an `.frm` extension. The storage engine also creates a data file. Its name begins with the table name and has a `.CSV` extension. The data file is a plain text file. When you store data into the table, the storage engine saves it into the data file in comma-separated values format.

```
mysql> CREATE TABLE test(i INT, c CHAR(10)) ENGINE = CSV;
Query OK, 0 rows affected (0.12 sec)

mysql> INSERT INTO test VALUES(1,'record one'),(2,'record two');
Query OK, 2 rows affected (0.00 sec)
Records: 2 Duplicates: 0 Warnings: 0

mysql> SELECT * FROM test;
+-----+-----+
| i   | c           |
+-----+-----+
| 1   | record one  |
| 2   | record two  |
+-----+-----+
2 rows in set (0.00 sec)
```

If you examine the `test.CSV` file in the database directory created by executing the preceding statements, its contents should look like this:

```
"1","record one"
"2","record two"
```

This format can be read, and even written, by spreadsheet applications such as Microsoft Excel or StarOffice Calc.

The `CSV` storage engine does not support indexing.

## 14.9. The `BLACKHOLE` Storage Engine

The `BLACKHOLE` storage engine was added in MySQL 4.1.11. This engine acts as a “black hole” that accepts data but throws it away and does not store it. Retrievals always return an empty result:

```
mysql> CREATE TABLE test(i INT, c CHAR(10)) ENGINE = BLACKHOLE;
Query OK, 0 rows affected (0.03 sec)

mysql> INSERT INTO test VALUES(1,'record one'),(2,'record two');
Query OK, 2 rows affected (0.00 sec)
Records: 2 Duplicates: 0 Warnings: 0

mysql> SELECT * FROM test;
Empty set (0.00 sec)
```

The `BLACKHOLE` storage engine is included in MySQL-Max binary distributions. To enable this storage engine if you build MySQL from source, invoke `configure` with the `-with-blackhole-storage-engine` option.

To examine the source for the `BLACKHOLE` engine, look in the `sql` directory of a MySQL source distribution.


When you create a `BLACKHOLE` table, the server creates a table format file in the database directory. The file begins with the table name and has an `.frm` extension. There are no other files associated with the table.

The `BLACKHOLE` storage engine supports all kinds of indexes. That is, you can include index declarations in the table definition.

You can check whether the `BLACKHOLE` storage engine is available with this statement:

```
mysql> SHOW VARIABLES LIKE 'have_blackhole_engine';
```

Inserts into a `BLACKHOLE` table do not store any data, but if the binary log is enabled, the SQL statements are logged (and replicated to slave servers). This can be useful as a repeater or filter mechanism. For example, suppose that your application requires slave-side filtering rules, but transferring all binary log data to the slave first results in too much traffic. In such a case, it is possible to set up on the master host a “dummy” slave process whose default storage engine is `BLACKHOLE`, depicted as follows:



The master writes to its binary log. The “dummy” `mysqld` process acts as a slave, applying the desired combination of `replicate-do-*` and `replicate-ignore-*` rules, and writes a new, filtered binary log of its own. (See [Section 6.8, “Replication Startup Options”](#).) This filtered log is provided to the slave.

The dummy process does not actually store any data, so there is little processing overhead incurred by running the additional `mysqld` process on the replication master host. This type of setup can be repeated with additional replication slaves.

Other possible uses for the `BLACKHOLE` storage engine include:

- Verification of dump file syntax.
- Measurement of the overhead from binary logging, by comparing performance using `BLACKHOLE` with and without binary logging enabled.
- `BLACKHOLE` is essentially a “no-op” storage engine, so it could be used for finding performance bottlenecks not related to the storage engine itself.

## 14.10. The `ISAM` Storage Engine

The original storage engine in MySQL was the `ISAM` engine. It was the only storage engine available until MySQL 3.23, when the improved `MyISAM` engine was introduced as the default. `ISAM` is deprecated. As of MySQL 4.1, it's included in the source but not enabled in binary distributions. It is not available in MySQL 5.0. Embedded MySQL server versions do not support `ISAM` tables by default.

Due to the deprecated status of `ISAM`, and because `MyISAM` is an improvement over `ISAM`, you are advised to convert any remaining `ISAM` tables to `MyISAM` as soon as possible. To convert an `ISAM` table to a `MyISAM` table, use an `ALTER TABLE` statement:

```
mysql> ALTER TABLE tbl_name TYPE = MYISAM;
```

For more information about `MyISAM`, see [Section 14.1, “The MyISAM Storage Engine”](#).

Each `ISAM` table is stored on disk in three files. The files have names that begin with the table name and have an extension to indicate the file type. An `.frm` file stores the table definition. The data file has an `.ISD` extension. The index file has an `.ISM` extension.

ISAM uses B-tree indexes.

You can check or repair ISAM tables with the `isamchk` utility. See [Section 5.9.4.1, “Using `myisamchk` for Crash Recovery”](#).

ISAM has the following properties:

- Compressed and fixed-length keys
- Fixed and dynamic record length
- 16 indexes per table, with 16 key parts per key
- Maximum key length 256 bytes (default)
- Data values are stored in machine format; this is fast, but machine/OS dependent

Many of the properties of MyISAM tables are also true for ISAM tables. However, there are also many differences. The following list describes some of the ways that ISAM is distinct from MyISAM:

- Not binary portable across OS/platforms.
- Can't handle tables larger than 4GB.
- Only supports prefix compression on strings.
- Smaller (more restrictive) key limits.
- Dynamic tables become more fragmented.
- Doesn't support MERGE tables.
- Tables are checked and repaired with `isamchk` rather than with `myisamchk`.
- Tables are compressed with `pack_isam` rather than with `myisampack`.
- Cannot be used with the `BACKUP TABLE` or `RESTORE TABLE` backup-related statements.
- Cannot be used with the `CHECK TABLE`, `REPAIR TABLE`, `OPTIMIZE TABLE`, or `ANALYZE TABLE` table-maintenance statements.
- No support for full-text searching or spatial data types.
- No support for multiple character sets per table.
- Indexes cannot be assigned to specific key caches.

---

# Chapter 15. MySQL Cluster

MySQL Cluster is a high-availability, high-redundancy version of MySQL adapted for the distributed computing environment. It uses the [NDB Cluster](#) storage engine to enable running several MySQL servers in a cluster. This storage engine is available and in binary releases from MySQL-Max 4.1.3. Beginning with MySQL 4.1.10a, it is also available in RPMs compatible with most modern Linux distributions. (If you install using RPM files, note that both the [mysql-server](#) and [mysql-max](#) RPMs must be installed to have MySQL Cluster capability.)

MySQL Cluster is currently available and supported on the following platforms:

- **Linux:** x86, AMD64, EMT64, s/390, PPC, Alpha, SPARC, UltraSparc
- **Solaris:** SPARC, UltraSparc, x86, AMD64, EMT64
- **BSD** (FreeBSD, NetBSD, OpenBSD): x86, AMD64, EMT64, PPC
- **Mac OS X:** PPC
- **HP-UX:** PA-RISC
- **Tru64:** Alpha
- **OpenVMS:** Alpha
- **IRIX:** MIPS
- **Novell Netware:** x86
- **QNX Neutrino:** x86
- **SCO OpenServer, OpenUnix, UnixWare:** x86

For exact levels of support available for MySQL Cluster on specific combinations of operating system versions, operating system distributions, and hardware platforms, please refer to the [Cluster Supported Platforms list](http://www.mysql.com/support/clusterplatforms.html) [http://www.mysql.com/support/clusterplatforms.html] maintained by the MySQL Support Team on the MySQL AB website.

MySQL Cluster is *not* currently supported on Microsoft Windows. We are working to make Cluster available on all operating systems supported by MySQL, including Windows, and will update the information provided here as this work continues.

This chapter represents a work in progress, and its contents are subject to revision as MySQL Cluster continues to evolve. Additional information regarding MySQL Cluster can be found on the MySQL AB Web site at <http://www.mysql.com/products/cluster/>.

## Additional resources

- Answers to some commonly asked questions about Cluster may be found in the [Section 15.12, “MySQL Cluster FAQ”](#).
- The MySQL Cluster mailing list: <http://lists.mysql.com/cluster>.
- The MySQL Cluster Forum: <http://forums.mysql.com/list.php?25>.

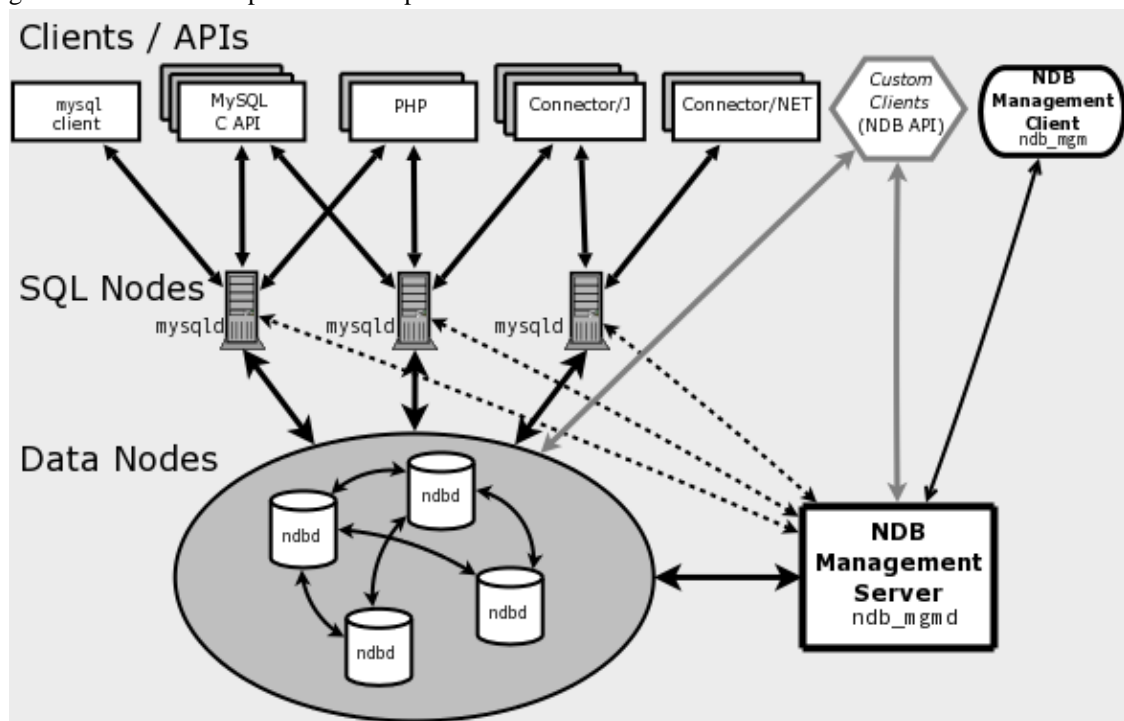
- Many MySQL Cluster users and some of the MySQL Cluster developers blog about their experiences with Cluster, and make feeds of these available through [PlanetMySQL](http://www.planetmysql.org/) [<http://www.planetmysql.org/>].
- If you are new to MySQL Cluster, you may find our Developer Zone article [How to set up a MySQL Cluster for two servers](http://dev.mysql.com/tech-resources/articles/mysql-cluster-for-two-servers.html) [<http://dev.mysql.com/tech-resources/articles/mysql-cluster-for-two-servers.html>] to be helpful.

## 15.1. MySQL Cluster Overview

*MySQL Cluster* is a technology that enables clustering of in-memory databases in a shared-nothing system. The shared-nothing architecture allows the system to work with very inexpensive hardware, and without any specific requirements on hardware or software. It also does not have any single point of failure because each component has its own memory and disk.

MySQL Cluster integrates the standard MySQL server with an in-memory clustered storage engine called **NDB**. In our documentation, the term **NDB** refers to the part of the setup that is specific to the storage engine, whereas “MySQL Cluster” refers to the combination of MySQL and the **NDB** storage engine.

A MySQL Cluster consists of a set of computers, each running a one or more processes which may include a MySQL server, a data node, a management server, and (possibly) a specialized data access programs. The relationship of these components in a cluster is shown here:



All these programs work together to form a MySQL Cluster. When data is stored in the **NDB Cluster** storage engine, the tables are stored in the data nodes. Such tables are directly accessible from all other MySQL servers in the cluster. Thus, in a payroll application storing data in a cluster, if one application updates the salary of an employee, all other MySQL servers that query this data can see this change immediately.

The data stored in the data nodes for MySQL Cluster can be mirrored; the cluster can handle failures of individual data nodes with no other impact than that a small number of transactions are aborted due to

losing the transaction state. Because transactional applications are expected to handle transaction failure, this should not be a source of problems.

## 15.2. Basic MySQL Cluster Concepts

*NDB* is an in-memory storage engine offering high-availability and data-persistence features.

The *NDB* storage engine can be configured with a range of failover and load-balancing options, but it is easiest to start with the storage engine at the cluster level. MySQL Cluster's *NDB* storage engine contains a complete set of data, dependent only on other data within the cluster itself.

The cluster portion of MySQL Cluster is currently configured independently of the MySQL servers. In a MySQL Cluster, each part of the cluster is considered to be a *node*.

**Note:** In many contexts, the term “node” is used to indicate a computer, but when discussing MySQL Cluster it means a *process*. It is possible to run any number of nodes on a single computer, for which we use the term *cluster host*.

(However, it should be noted MySQL does not currently support the use of multiple data nodes on a single computer in a production setting. See [Issues exclusive to MySQL Cluster](#).)

There are three types of cluster nodes, and in a minimal MySQL Cluster configuration, there will be at least three nodes, one of each of these types:

- *Management node* (MGM node): The role of this type of node is to manage the other nodes within the MySQL Cluster, performing such functions as providing configuration data, starting and stopping nodes, running backup, and so forth. Because this node type manages the configuration of the other nodes, a node of this type should be started first, before any other node. An MGM node is started with the command `ndb_mgmd`.
- *Data node*: This type of node stores cluster data. There are as many data nodes as there are replicas, times the number of fragments. For example, with two replicas, each having two fragments, you will need four data nodes. It is not necessary to have more than one replica. A data node is started with the command `ndbd`.
- *SQL node*: This is a node that accesses the cluster data. In the case of MySQL Cluster, an SQL node is a traditional MySQL server that uses the *NDB Cluster* storage engine. An SQL node is typically started with the command `mysqld --ndbcluster` or by using `mysqld` with the `ndb-cluster` option added to `my.cnf`.

An SQL node is actually just a specialised type of *API node*, which designates any application which accesses Cluster data. One example of an API node is the `ndb_restore` utility that is used to restore a cluster backup. It is possible to write such applications using the [NDB API](#) [<http://dev.mysql.com/doc/en/ndbapi/>].

**Important:** It is not realistic to expect to employ a three-node setup in a production environment. Such a configuration provides no redundancy; in order to benefit from MySQL Cluster's high-availability features, you must use multiple data and SQL nodes. The use of multiple management nodes is also highly recommended.

For a brief introduction to the relationships between nodes, node groups, replicas, and partitions in MySQL Cluster, see [Section 15.2.1, “MySQL Cluster Nodes, Node Groups, Replicas, and Partitions”](#).

Configuration of a cluster involves configuring each individual node in the cluster and setting up individual communication links between nodes. MySQL Cluster is currently designed with the intention that data nodes are homogeneous in terms of processor power, memory space, and bandwidth. In addition, to



provide a single point of configuration, all configuration data for the cluster as a whole is located in one configuration file.

The management server (MGM node) manages the cluster configuration file and the cluster log. Each node in the cluster retrieves the configuration data from the management server, and so requires a way to determine where the management server resides. When interesting events occur in the data nodes, the nodes transfer information about these events to the management server, which then writes the information to the cluster log.

In addition, there can be any number of cluster client processes or applications. These are of two types:

- **Standard MySQL clients:** These are no different for MySQL Cluster than they are for standard (non-Cluster) MySQL. In other words, MySQL Cluster can be accessed from existing MySQL applications written in PHP, Perl, C, C++, Java, Python, Ruby, and so on.
- **Management clients:** These clients connect to the management server and provide commands for starting and stopping nodes gracefully, starting and stopping message tracing (debug versions only), showing node versions and status, starting and stopping backups, and so on.

## 15.2.1. MySQL Cluster Nodes, Node Groups, Replicas, and Partitions

This section discusses the manner in which MySQL Cluster divides and duplicates data for storage.

Central to an understanding of this topic are the following concepts, listed here with brief definitions:

- **(Data) Node:** An `ndbd` process, which stores a *replica*—that is, a copy of the *partition* (see below) assigned to the node group of which the node is a member.

Each data node should be located on a separate computer. While it is also possible to host multiple `ndbd` processes on a single computer, such a configuration is not supported.

It is common for the terms “node” and “data node” to be used interchangeably when referring to an `ndbd` process; where mentioned, management (MGM) nodes (`ndb_mgmd` processes) and SQL nodes (`mysqld` processes) are specified as such in this discussion.

- **Node Group:** A node group consists of one or more nodes, and stores partitions, or sets of *replicas* (see next item).

**Note:** Currently, all node groups in a cluster must have the same number of nodes.

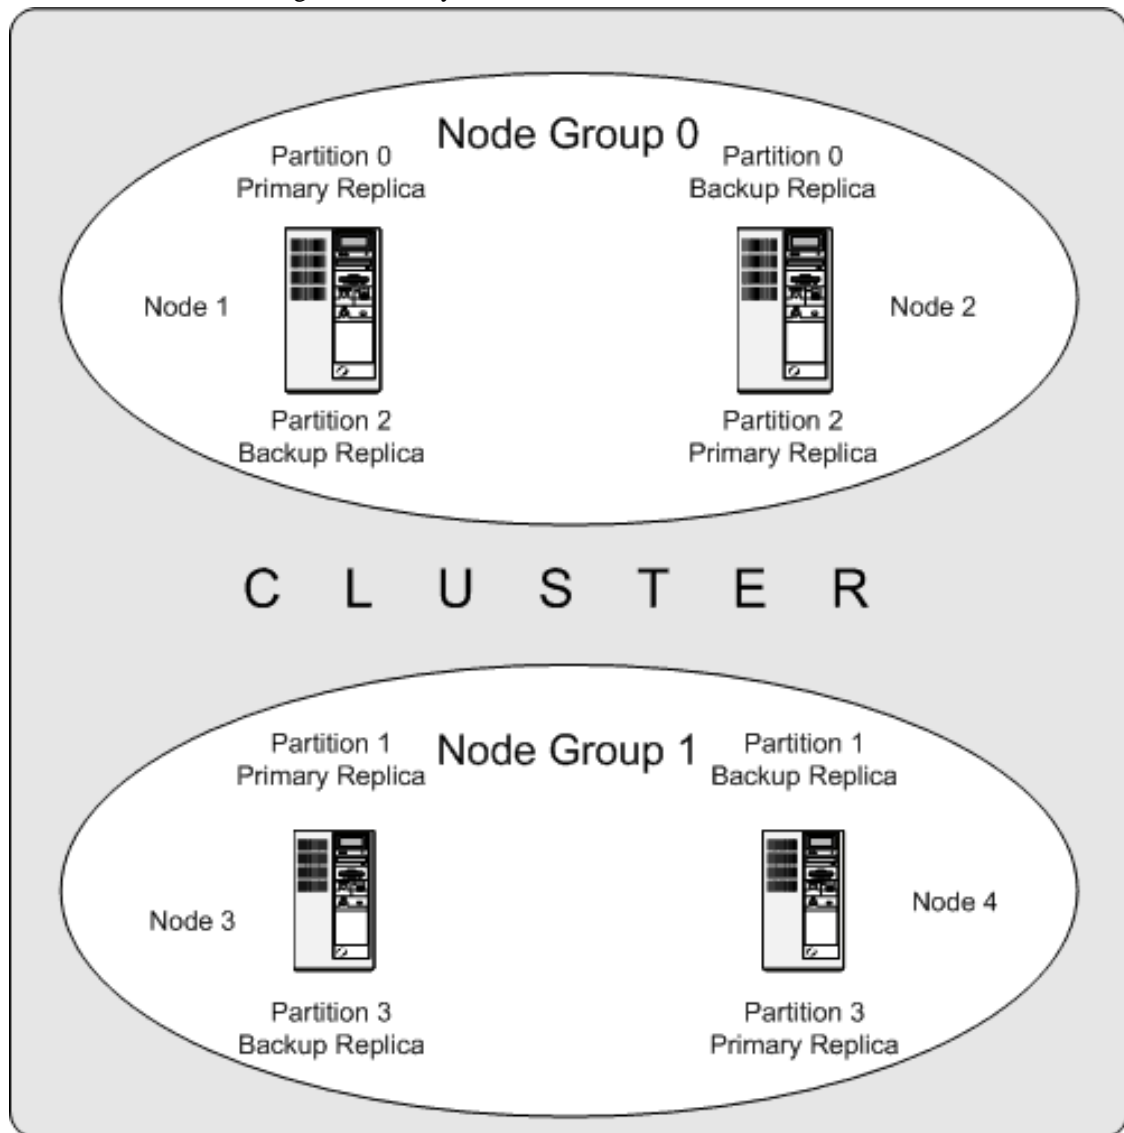
- **Partition:** This is a portion of the data stored by the cluster. There are as many cluster partitions as nodes participating in the cluster. Each node is responsible for keeping at least one copy of any partitions assigned to it (that is, at least one replica) available to the cluster.

A replica belongs entirely to a single node; a node can (and usually does) store several replicas.

- **Replica:** This is a copy of a cluster partition. Each node in a node group stores a replica. Also sometimes known as a *partition replica*. The number of replicas is equal to the number of nodes per node group.

The following diagram illustrates a MySQL Cluster with four data nodes, arranged in two node groups of two nodes each; nodes 1 and 2 belong to node group 0, and nodes 3 and 4 belong to node group 1. Note that only data (`ndbd`) nodes are shown here; although a working cluster requires an `ndb_mgm`


process for cluster management and at least one SQL node to access the data stored by the cluster, these have been omitted in the figure for clarity.



The data stored by the cluster is divided into four partitions, numbered 0, 1, 2, and 3. Each partition is stored — in multiple copies — on the same node group. Partitions are stored on alternate node groups:

- Partition 0 is stored on node group 0; a *primary replica* (primary copy) is stored on node 1, and a *backup replica* (backup copy of the partition) is stored on node 2.
- Partition 1 is stored on the other node group (node group 1); this partition's primary replica is on node 3, and its backup replica is on node 4.
- Partition 2 is stored on node group 0. However, the placing of its two replicas is reversed from that of Partition 0; for Partition 2, the primary replica is stored on node 2, and the backup on node 1.
- Partition 3 is stored on node group 1, and the placement of its two replicas are reversed from those of partition 1. That is, its primary replica is located on node 4, with the backup on node 3.

What this means regarding the continued operation of a MySQL Cluster is this: so long as each node group participating in the cluster has at least one node operating, the cluster has a complete copy of all data and remains viable. This is illustrated in the next diagram.



In this example, where the cluster consists of two node groups of two nodes each, any combination of at least one node in node group 0 and at least one node in node group 1 is sufficient to keep the cluster “alive” (indicated by arrows in the diagram). However, if *both* nodes from *either* node group fail, the remaining two nodes are not sufficient (shown by the arrows marked out with an **X**); in either case, the cluster has lost an entire partition and so can no longer provide access to a complete set of all cluster data.

## 15.3. Simple Multi-Computer How-To

This section is a “How-To” that describes the basics for how to plan, install, configure, and run a MySQL Cluster. Whereas the examples in [Section 15.4, “MySQL Cluster Configuration”](#) provide more in-depth information on a variety of clustering options and configuration, the result of following the guidelines and procedures outlined here should be a usable MySQL Cluster which meets the *minimum* requirements for availability and safeguarding of data.

This section covers hardware and software requirements; networking issues; installation of MySQL Cluster; configuration issues; starting, stopping, and restarting the cluster; loading of a sample database; and performing queries.


### Basic Assumptions

This How-To makes the following assumptions:

- The cluster setup has four nodes, each on a separate host, and each with a fixed network address on a typical Ethernet as shown here:

Node	IP Address
Management (MGM) node	192.168.0.10
MySQL server (SQL) node	192.168.0.20
Data (NDBD) node "A"	192.168.0.30
Data (NDBD) node "B"	192.168.0.40

This may be made clearer in the following diagram:



**Note:** In the interest of simplicity (and reliability), this How-To uses only numeric IP addresses. However, if DNS resolution is available on your network, it is possible to use hostnames in lieu of IP addresses in configuring Cluster. Alternatively, you can use the `/etc/hosts` file or your operating system's equivalent for providing a means to do host lookup if such is available.

- Each host in our scenario is an Intel-based desktop PC running a common, generic Linux distribution installed to disk in a standard configuration, and running no unnecessary services. The core OS with standard TCP/IP networking capabilities should be sufficient. Also for the sake of simplicity, we also assume that the filesystems on all hosts are set up identically. In the event that they are not, you will need to adapt these instructions accordingly.
- Standard 100 Mbps or 1 gigabit Ethernet cards are installed on each machine, along with the proper drivers for the cards, and that all four hosts are connected via a standard-issue Ethernet networking appliance such as a switch. (All machines should use network cards with the same throughput. That is, all four machines in the cluster should have 100 Mbps cards *or* all four machines should have 1

Gbps cards.) MySQL Cluster will work in a 100 Mbps network; however, gigabit Ethernet will provide better performance.

Note that MySQL Cluster is *not* intended for use in a network for which throughput is less than 100 Mbps. For this reason (among others), attempting to run a MySQL Cluster over a public network such as the Internet is not likely to be successful, and is not recommended.

4. For our sample data, we will use the `world` database which is available for download from the MySQL AB Web site. As this database takes up a relatively small amount of space, we assume that each machine has 256MB RAM, which should be sufficient for running the operating system, host NDB process, and (for the data nodes) for storing the database.

Although we refer to a Linux operating system in this How-To, the instructions and procedures that we provide here should be easily adaptable to other supported operating systems. We also assume that you already know how to perform a minimal installation and configuration of the operating system with networking capability, or that you are able to obtain assistance in this elsewhere if needed.

We discuss MySQL Cluster hardware, software, and networking requirements in somewhat greater detail in the next section. (See [Section 15.3.1, “Hardware, Software, and Networking”](#).)

### 15.3.1. Hardware, Software, and Networking

One of the strengths of MySQL Cluster is that it can be run on commodity hardware and has no unusual requirements in this regard, other than for large amounts of RAM, due to the fact that all live data storage is done in memory. (Note that this is subject to change, and that we intend to implement disk-based storage in a future MySQL Cluster release.) Naturally, multiple and faster CPUs will enhance performance. Memory requirements for Cluster processes are relatively small.

The software requirements for Cluster are also modest. Host operating systems do not require any unusual modules, services, applications, or configuration to support MySQL Cluster. For supported operating systems, a standard installation should be sufficient. The MySQL software requirements are simple: all that is needed is a production release of MySQL-max 4.1.3 or newer; you must use the `-max` version of MySQL to have Cluster support. (See [Section 5.3, “The `mysqld-max` Extended MySQL Server”](#).) It is not necessary to compile MySQL yourself merely to be able to use Cluster. In this How-To, we assume that you are using the `-max` binary appropriate to your operating system, available via the MySQL software downloads page at <http://dev.mysql.com/downloads/>.

For inter-node communication, Cluster supports TCP/IP networking in any standard topology, and the minimum expected for each host is a standard 100 Mbps Ethernet card, plus a switch, hub, or router to provide network connectivity for the cluster as a whole. We strongly recommend that a MySQL Cluster be run on its own subnet which is not shared with non-Cluster machines for the following reasons:

- **Security:** Communications between Cluster nodes are not encrypted or shielded in any way. The only means of protecting transmissions within a MySQL Cluster is to run your Cluster on a protected network. If you intend to use MySQL Cluster for Web applications, the cluster should definitely reside behind your firewall and not in your network's De-Militarized Zone (DMZ [[http://compnetworking.about.com/cs/networksecurity/g/bldef\\_dmz.htm](http://compnetworking.about.com/cs/networksecurity/g/bldef_dmz.htm)]) or elsewhere.
- **Efficiency:** Setting up a MySQL Cluster on a private or protected network allows the cluster to make exclusive use of bandwidth between cluster hosts. Using a separate switch for your MySQL Cluster not only helps protect against unauthorized access to Cluster data, it also ensures that Cluster nodes are shielded from interference caused by transmissions between other computers on the network. For enhanced reliability, you can use dual switches and dual cards to remove the network as a single point of failure; many device drivers support failover for such communication links.

It is also possible to use the high-speed Scalable Coherent Interface (SCI) with MySQL Cluster, but this is not a requirement. See [Section 15.10, “Using High-Speed Interconnects with MySQL Cluster”](#), for more about this protocol and its use with MySQL Cluster.

## 15.3.2. Multi-Computer Installation

Each MySQL Cluster host computer running data or SQL nodes must have installed on it a MySQL-max binary. For management nodes, it is not necessary to install the MySQL server binary, but you do have to install the MGM server daemon and client binaries (`ndb_mgmd` and `ndb_mgm`, respectively). This section covers the steps necessary to install the correct binaries for each type of Cluster node.

MySQL AB provides precompiled binaries that support Cluster, and there is generally no need to compile these yourself. Therefore, the first step in the installation process for each cluster host is to download the file `mysql-max-4.1.22-pc-linux-gnu-i686.tar.gz` from the [MySQL downloads area](#) [<http://dev.mysql.com/downloads/>]. We assume that you have placed it in each machine's `/var/tmp` directory. (If you do require a custom binary, see [Section 2.9.3, “Installing from the Development Source Tree”](#).)

RPMs are also available for both 32-bit and 64-bit Linux platforms; as of MySQL 4.1.10a, the `-max` binaries installed by the RPMs support the `NDBCluster` storage engine. If you choose to use these rather than the binary files, be aware that you must install *both* the `-server` and `-max` packages on all machines that are to host cluster nodes. (See [Section 2.4, “Installing MySQL on Linux”](#), for more information about installing MySQL using the RPMs.) After installing from RPM, you will still need to configure the cluster as discussed in [Section 15.3.3, “Multi-Computer Configuration”](#).

**Note:** After completing the installation, do not yet start any of the binaries. We will show you how to do so following the configuration of all nodes.

### Storage and SQL Node Installation

On each of the three machines designated to host storage or SQL nodes, perform the following steps as the system `root` user:

1. Check your `/etc/passwd` and `/etc/group` files (or use whatever tools are provided by your operating system for managing users and groups) to see whether there is already a `mysql` group and `mysql` user on the system. Some OS distributions create these as part of the operating system installation process. If they are not already present, create a new `mysql` user group, and then add a `mysql` user to this group:

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

The syntax for `useradd` and `groupadd` may differ slightly on different versions of Unix, or they may have different names such as `adduser` and `addgroup`.

2. Change location to the directory containing the downloaded file, unpack the archive, and create a symlink to the `mysql-max` directory named `mysql`. Note that the actual file and directory names will vary according to the MySQL version number.

```
shell> cd /var/tmp
shell> tar -xzvf -C /usr/local mysql-max-4.1.22-pc-linux-gnu-i686.tar.gz
shell> ln -s /usr/local/mysql-max-4.1.22-pc-linux-gnu-i686 /usr/local/mysql
```

3. Change location to the `mysql` directory and run the supplied script for creating the system databases:

```
shell> cd mysql
shell> scripts/mysql_install_db --user=mysql
```

4. Set the necessary permissions for the MySQL server and data directories:

```
shell> chown -R root .
shell> chown -R mysql data
shell> chgrp -R mysql .
```

Note that the data directory on each machine hosting a data node is `/usr/local/mysql/data`. We will use this piece of information when we configure the management node. (See [Section 15.3.3, “Multi-Computer Configuration”](#).)

5. Copy the MySQL startup script to the appropriate directory, make it executable, and set it to start when the operating system is booted up:

```
shell> cp support-files/mysql.server /etc/rc.d/init.d/
shell> chmod +x /etc/rc.d/init.d/mysql.server
shell> chkconfig --add mysql.server
```

(The startup scripts directory may vary depending on your operating system and version — for example, in some Linux distributions, it is `/etc/init.d`.)

Here we use Red Hat's `chkconfig` for creating links to the startup scripts; use whatever means is appropriate for this purpose on your operating system and distribution, such as `update-rc.d` on Debian.

Remember that the preceding steps must be performed separately for each machine on which a storage or SQL node is to reside.

### Management Node Installation

Installation for the management (MGM) node does not require installation of the `mysqld` binary. Only the binaries for the MGM server and client are required, which can be found in the downloaded archive. Again, we assume that you have placed this file in `/var/tmp`.

As system `root` (that is, after using `sudo`, `su root`, or your system's equivalent for temporarily assuming the system administrator account's privileges), perform the following steps to install `ndb_mgmd` and `ndb_mgm` on the Cluster management node host:

1. Change location to the `/var/tmp` directory, and extract the `ndb_mgm` and `ndb_mgmd` from the archive into a suitable directory such as `/usr/local/bin`:

```
shell> cd /var/tmp
shell> tar -zxvf mysql-4.1.22-pc-linux-gnu-i686.tar.gz
shell> cd mysql-4.1.22-pc-linux-gnu-i686
shell> cp /bin/ndb_mgm* /usr/local/bin
```

(You can safely delete the directory created by unpacking the downloaded archive, and the files it contains, from `/var/tmp` once `ndb_mgm` and `ndb_mgmd` have been copied to the executables directory.)

2. Change location to the directory into which you copied the files, and then make both of them executable:

```
shell> cd /usr/local/bin
shell> chmod +x ndb_mgm*
```

In [Section 15.3.3, “Multi-Computer Configuration”](#), we will create and write configuration files for all of the nodes in our example Cluster.

### 15.3.3. Multi-Computer Configuration

For our four-node, four-host MySQL Cluster, we will need to write four configuration files, one per node/host.

- Each data node or SQL node requires a `my.cnf` file that provides two pieces of information: a **connectstring** telling the node where to find the MGM node, and a line telling the MySQL server on this host (the machine hosting the data node) to run in NDB mode.

For more information on connectstrings, see [Section 15.4.4.2, “The Cluster connectstring”](#).

- The management node needs a `config.ini` file telling it how many replicas to maintain, how much memory to allocate for data and indexes on each data node, where to find the data nodes, where to save data to disk on each data node, and where to find any SQL nodes.

#### Configuring the Storage and SQL Nodes

The `my.cnf` file needed for the data nodes is fairly simple. The configuration file should be located in the `/etc` directory and can be edited using any text editor. (Create the file if it does not exist.) For example:

```
shell> vi /etc/my.cnf
```

We show `vi` being used here to create the file, but any text editor should work just as well.

For each data node and SQL node in our example setup, `my.cnf` should look like this:

```
# Options for mysqld process:
[MYSQLD]
ndbcluster                # run NDB engine
ndb-connectstring=192.168.0.10 # location of MGM node

# Options for ndbd process:
[MYSQL_CLUSTER]
ndb-connectstring=192.168.0.10 # location of MGM node
```

After entering the preceding information, save this file and exit the text editor. Do this for the machines hosting data node “A”, data node “B”, and the SQL node.

**Important:** Once you have started a `mysqld` process with the `ndbcluster` and `ndb-connectstring` parameters in the `[MYSQLD]` in the `my.cnf` file as shown previously, you cannot execute any `CREATE TABLE` or `ALTER TABLE` statements without having actually started the cluster. Otherwise, these statements will fail with an error. *This is by design.*

#### Configuring the Management Node

The first step in configuring the MGM node is to create the directory in which the configuration file can be found and then to create the file itself. For example (running as `root`):

```
shell> mkdir /var/lib/mysql-cluster
shell> cd /var/lib/mysql-cluster
shell> vi config.ini
```

For our representative setup, the `config.ini` file should read as follows:



```

# Options affecting ndbd processes on all data nodes:
[NDBD DEFAULT]
NoOfReplicas=2      # Number of replicas
DataMemory=80M     # How much memory to allocate for data storage
IndexMemory=18M    # How much memory to allocate for index storage
                   # For DataMemory and IndexMemory, we have used the
                   # default values. Since the "world" database takes up
                   # only about 500KB, this should be more than enough for
                   # this example Cluster setup.

# TCP/IP options:
[TCP DEFAULT]
portnumber=2202    # This the default; however, you can use any
                   # port that is free for all the hosts in cluster
                   # Note: It is recommended beginning with MySQL 5.0 that
                   # you do not specify the portnumber at all and simply allow
                   # the default value to be used instead

# Management process options:
[NDB_MGMD]
hostname=192.168.0.10 # Hostname or IP address of MGM node
datadir=/var/lib/mysql-cluster # Directory for MGM node logfiles

# Options for data node "A":
[NDBD]
hostname=192.168.0.30 # (one [NDBD] section per data node)
                       # Hostname or IP address
datadir=/usr/local/mysql/data # Directory for this data node's datafiles

# Options for data node "B":
[NDBD]
hostname=192.168.0.40 # Hostname or IP address
datadir=/usr/local/mysql/data # Directory for this data node's datafiles

# SQL node options:
[MYSQLD]
hostname=192.168.0.20 # Hostname or IP address
                       # (additional mysqld connections can be
                       # specified for this node for various
                       # purposes such as running ndb_restore)

```

(**Note:** The `world` database can be downloaded from <http://dev.mysql.com/doc/>, where it can be found listed under “Examples.”)

After all the configuration files have been created and these minimal options have been specified, you are ready to proceed with starting the cluster and verifying that all processes are running. We discuss how this is done in [Section 15.3.4, “Initial Startup”](#).

For more detailed information about the available MySQL Cluster configuration parameters and their uses, see [Section 15.4.4, “Configuration File”](#), and [Section 15.4, “MySQL Cluster Configuration”](#). For configuration of MySQL Cluster as relates to making backups, see [Section 15.8.4, “Configuration for Cluster Backup”](#).

**Note:** The default port for Cluster management nodes is 1186; the default port for data nodes is 2202. In MySQL 4.1, ports for data nodes are allocated sequentially beginning with port 2202 and these ports must be available for the cluster to use.

## 15.3.4. Initial Startup

Starting the cluster is not very difficult after it has been configured. Each cluster node process must be started separately, and on the host where it resides. Although it is possible to start the nodes in any order, it is recommended that the management node be started first, followed by the storage nodes, and then finally by any SQL nodes:

1. On the management host, issue the following command from the system shell to start the MGM node process:

```
shell> ndb_mgmd -f /var/lib/mysql-cluster/config.ini
```

Note that `ndb_mgmd` must be told where to find its configuration file, using the `-f` or `-config-file` option. (See [Section 15.6.3, “ndb\\_mgmd, the Management Server Process”](#), for details.)

2. On each of the data node hosts, run this command to start the `ndbd` process for the first time:

```
shell> ndbd --initial
```

Note that it is very important to use the `--initial` parameter *only* when starting `ndbd` for the first time, or when restarting after a backup/restore operation or a configuration change. This is because the `--initial` option causes the node to delete any files created by earlier `ndbd` instances that are needed for recovery, including the recovery log files.

3. If you used RPM files to install MySQL on the cluster host where the SQL node is to reside, you can (and should) use the startup script installed in `/etc/init.d` to start the MySQL server process on the SQL node. Note that you need to install the `-max` server RPM *in addition to* the Standard server RPM to run the `-max` server binary.

If all has gone well, and the cluster has been set up correctly, the cluster should now be operational. You can test this by invoking the `ndb_mgm` management node client. The output should look like that shown here, although you might see some slight differences in the output depending upon the exact version of MySQL that you are using:

```
shell> ndb_mgm
-- NDB Cluster -- Management Client --
ndb_mgm> SHOW
Connected to Management Server at: localhost:1186
Cluster Configuration
-----
[ndbd(NDB)]      2 node(s)
id=2      @192.168.0.30  (Version: 4.1.22, Nodegroup: 0, Master)
id=3      @192.168.0.40  (Version: 4.1.22, Nodegroup: 0)

[ndb_mgmd(MGM)] 1 node(s)
id=1      @192.168.0.10  (Version: 4.1.22)

[mysqld(SQL)]   1 node(s)
id=4      (Version: 4.1.22)
```

**Note:** The SQL node is referenced here as `[mysqld(API)]`. This is perfectly normal, and reflects the fact that the `mysqld` process is acting as a cluster API node.

You should now be ready to work with databases, tables, and data in MySQL Cluster. See [Section 15.3.5, “Loading Sample Data and Performing Queries”](#), for a brief discussion.

## 15.3.5. Loading Sample Data and Performing Queries

Working with data in MySQL Cluster is not much different from doing so in MySQL without Cluster. There are two points to keep in mind:

- For a table to be replicated in the cluster, it must use the `NDB Cluster` storage engine. To specify this, use the `ENGINE=NDB` or `ENGINE=NDBCLUSTER` table option. You can add this option when creating the table:

```
CREATE TABLE tbl_name ( ... ) ENGINE=NDBCLUSTER;
```

Alternatively, for an existing table that uses a different storage engine, use `ALTER TABLE` to

change the table to use **NDB Cluster**:

```
ALTER TABLE tbl_name ENGINE=NDBCLUSTER;
```

- Each **NDB** table *must* have a primary key. If no primary key is defined by the user when a table is created, the **NDB Cluster** storage engine automatically generates a hidden one. (**Note:** This hidden key takes up space just as does any other table index. It is not uncommon to encounter problems due to insufficient memory for accommodating these automatically created indexes.)

If you are importing tables from an existing database using the output of `mysqldump`, you can open the SQL script in a text editor and add the **ENGINE** option to any table creation statements, or replace any existing **ENGINE** (or **TYPE**) options. Suppose that you have the `world` sample database on another MySQL server that does not support MySQL Cluster, and you want to export the `City` table:

```
shell> mysqldump --add-drop-table world City > city_table.sql
```

The resulting `city_table.sql` file will contain this table creation statement (and the **INSERT** statements necessary to import the table data):

```
DROP TABLE IF EXISTS `City`;
CREATE TABLE `City` (
  `ID` int(11) NOT NULL auto_increment,
  `Name` char(35) NOT NULL default '',
  `CountryCode` char(3) NOT NULL default '',
  `District` char(20) NOT NULL default '',
  `Population` int(11) NOT NULL default '0',
  PRIMARY KEY (`ID`)
) ENGINE=MyISAM DEFAULT CHARSET=latin1;

INSERT INTO `City` VALUES (1,'Kabul','AFG','Kabol',1780000);
INSERT INTO `City` VALUES (2,'Qandahar','AFG','Qandahar',237500);
INSERT INTO `City` VALUES (3,'Herat','AFG','Herat',186800);
(remaining INSERT statements omitted)
```

You will need to make sure that MySQL uses the **NDB** storage engine for this table. There are two ways that this can be accomplished. One of these is to modify the table definition *before* importing it into the Cluster database. Using the `City` table as an example, modify the **ENGINE** option of the definition as follows:

```
DROP TABLE IF EXISTS `City`;
CREATE TABLE `City` (
  `ID` int(11) NOT NULL auto_increment,
  `Name` char(35) NOT NULL default '',
  `CountryCode` char(3) NOT NULL default '',
  `District` char(20) NOT NULL default '',
  `Population` int(11) NOT NULL default '0',
  PRIMARY KEY (`ID`)
) ENGINE=NDBCLUSTER DEFAULT CHARSET=latin1;

INSERT INTO `City` VALUES (1,'Kabul','AFG','Kabol',1780000);
INSERT INTO `City` VALUES (2,'Qandahar','AFG','Qandahar',237500);
INSERT INTO `City` VALUES (3,'Herat','AFG','Herat',186800);
(remaining INSERT statements omitted)
```

This must be done for the definition of each table that is to be part of the clustered database. The easiest way to accomplish this is to do a search-and-replace on the file that contains the definitions and replace all instances of **TYPE=engine\_name** or **ENGINE=engine\_name** with **ENGINE=NDBCLUSTER**. If you do not want to modify the file, you can use the unmodified file to create the tables, and then use **ALTER TABLE** to change their storage engine. The particulars are given later in this section.

Assuming that you have already created a database named `world` on the SQL node of the cluster, you can then use the `mysql` command-line client to read `city_table.sql`, and create and populate the corresponding table in the usual manner:

```
shell> mysql world < city_table.sql
```

It is very important to keep in mind that the preceding command must be executed on the host where the SQL node is running (in this case, on the machine with the IP address 192.168.0.20).

To create a copy of the entire `world` database on the SQL node, use `mysqldump` on the non-cluster server to export the database to a file named `world.sql`; for example, in the `/tmp` directory. Then modify the table definitions as just described and import the file into the SQL node of the cluster like this:

```
shell> mysql world < /tmp/world.sql
```

If you save the file to a different location, adjust the preceding instructions accordingly.

It is important to note that `NDB Cluster` in MySQL 4.1 does not support autodiscovery of databases. (See [Section 15.11, “Known Limitations of MySQL Cluster”](#).) This means that, once the `world` database and its tables have been created on one data node, you need to issue the `CREATE DATABASE world` statement followed by `FLUSH TABLES` on each SQL node in the cluster. This causes the node to recognize the database and read its table definitions.

Running `SELECT` queries on the SQL node is no different from running them on any other instance of a MySQL server. To run queries from the command line, you first need to log in to the MySQL Monitor in the usual way (specify the `root` password at the `Enter password:` prompt):

```
shell> mysql -u root -p
Enter password:
Welcome to the MySQL monitor.  Commands end with ; or \g.
Your MySQL connection id is 1 to server version: 4.1.22

Type 'help;' or '\h' for help. Type '\c' to clear the buffer.

mysql>
```

We simply use the MySQL server's `root` account and assume that you have followed the standard security precautions for installing a MySQL server, including setting a strong `root` password. For more information, see [Section 2.10.3, “Securing the Initial MySQL Accounts”](#).

It is worth taking into account that Cluster nodes do not make use of the MySQL privilege system when accessing one another. Setting or changing MySQL user accounts (including the `root` account) effects only applications that access the SQL node, not interaction between nodes.

If you did not modify the `ENGINE` clauses in the table definitions prior to importing the SQL script, you should run the following statements at this point:

```
mysql> USE world;
mysql> ALTER TABLE City ENGINE=NDBCLUSTER;
mysql> ALTER TABLE Country ENGINE=NDBCLUSTER;
mysql> ALTER TABLE CountryLanguage ENGINE=NDBCLUSTER;
```

Selecting a database and running a `SELECT` query against a table in that database is also accomplished in the usual manner, as is exiting the MySQL Monitor:

```
mysql> USE world;
mysql> SELECT Name, Population FROM City ORDER BY Population DESC LIMIT 5;
+-----+-----+
| Name      | Population |
+-----+-----+
| Bombay    | 10500000  |
| Seoul     | 9981619   |
| São Paulo | 9968485   |
| Shanghai  | 9696300   |
| Jakarta   | 9604900   |
+-----+-----+
```

```
5 rows in set (0.34 sec)
```

```
mysql> \q
Bye

shell>
```

Applications that use MySQL can employ standard APIs to access NDB tables. It is important to remember that your application must access the SQL node, and not the MGM or data nodes. This brief example shows how we might execute the `SELECT` statement just shown by using PHP 5's `mysqli` extension running on a Web server elsewhere on the network:

```
<!DOCTYPE HTML PUBLIC "-//W3C//DTD HTML 4.01 Transitional//EN"
"http://www.w3.org/TR/html4/loose.dtd">
<html>
<head>
<meta http-equiv="Content-Type"
content="text/html; charset=iso-8859-1">
<title>SIMPLE mysqli SELECT</title>
</head>
<body>
<?php
# connect to SQL node:
$link = new mysqli('192.168.0.20', 'root', 'root_password', 'world');
# parameters for mysqli constructor are:
# host, user, password, database

if( mysqli_connect_errno() )
die("Connect failed: " . mysqli_connect_error());

$query = "SELECT Name, Population
FROM City
ORDER BY Population DESC
LIMIT 5";

# if no errors...
if( $result = $link->query($query) )
{
?>
<table border="1" width="40%" cellpadding="4" cellspacing="1">
<tbody>
<tr>
<th width="10%">City</th>
<th>Population</th>
</tr>
<?
# then display the results...
while($row = $result->fetch_object())
printf("<tr>\n <td align=\"center\">%s</td><td>%d</td>\n</tr>\n",
$row->Name, $row->Population);
?>
</tbody>
</table>
<?
# ...and verify the number of rows that were retrieved
printf("<p>Affected rows: %d</p>\n", $link->affected_rows);
}
else
# otherwise, tell us what went wrong
echo mysqli_error();

# free the result set and the mysqli connection object
$result->close();
$link->close();
?>
</body>
</html>
```

We assume that the process running on the Web server can reach the IP address of the SQL node.

In a similar fashion, you can use the MySQL C API, Perl-DBI, Python-mysql, or MySQL AB's own Connectors to perform the tasks of data definition and manipulation just as you would normally with MySQL.

## 15.3.6. Safe Shutdown and Restart

To shut down the cluster, enter the following command in a shell on the machine hosting the MGM node:

```
shell> ndb_mgm -e shutdown
```

The `-e` option here is used to pass a command to the `ndb_mgm` client from the shell. See [Section 4.3.1, “Using Options on the Command Line”](#). The command causes the `ndb_mgm`, `ndb_mgmd`, and any `ndbd` processes to terminate gracefully. Any SQL nodes can be terminated using `mysqladmin shutdown` and other means.

To restart the cluster, run these commands:

- On the management host (192.168.0.10 in our example setup):

```
shell> ndb_mgmd -f /var/lib/mysql-cluster/config.ini
```

- On each of the data node hosts (192.168.0.30 and 192.168.0.40):

```
shell> ndbd
```

Remember *not* to invoke this command with the `--initial` option when restarting an NDBD node normally.

- On the SQL host (192.168.0.20):

```
shell> mysqld &
```

For information on making Cluster backups, see [Section 15.8.2, “Using The Management Client to Create a Backup”](#).

To restore the cluster from backup requires the use of the `ndb_restore` command. This is covered in [Section 15.8.3, “How to Restore a Cluster Backup”](#).

More information on configuring MySQL Cluster can be found in [Section 15.4, “MySQL Cluster Configuration”](#).

## 15.4. MySQL Cluster Configuration

A MySQL server that is part of a MySQL Cluster differs in only one respect from a normal (non-clustered) MySQL server, in that it employs the `NDB Cluster` storage engine. This engine is also referred to simply as `NDB`, and the two forms of the name are synonymous.

To avoid unnecessary allocation of resources, the server is configured by default with the `NDB` storage engine disabled. To enable `NDB`, you must modify the server's `my.cnf` configuration file, or start the server with the `--ndbcluster` option.

The MySQL server is a part of the cluster, so it also must know how to access an MGM node to obtain the cluster configuration data. The default behavior is to look for the MGM node on `localhost`. However, should you need to specify that its location is elsewhere, this can be done in `my.cnf` or on the MySQL server command line. Before the `NDB` storage engine can be used, at least one MGM node must be operational, as well as any desired data nodes.

## 15.4.1. Building MySQL Cluster from Source Code

**NDB**, the Cluster storage engine, is available in binary distributions for Linux, Mac OS X, and Solaris. We are working to make Cluster run on all operating systems supported by MySQL, including Windows.

If you choose to build from a source tarball or the MySQL 4.1 BitKeeper tree, be sure to use the `--with-ndbcluster` option when running `configure`. You can also use the `BUILD/compile-pentium-max` build script. Note that this script includes OpenSSL, so you must either have or obtain OpenSSL to build successfully, or else modify `compile-pentium-max` to exclude this requirement. Of course, you can also just follow the standard instructions for compiling your own binaries, and then perform the usual tests and installation procedure. See [Section 2.9.3, “Installing from the Development Source Tree”](#).

## 15.4.2. Installing the Software

In the next few sections, we assume that you are already familiar with installing MySQL, and here we cover only the differences between configuring MySQL Cluster and configuring MySQL without clustering. (See [Chapter 2, \*Installing and Upgrading MySQL\*](#), if you require more information about the latter.)

You will find Cluster configuration easiest if you have already have all management and data nodes running first; this is likely to be the most time-consuming part of the configuration. Editing the `my.cnf` file is fairly straightforward, and this section will cover only any differences from configuring MySQL without clustering.

## 15.4.3. Quick Test Setup of MySQL Cluster

To familiarize you with the basics, we will describe the simplest possible configuration for a functional MySQL Cluster. After this, you should be able to design your desired setup from the information provided in the other relevant sections of this chapter.

First, you need to create a configuration directory such as `/var/lib/mysql-cluster`, by executing the following command as the system `root` user:

```
shell> mkdir /var/lib/mysql-cluster
```

In this directory, create a file named `config.ini` that contains the following information. Substitute appropriate values for `HostName` and `DataDir` as necessary for your system.

```
# file "config.ini" - showing minimal setup consisting of 1 data node,
# 1 management server, and 3 MySQL servers.
# The empty default sections are not required, and are shown only for
# the sake of completeness.
# Data nodes must provide a hostname but MySQL Servers are not required
# to do so.
# If you don't know the hostname for your machine, use localhost.
# The DataDir parameter also has a default value, but it is recommended to
# set it explicitly.
# Note: DB, API, and MGM are aliases for NDBD, MYSQLD, and NDB_MGMD
# respectively. DB and API are deprecated and should not be used in new
# installations.
[NDBD DEFAULT]
NoOfReplicas= 1

[MYSQLD DEFAULT]
[NDB_MGMD DEFAULT]
[TCP DEFAULT]

[NDB_MGMD]
HostName= myhost.example.com

[NDBD]
HostName= myhost.example.com
```

```
DataDir= /var/lib/mysql-cluster
[MYSQLD]
[MYSQLD]
[MYSQLD]
```

You can now start the `ndb_mgmd` management server. By default, it attempts to read the `config.ini` file in its current working directory, so change location into the directory where the file is located and then invoke `ndb_mgmd`:

```
shell> cd /var/lib/mysql-cluster
shell> ndb_mgmd
```

Then start a single data node by running `ndbd`. When starting `ndbd` for a given data node for the very first time, you should use the `--initial` option as shown here:

```
shell> ndbd --initial
```

For subsequent `ndbd` starts, you will generally want to *omit* the `--initial` option:

```
shell> ndbd
```

The reason for omitting `--initial` on subsequent restarts is that this option causes `ndbd` to delete and re-create all existing data and log files (as well as all table metadata) for this data node. One exception to this rule about not using `--initial` except for the first `ndbd` invocation is that you use it when restarting the cluster and restoring from backup after adding new data nodes.

By default, `ndbd` looks for the management server at `localhost` on port 1186. (Prior to MySQL 4.1.8, the default port was 2200.)

**Note:** If you have installed MySQL from a binary tarball, you will need to specify the path of the `ndb_mgmd` and `ndbd` servers explicitly. (Normally, these will be found in `/usr/local/mysql/bin`.)

Finally, change location to the MySQL data directory (usually `/var/lib/mysql` or `/usr/local/mysql/data`), and make sure that the `my.cnf` file contains the option necessary to enable the NDB storage engine:

```
[mysqld]
ndbcluster
```

You can now start the MySQL server as usual:

```
shell> mysqld_safe --user=mysql &
```

Wait a moment to make sure the MySQL server is running properly. If you see the notice `mysql ended`, check the server's `.err` file to find out what went wrong.

If all has gone well so far, you now can start using the cluster. Connect to the server and verify that the `NDBCLUSTER` storage engine is enabled:

```
shell> mysql
Welcome to the MySQL monitor.  Commands end with ; or \g.
Your MySQL connection id is 1 to server version: 4.1.22-Max

Type 'help;' or '\h' for help. Type '\c' to clear the buffer.

mysql> SHOW ENGINES\G
...
***** 12. row *****
Engine: NDBCLUSTER
```



```
Support: YES
Comment: Clustered, fault-tolerant, memory-based tables
***** 13. row *****
Engine: NDB
Support: YES
Comment: Alias for NDBCLUSTER
...
```

The row numbers shown in the preceding example output may be different from those shown on your system, depending upon how your server is configured.

Try to create an `NDBCLUSTER` table:

```
shell> mysql
mysql> USE test;
Database changed

mysql> CREATE TABLE ctest (i INT) ENGINE=NDBCLUSTER;
Query OK, 0 rows affected (0.09 sec)

mysql> SHOW CREATE TABLE ctest \G
***** 1. row *****
      Table: ctest
Create Table: CREATE TABLE `ctest` (
  `i` int(11) default NULL
) ENGINE=ndbcluster DEFAULT CHARSET=latin1
1 row in set (0.00 sec)
```

To check that your nodes were set up properly, start the management client:

```
shell> ndb_mgm
```

Use the `SHOW` command from within the management client to obtain a report on the cluster's status:

```
NDB> SHOW
Cluster Configuration
-----
[nbdb(NDB)]      1 node(s)
id=2      @127.0.0.1 (Version: 3.5.3, Nodegroup: 0, Master)

[ndb_mgmd(MGM)] 1 node(s)
id=1      @127.0.0.1 (Version: 3.5.3)

[mysqld(API)]   3 node(s)
id=3      @127.0.0.1 (Version: 3.5.3)
id=4 (not connected, accepting connect from any host)
id=5 (not connected, accepting connect from any host)
```

At this point, you have successfully set up a working MySQL Cluster. You can now store data in the cluster by using any table created with `ENGINE=NDBCLUSTER` or its alias `ENGINE=NDB`.

## 15.4.4. Configuration File

Configuring MySQL Cluster requires working with two files:

- `my.cnf`: Specifies options for all MySQL Cluster executables. This file, with which you should be familiar with from previous work with MySQL, must be accessible by each executable running in the cluster.
- `config.ini`: This file is read only by the MySQL Cluster management server, which then distributes the information contained therein to all processes participating in the cluster. `config.ini` contains a description of each node involved in the cluster. This includes configuration parameters for data nodes and configuration parameters for connections between all nodes in the cluster. For a quick reference to the sections that can appear in this file, and what sorts of configuration parameters may be placed in each section, see [Sections of the config.ini File](#).

We are continuously making improvements in Cluster configuration and attempting to simplify this process. Although we strive to maintain backward compatibility, there may be times when introduce an incompatible change. In such cases we will try to let Cluster users know in advance if a change is not backward compatible. If you find such a change and we have not documented it, please report it in the MySQL bugs database using the instructions given in [Section 1.8, “How to Report Bugs or Problems”](#).

### 15.4.4.1. Basic Example Configuration

To support MySQL Cluster, you will need to update `my.cnf` as shown in the following example. Note that the options shown here should not be confused with those that are used in `config.ini` files. You may also specify these parameters on the command line when invoking the executables.

From version 4.1.8 some simplifications in `my.cnf` were made, including new sections for `ndb-cluster` executables. However, these should not be confused with those occurring in `config.ini` files. As always, you may specify these parameters when invoking those executables from the command line.

```
# my.cnf
# example additions to my.cnf for MySQL Cluster
# (valid from 4.1.8)

# enable ndbcluster storage engine, and provide connectstring for
# management server host (default port is 1186)
[mysqld]
ndbcluster
ndb-connectstring=ndb_mgmd.mysql.com

# provide connectstring for management server host (default port: 1186)
[ndbd]
connect-string=ndb_mgmd.mysql.com

# provide connectstring for management server host (default port: 1186)
[ndb_mgm]
connect-string=ndb_mgmd.mysql.com

# provide location of cluster configuration file
[ndb_mgmd]
config-file=/etc/config.ini
```

(For more information on connectstrings, see [Section 15.4.4.2, “The Cluster connectstring”](#).)

```
# my.cnf
# example additions to my.cnf for MySQL Cluster
# (will work on all versions)

# enable ndbcluster storage engine, and provide connectstring for management
# server host to the default port 1186
[mysqld]
ndbcluster
ndb-connectstring=ndb_mgmd.mysql.com:1186
```

**Important:** Once you have started a `mysqld` process with the `ndbcluster` and `ndb-connectstring` parameters in the `[MYSQLD]` in the `my.cnf` file as shown previously, you cannot execute any `CREATE TABLE` or `ALTER TABLE` statements without having actually started the cluster. Otherwise, these statements will fail with an error. *This is by design.*

Starting with MySQL 4.1.8, you may also use a separate `[mysql_cluster]` section in the cluster `my.cnf` file for settings to be read and used by all executables:

```
# cluster-specific settings
[mysql_cluster]
ndb-connectstring=ndb_mgmd.mysql.com:1186
```

For additional `NDB` variables that can be set in the `my.cnf` file, see [Section 5.2.2, “System Variables”](#).

The configuration file is named `config.ini` by default. It is read by `ndb_mgmd` at startup and can be placed anywhere. Its location and name are specified by using `--config-file=path_name` on the `ndb_mgmd` command line. If the configuration file is not specified, `ndb_mgmd` by default tries to read a file named `config.ini` located in the current working directory.

Currently, the configuration file is in INI format, which consists of sections preceded by section headings (surrounded by square brackets), followed by the appropriate parameter names and values. One deviation from the standard INI format is that the parameter name and value can be separated by a colon (':') as well as the equals sign ('='). Another deviation is that sections are not uniquely identified by section name. Instead, unique sections (such as two different nodes of the same type) are identified by a unique ID specified as a parameter within the section.

Default values are defined for most parameters, and can also be specified in `config.ini`. To create a default value section, simply add the word `DEFAULT` to the section name. For example, an `[NDBD]` section contains parameters that apply to a particular data node, whereas an `[NDBD DEFAULT]` section contains parameters that apply to all data nodes. Suppose that all data nodes should use the same data memory size. To configure them all, create an `[NDBD DEFAULT]` section that contains a `DataMemory` line to specify the data memory size.

At a minimum, the configuration file must define the computers and nodes involved in the cluster and on which computers these nodes are located. An example of a simple configuration file for a cluster consisting of one management server, two data nodes and two MySQL servers is shown here:

```
# file "config.ini" - 2 data nodes and 2 SQL nodes
# This file is placed in the startup directory of ndb_mgmd (the
# management server)
# The first MySQL Server can be started from any host. The second
# can be started only on the host mysqld_5.mysql.com

[NDBD DEFAULT]
NoOfReplicas= 2
DataDir= /var/lib/mysql-cluster

[NDB_MGMD]
Hostname= ndb_mgmd.mysql.com
DataDir= /var/lib/mysql-cluster

[NDBD]
HostName= ndbd_2.mysql.com

[NDBD]
HostName= ndbd_3.mysql.com

[MYSQLD]
[MYSQLD]
HostName= mysqld_5.mysql.com
```

Note that each node has its own section in the `config.ini`. For instance, this cluster has two data nodes, so the preceding configuration file contains two `[NDBD]` sections defining these nodes.

### Sections of the `config.ini` File

There are six different sections that you can use in the `config.ini` configuration file, as described in the following list:

- `[COMPUTER]`: Defines cluster hosts. This is not required to configure a viable MySQL Cluster, but be may used as a convenience when setting up a large cluster. See [Section 15.4.4.3, “Defining Cluster Computers”](#), for more information.
- `[NDBD]`: Defines a cluster data node (`ndbd` process). See [Section 15.4.4.5, “Defining Data Nodes”](#), for details.
- `[MYSQLD]`: Defines the cluster's MySQL server nodes (also called SQL or API nodes). For a dis-

discussion of SQL node configuration, see [Section 15.4.4.6, “Defining SQL and Other API Nodes”](#).

- `[MGM]` or `[NDB_MGMD]`: Defines a cluster management server (MGM) node. For information concerning the configuration of MGM nodes, see [Section 15.4.4.4, “Defining the Management Server”](#).
- `[TCP]`: Defines a TCP/IP connection between cluster nodes, with TCP/IP being the default connection protocol. Normally, `[TCP]` or `[TCP DEFAULT]` sections are not required to set up a MySQL Cluster, as the cluster handles this automatically; however, it may be necessary in some situations to override the defaults provided by the cluster. See [Section 15.4.4.7, “Cluster TCP/IP Connections”](#), for information about available TCP/IP configuration parameters and how to use them. (You may also find [Section 15.4.4.8, “TCP/IP Connections Using Direct Connections”](#) to be of interest in some cases.)
- `[SHM]`: Defines shared-memory connections between nodes. Prior to MySQL 4.1.9, this type of connection was available only in binaries that were built using the `--with-ndb-shm` option. Beginning with MySQL 4.1.9-max, it is enabled by default, but should still be considered experimental. For a discussion of SHM interconnects, see [Section 15.4.4.9, “Shared-Memory Connections”](#).
- `[SCI]`: Defines *Scalable Coherent Interface* connections between cluster data nodes. Such connections require software which, while freely available, is not part of the MySQL Cluster distribution, as well as specialised hardware. See [Section 15.4.4.10, “SCI Transport Connections”](#) for detailed information about SCI interconnects.

You can define `DEFAULT` values for each section. As of MySQL 4.1.5, all parameter names are case-insensitive, which differs from parameters specified in `my.cnf` or `my.ini` files.

### 15.4.4.2. The Cluster `connectstring`

With the exception of the MySQL Cluster management server (`ndb_mgmd`), each node that is part of a MySQL Cluster requires a *connectstring* that points to the management server's location. This connectstring is used in establishing a connection to the management server as well as in performing other tasks depending on the node's role in the cluster. The syntax for a connectstring is as follows:

```
<connectstring> :=
  [<nodeid-specification>,<host-specification>[,<host-specification>]
<nodeid-specification> := node_id
<host-specification> := host_name[:port_num]
```

`node_id` is an integer larger than 1 which identifies a node in `config.ini`. `host_name` is a string representing a valid Internet host name or IP address. `port_num` is an integer referring to a TCP/IP port number.

```
example 1 (long):  "nodeid=2,myhost1:1100,myhost2:1100,192.168.0.3:1200"
example 2 (short): "myhost1"
```

All nodes will use `localhost:1186` as the default connectstring value if none is provided. If `port_num` is omitted from the connectstring, the default port is 1186. (**Note:** Prior to MySQL 4.1.8, the default port was 2200.) This port should always be available on the network because it has been assigned by IANA for this purpose (see <http://www.iana.org/assignments/port-numbers> for details).

By listing multiple `<host-specification>` values, it is possible to designate several redundant management servers. A cluster node will attempt to contact successive management servers on each host in the order specified, until a successful connection has been established.

There are a number of different ways to specify the connectstring:

- Each executable has its own command-line option which enables specifying the management server at startup. (See the documentation for the respective executable.)
- Beginning with MySQL 4.1.8, it is also possible to set the connectstring for all nodes in the cluster at once by placing it in a `[mysql_cluster]` section in the management server's `my.cnf` file.
- For backward compatibility, two other options are available, using the same syntax:
  1. Set the `NDB_CONNECTSTRING` environment variable to contain the connectstring.
  2. Write the connectstring for each executable into a text file named `Ndb.cfg` and place this file in the executable's startup directory.

However, these are now deprecated and should not be used for new installations.

The recommended method for specifying the connectstring is to set it on the command line or in the `my.cnf` file for each executable.

### 15.4.4.3. Defining Cluster Computers

The `[COMPUTER]` section has no real significance other than serving as a way to avoid the need of defining host names for each node in the system. All parameters mentioned here are required.

- `Id`

This is an integer value, used to refer to the host computer elsewhere in the configuration file. This is not the same as the node ID.

- `HostName`

This is the computer's hostname or IP address.

### 15.4.4.4. Defining the Management Server

The `[NDB_MGMD]` section is used to configure the behavior of the management server. `[MGM]` can be used as an alias; the two section names are equivalent. All parameters in the following list are optional and assume their default values if omitted. **Note:** If neither the `ExecuteOnComputer` nor the `HostName` parameter is present, the default value `localhost` will be assumed for both.

- `Id`

Each node in the cluster has a unique identity, which is represented by an integer value in the range 1 to 63 inclusive. This ID is used by all internal cluster messages for addressing the node.

- `ExecuteOnComputer`

This refers to the `Id` set for one of the computers defined in a `[COMPUTER]` section of the `config.ini` file.

- `PortNumber`

This is the port number on which the management server listens for configuration requests and man-

agement commands.

- `HostName`

Specifying this parameter defines the hostname of the computer on which the management node is to reside. To specify a hostname other than `localhost`, either this parameter or `ExecuteOnComputer` is required.

- `LogDestination`

This parameter specifies where to send cluster logging information. There are three options in this regard: `CONSOLE`, `SYSLOG`, and `FILE`:

- `CONSOLE` outputs the log to `stdout`:

```
CONSOLE
```

- `SYSLOG` sends the log to a `syslog` facility, possible values being one of `auth`, `authpriv`, `cron`, `daemon`, `ftp`, `kern`, `lpr`, `mail`, `news`, `syslog`, `user`, `uucp`, `local0`, `local1`, `local2`, `local3`, `local4`, `local5`, `local6`, or `local7`.

**Note:** Not every facility is necessarily supported by every operating system.

```
SYSLOG:facility=syslog
```

- `FILE` pipes the cluster log output to a regular file on the same machine. The following values can be specified:
  - `filename`: The name of the logfile.
  - `maxsize`: The maximum size (in bytes) to which the file can grow before logging rolls over to a new file. When this occurs, the old logfile is renamed by appending `.N` to the filename, where `N` is the next number not yet used with this name.
  - `maxfiles`: The maximum number of logfiles.

```
FILE:filename=cluster.log,maxsize=1000000,maxfiles=6
```

It is possible to specify multiple log destinations separated by semicolons as shown here:

```
CONSOLE;SYSLOG:facility=local0;FILE:filename=/var/log/mgmd
```

The default value for the `FILE` parameter is

```
FILE:filename=ndb_node_id_cluster.log,maxsize=1000000,maxfiles=6,
```

where `node_id` is the ID of the node.

- `ArbitrationRank`

This parameter is used to define which nodes can act as arbitrators. Only MGM nodes and SQL nodes can be arbitrators. `ArbitrationRank` can take one of the following values:

- `0`: The node will never be used as an arbitrator.
- `1`: The node has high priority; that is, it will be preferred as an arbitrator over low-priority nodes.

- **2**: Indicates a low-priority node which be used as an arbitrator only if a node with a higher priority is not available for that purpose.

Normally, the management server should be configured as an arbitrator by setting its `ArbitrationRank` to 1 (the default value) and that of all SQL nodes to 0.

- `ArbitrationDelay`

An integer value which causes the management server's responses to arbitration requests to be delayed by that number of milliseconds. By default, this value is 0; it is normally not necessary to change it.

- `DataDir`

This specifies the directory where output files from the management server will be placed. These files include cluster log files, process output files, and the daemon's process ID (PID) file. (For log files, this location can be overridden by setting the `FILE` parameter for `LogDestination` as discussed previously in this section.)

### 15.4.4.5. Defining Data Nodes

The `[NDBD]` and `[NDBD DEFAULT]` sections are used to configure the behavior of the cluster's data nodes. There are many parameters which control buffer sizes, pool sizes, timeouts, and so forth. The only mandatory parameters are:

- Either `ExecuteOnComputer` or `HostName`, which must be defined in the local `[NDBD]` section.
- The parameter `NoOfReplicas`, which must be defined in the `[NDBD DEFAULT]` section, as it is common to all Cluster data nodes.

Most data node parameters are set in the `[NDBD DEFAULT]` section. Only those parameters explicitly stated as being able to set local values are allowed to be changed in the `[NDBD]` section. Where present, `HostName`, `Id` and `ExecuteOnComputer` *must* be defined in the local `[NDBD]` section, and not in any other section of `config.ini`. In other words, settings for these parameters are specific to one data node.

For those parameters affecting memory usage or buffer sizes, it is possible to use `K`, `M`, or `G` as a suffix to indicate units of 1024, 1024×1024, or 1024×1024×1024. (For example, `100K` means 100 × 1024 = 102400.) Parameter names and values are currently case-sensitive.

#### Identifying Data Nodes

The `Id` value (that is, the data node identifier) can be allocated on the command line when the node is started or in the configuration file.

- `Id`

This is the node ID used as the address of the node for all cluster internal messages. This is an integer in the range 1 to 63 inclusive. Each node in the cluster must have a unique identity.

-

### `ExecuteOnComputer`

This refers to the `Id` set for one of the computers defined in a `[COMPUTER]` section.

- 

### `HostName`

Specifying this parameter defines the hostname of the computer on which the data node is to reside. To specify a hostname other than `localhost`, either this parameter or `ExecuteOnComputer` is required.

- 

### `ServerPort` (*OBSOLETE*)

Each node in the cluster uses a port to connect to other nodes. This port is used also for non-TCP transporters in the connection setup phase. The default port is allocated dynamically in such a way as to ensure that no two nodes on the same computer receive the same port number, so it should not normally be necessary to specify a value for this parameter.

- 

### `NoOfReplicas`

This global parameter can be set only in the `[NDBD DEFAULT]` section, and defines the number of replicas for each table stored in the cluster. This parameter also specifies the size of node groups. A node group is a set of nodes all storing the same information.

Node groups are formed implicitly. The first node group is formed by the set of data nodes with the lowest node IDs, the next node group by the set of the next lowest node identities, and so on. By way of example, assume that we have 4 data nodes and that `NoOfReplicas` is set to 2. The four data nodes have node IDs 2, 3, 4 and 5. Then the first node group is formed from nodes 2 and 3, and the second node group by nodes 4 and 5. It is important to configure the cluster in such a manner that nodes in the same node groups are not placed on the same computer because a single hardware failure would cause the entire cluster to crash.

If no node IDs are provided, the order of the data nodes will be the determining factor for the node group. Whether or not explicit assignments are made, they can be viewed in the output of the management client's `SHOW` statement.

There is no default value for `NoOfReplicas`; the maximum possible value is 4.

- 

### `DataDir`

This parameter specifies the directory where trace files, log files, pid files and error logs are placed.

- 

### `FileSystemPath`

This parameter specifies the directory where all files created for metadata, REDO logs, UNDO logs and data files are placed. The default is the directory specified by `DataDir`. **Note:** This directory must exist before the `ndbd` process is initiated.

The recommended directory hierarchy for MySQL Cluster includes `/var/lib/mysql-cluster`, under which a directory for the node's filesystem is created. The name of this subdirectory contains the node ID. For example, if the node ID is 2, this subdirectory is named `ndb_2_fs`.

- 

### `BackupDataDir`



This parameter specifies the directory in which backups are placed. If omitted, the default backup location is the directory named `BACKUP` under the location specified by the `FileSystemPath` parameter. (See above.)

### Data Memory, Index Memory, and String Memory

`DataMemory` and `IndexMemory` are [NDBD] parameters specifying the size of memory segments used to store the actual records and their indexes. In setting values for these, it is important to understand how `DataMemory` and `IndexMemory` are used, as they usually need to be updated to reflect actual usage by the cluster:

- 

#### `DataMemory`

This parameter defines the amount of space (in bytes) available for storing database records. The entire amount specified by this value is allocated in memory, so it is extremely important that the machine has sufficient physical memory to accommodate it.

The memory allocated by `DataMemory` is used to store both the actual records and indexes. Each record is currently of fixed size. (Even `VARCHAR` columns are stored as fixed-width columns.) There is a 16-byte overhead on each record; an additional amount for each record is incurred because it is stored in a 32KB page with 128 byte page overhead (see below). There is also a small amount wasted per page due to the fact that each record is stored in only one page. The maximum record size is currently 8052 bytes.

The memory space defined by `DataMemory` is also used to store ordered indexes, which use about 10 bytes per record. Each table row is represented in the ordered index. A common error among users is to assume that all indexes are stored in the memory allocated by `IndexMemory`, but this is not the case: Only primary key and unique hash indexes use this memory; ordered indexes use the memory allocated by `DataMemory`. However, creating a primary key or unique hash index also creates an ordered index on the same keys, unless you specify `USING HASH` in the index creation statement. This can be verified by running `ndb_desc -d db_name table_name` in the management client.

The memory space allocated by `DataMemory` consists of 32KB pages, which are allocated to table fragments. Each table is normally partitioned into the same number of fragments as there are data nodes in the cluster. Thus, for each node, there are the same number of fragments as are set in `NoOfReplicas`.

Once a page has been allocated, it is currently not possible to return it to the pool of free pages, except by deleting the table. (This also means that `DataMemory` pages, once allocated to a given table, cannot be used by other tables.) Performing a node recovery also compresses the partition because all records are inserted into empty partitions from other live nodes.

The `DataMemory` memory space also contains UNDO information: For each update, a copy of the unaltered record is allocated in the `DataMemory`. There is also a reference to each copy in the ordered table indexes. Unique hash indexes are updated only when the unique index columns are updated, in which case a new entry in the index table is inserted and the old entry is deleted upon commit. For this reason, it is also necessary to allocate enough memory to handle the largest transactions performed by applications using the cluster. In any case, performing a few large transactions holds no advantage over using many smaller ones, for the following reasons:

- Large transactions are not any faster than smaller ones
- Large transactions increase the number of operations that are lost and must be repeated in event of transaction failure

- Large transactions use more memory

The default value for `DataMemory` is 80MB; the minimum is 1MB. There is no maximum size, but in reality the maximum size has to be adapted so that the process does not start swapping when the limit is reached. This limit is determined by the amount of physical RAM available on the machine and by the amount of memory that the operating system may commit to any one process. 32-bit operating systems are generally limited to 2–4GB per process; 64-bit operating systems can use more. For large databases, it may be preferable to use a 64-bit operating system for this reason. In addition, it is also possible to run more than one `ndbd` process per machine, and this may prove advantageous on machines with multiple CPUs.

- `IndexMemory`

This parameter controls the amount of storage used for hash indexes in MySQL Cluster. Hash indexes are always used for primary key indexes, unique indexes, and unique constraints. Note that when defining a primary key and a unique index, two indexes will be created, one of which is a hash index used for all tuple accesses as well as lock handling. It is also used to enforce unique constraints.

The size of the hash index is 25 bytes per record, plus the size of the primary key. For primary keys larger than 32 bytes another 8 bytes is added.

The default value for `IndexMemory` is 18MB. The minimum is 1MB.

- `StringMemory`

This parameter determines how much memory is allocated for strings such as table names, and is specified in an `[NDBD]` or `[NDBD DEFAULT]` section of the `config.ini` file. A value between 0 and 100 inclusive is interpreted as a percent of the maximum default value, which is calculated based on a number of factors including the number of tables, maximum table name size, maximum size of `.FRM` files, `MaxNoOfTriggers`, maximum column name size, and maximum default column value. In general it is safe to assume that the maximum default value is approximately 5 MB for a MySQL Cluster having 1000 tables.

A value greater than 100 is interpreted as a number of bytes.

In MySQL 4.1, the default value is 100 — that is, 100 percent of the default maximum, or roughly 5 MB. It is possible to reduce this value safely, but it should never be less than 5 percent. If you encounter Error 773 Out of string memory, please modify `StringMemory` config parameter: Permanent error: Schema error, this means that means that you have set the `StringMemory` value too low. 25 (25 percent) is not excessive, and should prevent this error from recurring in all but the most extreme conditions, as when there are hundreds or thousands of `NDB` tables with names whose lengths and columns whose number approach their permitted maximums.

The following example illustrates how memory is used for a table. Consider this table definition:

```
CREATE TABLE example (
  a INT NOT NULL,
  b INT NOT NULL,
  c INT NOT NULL,
  PRIMARY KEY(a),
  UNIQUE(b)
) ENGINE=NDBCLUSTER;
```

For each record, there are 12 bytes of data plus 12 bytes overhead. Having no nullable columns saves 4 bytes of overhead. In addition, we have two ordered indexes on columns `a` and `b` consuming roughly 10

bytes each per record. There is a primary key hash index on the base table using roughly 29 bytes per record. The unique constraint is implemented by a separate table with `b` as primary key and `a` as a column. This other table consumes an additional 29 bytes of index memory per record in the `example` table as well 8 bytes of record data plus 12 bytes of overhead.

Thus, for one million records, we need 58MB for index memory to handle the hash indexes for the primary key and the unique constraint. We also need 64MB for the records of the base table and the unique index table, plus the two ordered index tables.

You can see that hash indexes takes up a fair amount of memory space; however, they provide very fast access to the data in return. They are also used in MySQL Cluster to handle uniqueness constraints.

Currently, the only partitioning algorithm is hashing and ordered indexes are local to each node. Thus, ordered indexes cannot be used to handle uniqueness constraints in the general case.

An important point for both `IndexMemory` and `DataMemory` is that the total database size is the sum of all data memory and all index memory for each node group. Each node group is used to store replicated information, so if there are four nodes with two replicas, there will be two node groups. Thus, the total data memory available is  $2 \times \text{DataMemory}$  for each data node.

It is highly recommended that `DataMemory` and `IndexMemory` be set to the same values for all nodes. Data distribution is even over all nodes in the cluster, so the maximum amount of space available for any node can be no greater than that of the smallest node in the cluster.

`DataMemory` and `IndexMemory` can be changed, but decreasing either of these can be risky; doing so can easily lead to a node or even an entire MySQL Cluster that is unable to restart due to there being insufficient memory space. Increasing these values should be acceptable, but it is recommended that such upgrades are performed in the same manner as a software upgrade, beginning with an update of the configuration file, and then restarting the management server followed by restarting each data node in turn.

Updates do not increase the amount of index memory used. Inserts take effect immediately; however, rows are not actually deleted until the transaction is committed.

### Transaction Parameters

The next three `[NDBD]` parameters that we discuss are important because they affect the number of parallel transactions and the sizes of transactions that can be handled by the system. `MaxNoOfConcurrentTransactions` sets the number of parallel transactions possible in a node. `MaxNoOfConcurrentOperations` sets the number of records that can be in update phase or locked simultaneously.

Both of these parameters (especially `MaxNoOfConcurrentOperations`) are likely targets for users setting specific values and not using the default value. The default value is set for systems using small transactions, to ensure that these do not use excessive memory.

- `MaxNoOfConcurrentTransactions`

For each active transaction in the cluster there must be a record in one of the cluster nodes. The task of coordinating transactions is spread among the nodes. The total number of transaction records in the cluster is the number of transactions in any given node times the number of nodes in the cluster.

Transaction records are allocated to individual MySQL servers. Normally, there is at least one transaction record allocated per connection that using any table in the cluster. For this reason, one should ensure that there are more transaction records in the cluster than there are concurrent connections to all MySQL servers in the cluster.

This parameter must be set to the same value for all cluster nodes.

Changing this parameter is never safe and doing so can cause a cluster to crash. When a node crashes, one of the nodes (actually the oldest surviving node) will build up the transaction state of all transactions ongoing in the crashed node at the time of the crash. It is thus important that this node has as many transaction records as the failed node.

The default value is 4096.

- 

#### `MaxNoOfConcurrentOperations`

It is a good idea to adjust the value of this parameter according to the size and number of transactions. When performing transactions of only a few operations each and not involving a great many records, there is no need to set this parameter very high. When performing large transactions involving many records need to set this parameter higher.

Records are kept for each transaction updating cluster data, both in the transaction coordinator and in the nodes where the actual updates are performed. These records contain state information needed to find UNDO records for rollback, lock queues, and other purposes.

This parameter should be set to the number of records to be updated simultaneously in transactions, divided by the number of cluster data nodes. For example, in a cluster which has four data nodes and which is expected to handle 1,000,000 concurrent updates using transactions, you should set this value to  $1000000 / 4 = 250000$ .

Read queries which set locks also cause operation records to be created. Some extra space is allocated within individual nodes to accommodate cases where the distribution is not perfect over the nodes.

When queries make use of the unique hash index, there are actually two operation records used per record in the transaction. The first record represents the read in the index table and the second handles the operation on the base table.

The default value is 32768.

This parameter actually handles two values that can be configured separately. The first of these specifies how many operation records are to be placed with the transaction coordinator. The second part specifies how many operation records are to be local to the database.

A very large transaction performed on an eight-node cluster requires as many operation records in the transaction coordinator as there are reads, updates, and deletes involved in the transaction. However, the operation records of the are spread over all eight nodes. Thus, if it is necessary to configure the system for one very large transaction, it is a good idea to configure the two parts separately. `MaxNoOfConcurrentOperations` will always be used to calculate the number of operation records in the transaction coordinator portion of the node.

It is also important to have an idea of the memory requirements for operation records. These consume about 1KB per record.

- 

#### `MaxNoOfLocalOperations`

By default, this parameter is calculated as  $1.1 \times \text{MaxNoOfConcurrentOperations}$ . This fits systems with many simultaneous transactions, none of them being very large. If there is a need to handle one very large transaction at a time and there are many nodes, it is a good idea to override the default value by explicitly specifying this parameter.

### Transaction Temporary Storage

The next set of [NDBD] parameters is used to determine temporary storage when executing a statement that is part of a Cluster transaction. All records are released when the statement is completed and the cluster is waiting for the commit or rollback.

The default values for these parameters are adequate for most situations. However, users with a need to support transactions involving large numbers of rows or operations may need to increase these values to enable better parallelism in the system, whereas users whose applications require relatively small transactions can decrease the values to save memory.

- `MaxNoOfConcurrentIndexOperations`

For queries using a unique hash index, another temporary set of operation records is used during a query's execution phase. This parameter sets the size of that pool of records. Thus, this record is allocated only while executing a part of a query. As soon as this part has been executed, the record is released. The state needed to handle aborts and commits is handled by the normal operation records, where the pool size is set by the parameter `MaxNoOfConcurrentOperations`.

The default value of this parameter is 8192. Only in rare cases of extremely high parallelism using unique hash indexes should it be necessary to increase this value. Using a smaller value is possible and can save memory if the DBA is certain that a high degree of parallelism is not required for the cluster.

- `MaxNoOfFiredTriggers`

The default value of `MaxNoOfFiredTriggers` is 4000, which is sufficient for most situations. In some cases it can even be decreased if the DBA feels certain the need for parallelism in the cluster is not high.

A record is created when an operation is performed that affects a unique hash index. Inserting or deleting a record in a table with unique hash indexes or updating a column that is part of a unique hash index fires an insert or a delete in the index table. The resulting record is used to represent this index table operation while waiting for the original operation that fired it to complete. This operation is short-lived but can still require a large number of records in its pool for situations with many parallel write operations on a base table containing a set of unique hash indexes.

- `TransactionBufferMemory`

The memory affected by this parameter is used for tracking operations fired when updating index tables and reading unique indexes. This memory is used to store the key and column information for these operations. It is only very rarely that the value for this parameter needs to be altered from the default.

The default value for `TransactionBufferMemory` is 1MB.

Normal read and write operations use a similar buffer, whose usage is even more short-lived. The compile-time parameter `ZATTRBUF_FILESIZE` (found in `ndb/src/kernel/blocks/Dbtc/Dbtc.hpp`) set to  $4000 \times 128$  bytes (500KB). A similar buffer for key information, `ZDATABUF_FILESIZE` (also in `Dbtc.hpp`) contains  $4000 \times 16 = 62.5$ KB of buffer space. `Dbtc` is the module that handles transaction coordination.

## Scans and Buffering

There are additional [NDBD] parameters in the `Dblqh` module (in `ndb/src/kernel/blocks/Dblqh/Dblqh.hpp`) that affect reads and updates. These include `ZAT-`

`TRINIBUF_FILESIZE`, set by default to 10000 × 128 bytes (1250KB) and `ZDATABUF_FILE_SIZE`, set by default to 10000\*16 bytes (roughly 156KB) of buffer space. To date, there have been neither any reports from users nor any results from our own extensive tests suggesting that either of these compile-time limits should be increased.

- `MaxNoOfConcurrentScans`

This parameter is used to control the number of parallel scans that can be performed in the cluster. Each transaction coordinator can handle the number of parallel scans defined for this parameter. Each scan query is performed by scanning all partitions in parallel. Each partition scan uses a scan record in the node where the partition is located, the number of records being the value of this parameter times the number of nodes. The cluster should be able to sustain `MaxNoOfConcurrentScans` scans concurrently from all nodes in the cluster.

Scans are actually performed in two cases. The first of these cases occurs when no hash or ordered indexes exists to handle the query, in which case the query is executed by performing a full table scan. The second case is encountered when there is no hash index to support the query but there is an ordered index. Using the ordered index means executing a parallel range scan. The order is kept on the local partitions only, so it is necessary to perform the index scan on all partitions.

The default value of `MaxNoOfConcurrentScans` is 256. The maximum value is 500.

This parameter specifies the number of scans possible in the transaction coordinator. If the number of local scan records is not provided, it is calculated as the product of `MaxNoOfConcurrentScans` and the number of data nodes in the system.

- `MaxNoOfLocalScans`

Specifies the number of local scan records if many scans are not fully parallelized.

- `BatchSizePerLocalScan`

This parameter is used to calculate the number of lock records which must be there to handle many concurrent scan operations.

The default value is 64; this value has a strong connection to the `ScanBatchSize` defined in the SQL nodes.

- `LongMessageBuffer`

This is an internal buffer used for passing messages within individual nodes and between nodes. Although it is highly unlikely that this would need to be changed, it is configurable. By default, it is set to 1MB.

## Logging and Checkpointing

These `[NDBD]` parameters control log and checkpoint behavior.

- `NoOfFragmentLogFiles`

This parameter sets the size of the node's REDO log files. REDO log files are organized in a ring. It is extremely important that the first and last log files (sometimes referred to as the “head” and “tail” log files, respectively) do not meet. When these approach one another too closely, the node begins

aborting all transactions encompassing updates due to a lack of room for new log records.

A REDO log record is not removed until three local checkpoints have been completed since that log record was inserted. Checkpointing frequency is determined by its own set of configuration parameters discussed elsewhere in this chapter.

How these parameters interact and proposals for how to configure them are discussed in [Section 15.4.6, “Configuring Parameters for Local Checkpoints”](#).

The default parameter value is 8, which means 8 sets of 4 16MB files for a total of 512MB. In other words, REDO log space must be allocated in blocks of 64MB. In scenarios requiring a great many updates, the value for `NoOfFragmentLogFiles` may need to be set as high as 300 or even higher to provide sufficient space for REDO logs.

If the checkpointing is slow and there are so many writes to the database that the log files are full and the log tail cannot be cut without jeopardizing recovery, all updating transactions are aborted with internal error code 410 (*Out of log file space temporarily*). This condition prevails until a checkpoint has completed and the log tail can be moved forward.

**Important:** This parameter cannot be changed “on the fly”; you must restart the node using `--initial`. If you wish to change this value for a running cluster, you can do so via a rolling node restart.

- `MaxNoOfOpenFiles`

This parameter sets a ceiling on how many internal threads to allocate for open files. *Any situation requiring a change in this parameter should be reported as a bug.*

The default value is 40.

- `MaxNoOfSavedMessages`

This parameter sets the maximum number of trace files that are kept before overwriting old ones. Trace files are generated when, for whatever reason, the node crashes.

The default is 25 trace files.

## Metadata Objects

The next set of `[NDBD]` parameters defines pool sizes for metadata objects, used to define the maximum number of attributes, tables, indexes, and trigger objects used by indexes, events, and replication between clusters. Note that these act merely as “suggestions” to the cluster, and any that are not specified revert to the default values shown.

- `MaxNoOfAttributes`

Defines the number of attributes that can be defined in the cluster.

The default value is 1000, with the minimum possible value being 32. The maximum is 4294967039. Each attribute consumes around 200 bytes of storage per node due to the fact that all metadata is fully replicated on the servers.

When setting `MaxNoOfAttributes`, it is important to prepare in advance for any `ALTER TABLE` statements that you might want to perform in the future. This is due to the fact, during the exe-

cution of `ALTER TABLE` on a Cluster table, 3 times the number of attributes as in the original table are used. For example, if a table requires 100 attributes, and you want to be able to alter it later, you need to set the value of `MaxNoOfAttributes` to 300. Assuming that you can create all desired tables without any problems, a good rule of thumb is to add two times the number of attributes in the largest table to `MaxNoOfAttributes` to be sure. You should also verify that this number is sufficient by trying an actual `ALTER TABLE` after configuring the parameter. If this is not successful, increase `MaxNoOfAttributes` by another multiple of the original value and test it again.

- `MaxNoOfTables`

A table object is allocated for each table, unique hash index, and ordered index. This parameter sets the maximum number of table objects for the cluster as a whole.

For each attribute that has a `BLOB` data type an extra table is used to store most of the `BLOB` data. These tables also must be taken into account when defining the total number of tables.

The default value of this parameter is 128. The minimum is 8 and the maximum is 1600. Each table object consumes approximately 20KB per node.

- `MaxNoOfOrderedIndexes`

For each ordered index in the cluster, an object is allocated describing what is being indexed and its storage segments. By default, each index so defined also defines an ordered index. Each unique index and primary key has both an ordered index and a hash index.

The default value of this parameter is 128. Each object consumes approximately 10KB of data per node.

- `MaxNoOfUniqueHashIndexes`

For each unique index that is not a primary key, a special table is allocated that maps the unique key to the primary key of the indexed table. By default, an ordered index is also defined for each unique index. To prevent this, you must specify the `USING HASH` option when defining the unique index.

The default value is 64. Each index consumes approximately 15KB per node.

- `MaxNoOfTriggers`

Internal update, insert, and delete triggers are allocated for each unique hash index. (This means that three triggers are created for each unique hash index.) However, an *ordered* index requires only a single trigger object. Backups also use three trigger objects for each normal table in the cluster.

This parameter sets the maximum number of trigger objects in the cluster.

The default value is 768.

- `MaxNoOfIndexes`

This parameter was deprecated in MySQL 4.1.5. In MySQL 4.1.6 and newer versions, you should use `MaxNoOfOrderedIndexes` and `MaxNoOfUniqueHashIndexes` instead.

This parameter is used only by unique hash indexes. There needs to be one record in this pool for each unique hash index defined in the cluster.

The default value of this parameter is 128.



## Boolean Parameters

The behavior of data nodes is also affected by a set of [NDBD] parameters taking on boolean values. These parameters can each be specified as `TRUE` by setting them equal to `1` or `Y`, and as `FALSE` by setting them equal to `0` or `N`.

- `LockPagesInMainMemory`

For a number of operating systems, including Solaris and Linux, it is possible to lock a process into memory and so avoid any swapping to disk. This can be used to help guarantee the cluster's real-time characteristics.

This feature is disabled by default.

- `StopOnError`

This parameter specifies whether an `ndbd` process should exit or perform an automatic restart when an error condition is encountered.

This feature is enabled by default.

- `Diskless`

It is possible to specify MySQL Cluster tables as *diskless*, meaning that tables are not checkpointed to disk and that no logging occurs. Such tables exist only in main memory. A consequence of using diskless tables is that neither the tables nor the records in those tables survive a crash. However, when operating in diskless mode, it is possible to run `ndbd` on a diskless computer.

**Important:** This feature causes the *entire* cluster to operate in diskless mode.

When this feature is enabled, Cluster online backup is disabled. In addition, a partial start of the cluster is not possible.

`Diskless` is disabled by default.

- `RestartOnErrorInsert`

This feature is accessible only when building the debug version where it is possible to insert errors in the execution of individual blocks of code as part of testing.

This feature is disabled by default.

## Controlling Timeouts, Intervals, and Disk Paging

There are a number of [NDBD] parameters specifying timeouts and intervals between various actions in Cluster data nodes. Most of the timeout values are specified in milliseconds. Any exceptions to this are mentioned where applicable.

- `TimeBetweenWatchDogCheck`

To prevent the main thread from getting stuck in an endless loop at some point, a “watchdog” thread checks the main thread. This parameter specifies the number of milliseconds between checks. If the

process remains in the same state after three checks, the watchdog thread terminates it.

This parameter can easily be changed for purposes of experimentation or to adapt to local conditions. It can be specified on a per-node basis although there seems to be little reason for doing so.

The default timeout is 4000 milliseconds (4 seconds).

- `StartPartialTimeout`

This parameter specifies how long the Cluster waits for all data nodes to come up before the cluster initialization routine is invoked. This timeout is used to avoid a partial Cluster startup whenever possible.

The default value is 30000 milliseconds (30 seconds). 0 disables the timeout. In other words, the cluster may start only if all nodes are available.

- `StartPartitionedTimeout`

If the cluster is ready to start after waiting for `StartPartialTimeout` milliseconds but is still possibly in a partitioned state, the cluster waits until this timeout has also passed.

The default timeout is 60000 milliseconds (60 seconds).

- `StartFailureTimeout`

If a data node has not completed its startup sequence within the time specified by this parameter, the node startup fails. Setting this parameter to 0 (the default value) means that no data node timeout is applied.

For nonzero values, this parameter is measured in milliseconds. For data nodes containing extremely large amounts of data, this parameter should be increased. For example, in the case of a data node containing several gigabytes of data, a period as long as 10–15 minutes (that is, 600000 to 1000000 milliseconds) might be required to to perform a node restart.

- `HeartbeatIntervalDbDb`

One of the primary methods of discovering failed nodes is by the use of heartbeats. This parameter states how often heartbeat signals are sent and how often to expect to receive them. After missing three heartbeat intervals in a row, the node is declared dead. Thus, the maximum time for discovering a failure through the heartbeat mechanism is four times the heartbeat interval.

The default heartbeat interval is 1500 milliseconds (1.5 seconds). This parameter must not be changed drastically and should not vary widely between nodes. If one node uses 5000 milliseconds and the node watching it uses 1000 milliseconds, obviously the node will be declared dead very quickly. This parameter can be changed during an online software upgrade, but only in small increments.

- `HeartbeatIntervalDbApi`

Each data node sends heartbeat signals to each MySQL server (SQL node) to ensure that it remains in contact. If a MySQL server fails to send a heartbeat in time it is declared “dead,” in which case all ongoing transactions are completed and all resources released. The SQL node cannot reconnect until all activities initiated by the previous MySQL instance have been completed. The three-heartbeat cri-

teria for this determination are the same as described for [HeartbeatIntervalDbDb](#).

The default interval is 1500 milliseconds (1.5 seconds). This interval can vary between individual data nodes because each data node watches the MySQL servers connected to it, independently of all other data nodes.

- [TimeBetweenLocalCheckpoints](#)

This parameter is an exception in that it does not specify a time to wait before starting a new local checkpoint; rather, it is used to ensure that local checkpoints are not performed in a cluster where relatively few updates are taking place. In most clusters with high update rates, it is likely that a new local checkpoint is started immediately after the previous one has been completed.

The size of all write operations executed since the start of the previous local checkpoints is added. This parameter is also exceptional in that it is specified as the base-2 logarithm of the number of 4-byte words, so that the default value 20 means 4MB ( $4 \times 2^{20}$ ) of write operations, 21 would mean 8MB, and so on up to a maximum value of 31, which equates to 8GB of write operations.

All the write operations in the cluster are added together. Setting [TimeBetweenLocalCheckpoints](#) to 6 or less means that local checkpoints will be executed continuously without pause, independent of the cluster's workload.

- [TimeBetweenGlobalCheckpoints](#)

When a transaction is committed, it is committed in main memory in all nodes on which the data is mirrored. However, transaction log records are not flushed to disk as part of the commit. The reasoning behind this behavior is that having the transaction safely committed on at least two autonomous host machines should meet reasonable standards for durability.

It is also important to ensure that even the worst of cases — a complete crash of the cluster — is handled properly. To guarantee that this happens, all transactions taking place within a given interval are put into a global checkpoint, which can be thought of as a set of committed transactions that has been flushed to disk. In other words, as part of the commit process, a transaction is placed in a global checkpoint group. Later, this group's log records are flushed to disk, and then the entire group of transactions is safely committed to disk on all computers in the cluster.

This parameter defines the interval between global checkpoints. The default is 2000 milliseconds.

- [TimeBetweenInactiveTransactionAbortCheck](#)

Timeout handling is performed by checking a timer on each transaction once for every interval specified by this parameter. Thus, if this parameter is set to 1000 milliseconds, every transaction will be checked for timing out once per second.

The default value is 1000 milliseconds (1 second).

- [TransactionInactiveTimeout](#)

This parameter states the maximum time that is permitted to lapse between operations in the same transaction before the transaction is aborted.

The default for this parameter is zero (no timeout). For a real-time database that needs to ensure that no transaction keeps locks for too long, this parameter should be set to a much smaller value. The unit is milliseconds.

- [TransactionDeadlockDetectionTimeout](#)

When a node executes a query involving a transaction, the node waits for the other nodes in the cluster to respond before continuing. A failure to respond can occur for any of the following reasons:

- The node is “dead”
- The operation has entered a lock queue
- The node requested to perform the action could be heavily overloaded.

This timeout parameter states how long the transaction coordinator waits for query execution by another node before aborting the transaction, and is important for both node failure handling and deadlock detection. Setting it too high can cause a undesirable behavior in situations involving deadlocks and node failure.

The default timeout value is 1200 milliseconds (1.2 seconds).

- [NoOfDiskPagesToDiskAfterRestartTUP](#)

When executing a local checkpoint, the algorithm flushes all data pages to disk. Merely doing so as quickly as possible without any moderation is likely to impose excessive loads on processors, networks, and disks. To control the write speed, this parameter specifies how many pages per 100 milliseconds are to be written. In this context, a “page” is defined as 8KB. This parameter is specified in units of 80KB per second, so , setting [NoOfDiskPagesToDiskAfterRestartTUP](#) to a value of 20 entails writing 1.6MB in data pages to disk each second during a local checkpoint. This value includes the writing of UNDO log records for data pages. That is, this parameter handles the limitation of writes from data memory. UNDO log records for index pages are handled by the parameter [NoOfDiskPagesToDiskAfterRestartACC](#). (See the entry for [IndexMemory](#) for information about index pages.)

In short, this parameter specifies how quickly to execute local checkpoints. It operates in conjunction with [NoOfFragmentLogFiles](#), [DataMemory](#), and [IndexMemory](#).

For more information about the interaction between these parameters and possible strategies for choosing appropriate values for them, see [Section 15.4.6, “Configuring Parameters for Local Checkpoints”](#).

The default value is 40 (3.2MB of data pages per second).

- [NoOfDiskPagesToDiskAfterRestartACC](#)

This parameter uses the same units as [NoOfDiskPagesToDiskAfterRestartTUP](#) and acts in a similar fashion, but limits the speed of writing index pages from index memory.

The default value of this parameter is 20 (1.6MB of index memory pages per second).

- [NoOfDiskPagesToDiskDuringRestartTUP](#)

This parameter is used in a fashion similar to [NoOfDiskPagesToDiskAfterRestartTUP](#) and [NoOfDiskPagesToDiskAfterRestartACC](#), only it does so with regard to local checkpoints executed in the node when a node is restarting. A local checkpoint is always performed as part of all node restarts. During a node restart it is possible to write to disk at a higher speed than at other times, because fewer activities are being performed in the node.

This parameter covers pages written from data memory.

The default value is 40 (3.2MB per second).

- [NoOfDiskPagesToDiskDuringRestartACC](#)

Controls the number of index memory pages that can be written to disk during the local checkpoint phase of a node restart.

As with [NoOfDiskPagesToDiskAfterRestartTUP](#) and [NoOfDiskPagesToDiskAfterRestartACC](#), values for this parameter are expressed in terms of 8KB pages written per 100 milliseconds (80KB/second).

The default value is 20 (1.6MB per second).

- [ArbitrationTimeout](#)

This parameter specifies how long data nodes wait for a response from the arbitrator to an arbitration message. If this is exceeded, the network is assumed to have split.

The default value is 1000 milliseconds (1 second).

## Buffering and Logging

Several [\[NDBD\]](#) configuration parameters corresponding to former compile-time parameters were introduced in MySQL 4.1.5. These enable the advanced user to have more control over the resources used by node processes and to adjust various buffer sizes at need.

These buffers are used as front ends to the file system when writing log records to disk. If the node is running in diskless mode, these parameters can be set to their minimum values without penalty due to the fact that disk writes are “faked” by the [NDB](#) storage engine's filesystem abstraction layer.

- [UndoIndexBuffer](#)

The UNDO index buffer, whose size is set by this parameter, is used during local checkpoints. The [NDB](#) storage engine uses a recovery scheme based on checkpoint consistency in conjunction with an operational REDO log. To produce a consistent checkpoint without blocking the entire system for writes, UNDO logging is done while performing the local checkpoint. UNDO logging is activated on a single table fragment at a time. This optimization is possible because tables are stored entirely in main memory.

The UNDO index buffer is used for the updates on the primary key hash index. Inserts and deletes rearrange the hash index; the [NDB](#) storage engine writes UNDO log records that map all physical changes to an index page so that they can be undone at system restart. It also logs all active insert operations for each fragment at the start of a local checkpoint.

Reads and updates set lock bits and update a header in the hash index entry. These changes are handled by the page-writing algorithm to ensure that these operations need no UNDO logging.

This buffer is 2MB by default. The minimum value is 1MB, which is sufficient for most applications. For applications doing extremely large or numerous inserts and deletes together with large transactions and large primary keys, it may be necessary to increase the size of this buffer. If this buffer is too small, the [NDB](#) storage engine issues internal error code 677 ([Index UNDO buf-](#)

fers overloaded).

**Important:** It is not safe to decrease the value of this parameter during a rolling restart.

- 

#### UndoDataBuffer

This parameter sets the size of the UNDO data buffer, which performs a function similar to that of the UNDO index buffer, except the UNDO data buffer is used with regard to data memory rather than index memory. This buffer is used during the local checkpoint phase of a fragment for inserts, deletes, and updates.

Because UNDO log entries tend to grow larger as more operations are logged, this buffer is also larger than its index memory counterpart, with a default value of 16MB.

This amount of memory may be unnecessarily large for some applications. In such cases, it is possible to decrease this size to a minimum of 1MB.

It is rarely necessary to increase the size of this buffer. If there is such a need, it is a good idea to check whether the disks can actually handle the load caused by database update activity. A lack of sufficient disk space cannot be overcome by increasing the size of this buffer.

If this buffer is too small and gets congested, the NDB storage engine issues internal error code 891 (Data UNDO buffers overloaded).

**Important:** It is not safe to decrease the value of this parameter during a rolling restart.

- 

#### RedoBuffer

All update activities also need to be logged. The REDO log makes it possible to replay these updates whenever the system is restarted. The NDB recovery algorithm uses a “fuzzy” checkpoint of the data together with the UNDO log, and then applies the REDO log to play back all changes up to the restoration point.

`RedoBuffer` sets the size of the buffer in which the REDO log is written, and is 8MB by default. The minimum value is 1MB.

If this buffer is too small, the NDB storage engine issues error code 1221 (`REDO log buffers overloaded`).

**Important:** It is not safe to decrease the value of this parameter during a rolling restart.

### Controlling Log Messages

In managing the cluster, it is very important to be able to control the number of log messages sent for various event types to `stdout`. For each event category, there are 16 possible event levels (numbered 0 through 15). Setting event reporting for a given event category to level 15 means all event reports in that category are sent to `stdout`; setting it to 0 means that there will be no event reports made in that category.

By default, only the startup message is sent to `stdout`, with the remaining event reporting level defaults being set to 0. The reason for this is that these messages are also sent to the management server's cluster log.

An analogous set of levels can be set for the management client to determine which event levels to record in the cluster log.

- `LogLevelStartup`  
The reporting level for events generated during startup of the process.  
The default level is 1.
- `LogLevelShutdown`  
The reporting level for events generated as part of graceful shutdown of a node.  
The default level is 0.
- `LogLevelStatistic`  
The reporting level for statistical events such as number of primary key reads, number of updates, number of inserts, information relating to buffer usage, and so on.  
The default level is 0.
- `LogLevelCheckpoint`  
The reporting level for events generated by local and global checkpoints.  
The default level is 0.
- `LogLevelNodeRestart`  
The reporting level for events generated during node restart.  
The default level is 0.
- `LogLevelConnection`  
The reporting level for events generated by connections between cluster nodes.  
The default level is 0.
- `LogLevelError`  
The reporting level for events generated by errors and warnings by the cluster as a whole. These errors do not cause any node failure but are still considered worth reporting.  
The default level is 0.
- `LogLevelInfo`  
The reporting level for events generated for information about the general state of the cluster.  
The default level is 0.

## Backup Parameters

The [NDBD] parameters discussed in this section define memory buffers set aside for execution of on-line backups.

- [BackupDataBufferSize](#)

In creating a backup, there are two buffers used for sending data to the disk. The backup data buffer is used to fill in data recorded by scanning a node's tables. Once this buffer has been filled to the level specified as [BackupWriteSize](#) (see below), the pages are sent to disk. While flushing data to disk, the backup process can continue filling this buffer until it runs out of space. When this happens, the backup process pauses the scan and waits until some disk writes have completed freed up memory so that scanning may continue.

The default value is 2MB.

- [BackupLogBufferSize](#)

The backup log buffer fulfills a role similar to that played by the backup data buffer, except that it is used for generating a log of all table writes made during execution of the backup. The same principles apply for writing these pages as with the backup data buffer, except that when there is no more space in the backup log buffer, the backup fails. For that reason, the size of the backup log buffer must be large enough to handle the load caused by write activities while the backup is being made. See [Section 15.8.4, “Configuration for Cluster Backup”](#).

The default value for this parameter should be sufficient for most applications. In fact, it is more likely for a backup failure to be caused by insufficient disk write speed than it is for the backup log buffer to become full. If the disk subsystem is not configured for the write load caused by applications, the cluster is unlikely to be able to perform the desired operations.

It is preferable to configure cluster nodes in such a manner that the processor becomes the bottleneck rather than the disks or the network connections.

The default value is 2MB.

- [BackupMemory](#)

This parameter is simply the sum of [BackupDataBufferSize](#) and [BackupLogBufferSize](#).

The default value is  $2\text{MB} + 2\text{MB} = 4\text{MB}$ .

**Important:** If [BackupDataBufferSize](#) and [BackupLogBufferSize](#) taken together exceed 4MB, then this parameter must be set explicitly in the `config.ini` file to their sum.

- [BackupWriteSize](#)

This parameter specifies the default size of messages written to disk by the backup log and backup data buffers.

- [BackupMaxWriteSize](#)

This parameter specifies the maximum size of messages written to disk by the backup log and backup data buffers.

The default value is 256KB.



**Important:** When specifying these parameters, the following relationships must hold true. Otherwise, the data node will be unable to start.

- `BackupDataBufferSize >= BackupWriteSize + 188KB`
- `BackupLogBufferSize >= BackupWriteSize + 16KB`
- `BackupMaxWriteSize >= BackupWriteSize`

### 15.4.4.6. Defining SQL and Other API Nodes

The `[MYSQLD]` and `[API]` sections in the `config.ini` file define the behavior of the MySQL servers (SQL nodes) and other applications (API nodes) used to access cluster data. None of the parameters shown is required. If no computer or host name is provided, any host can use this SQL or API node.

Generally speaking, a `[MYSQLD]` section is used to indicate a MySQL server providing an SQL interface to the cluster, and an `[API]` section is used for applications other than `mysqld` processes accessing cluster data, but the two designations are actually synonymous; you can, for instance, list parameters for a MySQL server acting as an SQL node in an `[API]` section.

- `Id`

The `Id` value is used to identify the node in all cluster internal messages. It must be an integer in the range 1 to 63 inclusive, and must be unique among all node IDs within the cluster.
- `ExecuteOnComputer`

This refers to the `Id` set for one of the computers (hosts) defined in a `[COMPUTER]` section of the configuration file.
- `HostName`

Specifying this parameter defines the hostname of the computer on which the SQL node (API node) is to reside. To specify a hostname other than `localhost`, either this parameter or `ExecuteOnComputer` is required.
- `ArbitrationRank`

This parameter defines which nodes can act as arbitrators. Both MGM nodes and SQL nodes can be arbitrators. A value of 0 means that the given node is never used as an arbitrator, a value of 1 gives the node high priority as an arbitrator, and a value of 2 gives it low priority. A normal configuration uses the management server as arbitrator, setting its `ArbitrationRank` to 1 (the default) and those for all SQL nodes to 0.
- `ArbitrationDelay`

Setting this parameter to any other value than 0 (the default) means that responses by the arbitrator to arbitration requests will be delayed by the stated number of milliseconds. It is usually not necessary to change this value.
-

### BatchByteSize

For queries that are translated into full table scans or range scans on indexes, it is important for best performance to fetch records in properly sized batches. It is possible to set the proper size both in terms of number of records ([BatchSize](#)) and in terms of bytes ([BatchByteSize](#)). The actual batch size is limited by both parameters.

The speed at which queries are performed can vary by more than 40% depending upon how this parameter is set. In future releases, MySQL Server will make educated guesses on how to set parameters relating to batch size, based on the query type.

This parameter is measured in bytes and by default is equal to 32KB.

- 

### BatchSize

This parameter is measured in number of records and is by default set to 64. The maximum size is 992.

- 

### MaxScanBatchSize

The batch size is the size of each batch sent from each data node. Most scans are performed in parallel to protect the MySQL Server from receiving too much data from many nodes in parallel; this parameter sets a limit to the total batch size over all nodes.

The default value of this parameter is set to 256KB. Its maximum size is 16MB.

You can obtain some information from a MySQL server running as a Cluster SQL node using [SHOW STATUS](#) in the `mysql` client, as shown here:

```
mysql> SHOW STATUS LIKE 'ndb%';
+-----+-----+
| Variable_name | Value |
+-----+-----+
| Ndb_cluster_node_id | 5 |
| Ndb_config_from_host | 192.168.0.112 |
| Ndb_config_from_port | 1186 |
| Ndb_number_of_storage_nodes | 4 |
+-----+-----+
4 rows in set (0.02 sec)
```

For information about these Cluster system status variables, see [Section 5.2.4, “Status Variables”](#).

## 15.4.4.7. Cluster TCP/IP Connections

TCP/IP is the default transport mechanism for establishing connections in MySQL Cluster. It is normally not necessary to define connections because Cluster automatically set ups a connection between each of the data nodes, between each data node and all MySQL server nodes, and between each data node and the management server. (For one exception to this rule, see [Section 15.4.4.8, “TCP/IP Connections Using Direct Connections”](#).) [TCP] sections in the `config.ini` file explicitly define TCP/IP connections between nodes in the cluster.

It is only necessary to define a connection to override the default connection parameters. In that case, it is necessary to define at least `NodeId1`, `NodeId2`, and the parameters to change.

It is also possible to change the default values for these parameters by setting them in the [TCP DEFAULT] section.

- `NodeId1, NodeId2`

To identify a connection between two nodes it is necessary to provide their node IDs in the `[TCP]` section of the configuration file. These are the same unique `Id` values for each of these nodes as described in [Section 15.4.4.6, “Defining SQL and Other API Nodes”](#).

- `SendBufferMemory`

TCP transporters use a buffer to store all messages before performing the send call to the operating system. When this buffer reaches 64KB its contents are sent; these are also sent when a round of messages have been executed. To handle temporary overload situations it is also possible to define a bigger send buffer. The default size of the send buffer is 256KB.

- `SendSignalId`

To be able to retrace a distributed message datagram, it is necessary to identify each message. When this parameter is set to `Y`, message IDs are transported over the network. This feature is disabled by default.

- `Checksum`

This parameter is a boolean parameter (enabled by setting it to `Y` or `1`, disabled by setting it to `N` or `0`). It is disabled by default. When it is enabled, checksums for all messages are calculated before they placed in the send buffer. This feature ensures that messages are not corrupted while waiting in the send buffer, or by the transport mechanism.

- `PortNumber` (*OBSOLETE*)

This formerly specified the port number to be used for listening for connections from other nodes. This parameter should no longer be used.

- `ReceiveBufferMemory`

Specifies the size of the buffer used when receiving data from the TCP/IP socket. There is seldom any need to change this parameter from its default value of 64KB, except possibly to save memory.

### 15.4.4.8. TCP/IP Connections Using Direct Connections

Setting up a cluster using direct connections between data nodes requires specifying explicitly the crossover IP addresses of the data nodes so connected in the `[TCP]` section of the cluster `config.ini` file.

In the following example, we envision a cluster with at least four hosts, one each for a management server, an SQL node, and two data nodes. The cluster as a whole resides on the `172.23.72.*` subnet of a LAN. In addition to the usual network connections, the two data nodes are connected directly using a standard crossover cable, and communicate with one another directly using IP addresses in the `1.1.0.*` address range as shown:

```
# Management Server
[NDB_MGMD]
Id=1
HostName=172.23.72.20

# SQL Node
[MYSQLD]
```

```

Id=2
HostName=172.23.72.21

# Data Nodes
[NDBD]
Id=3
HostName=172.23.72.22

[NDBD]
Id=4
HostName=172.23.72.23

# TCP/IP Connections
[TCP]
NodeId1=3
NodeId2=4
HostName1=1.1.0.1
HostName2=1.1.0.2

```

The `HostName $N$`  parameter, where  $N$  is an integer, is used only when specifying direct TCP/IP connections.

The use of direct connections between data nodes can improve the cluster's overall efficiency by allowing the data nodes to bypass an Ethernet device such as a switch, hub, or router, thus cutting down on the cluster's latency. It is important to note that to take the best advantage of direct connections in this fashion with more than two data nodes, you must have a direct connection between each data node and every other data node in the same node group.

### 15.4.4.9. Shared-Memory Connections

Beginning with MySQL 4.1.9-max, MySQL Cluster will attempt to use the shared memory transporter and configure it automatically where possible. (In previous versions of MySQL Cluster, shared memory segments functioned only when the `-max` binary was built using `--with-ndb-shm`.) [SHM] sections in the `config.ini` file explicitly define shared-memory connections between nodes in the cluster. When explicitly defining shared memory as the connection method, it is necessary to define at least `NodeId1`, `NodeId2` and `ShmKey`. All other parameters have default values that should work well in most cases.

**Important:** *SHM functionality is considered experimental only.* It is not officially supported in MySQL 4.1. This means that you must determine for yourself or by using our free resources (forums, mailing lists) whether it can be made to work correctly in your specific case.

- `NodeId1`, `NodeId2`

To identify a connection between two nodes it is necessary to provide node identifiers for each of them, as `NodeId1` and `NodeId2`.

- `ShmKey`

When setting up shared memory segments, a node ID, expressed as an integer, is used to identify uniquely the shared memory segment to use for the communication. There is no default value.

- `ShmSize`

Each SHM connection has a shared memory segment where messages between nodes are placed by the sender and read by the reader. The size of this segment is defined by `ShmSize`. The default value is 1MB.

-

### SendSignalId

To retrace the path of a distributed message, it is necessary to provide each message with a unique identifier. Setting this parameter to **Y** causes these message IDs to be transported over the network as well. This feature is disabled by default.

•

### Checksum

This parameter is a boolean (**Y/N**) parameter which is disabled by default. When it is enabled, checksums for all messages are calculated before being placed in the send buffer.

This feature prevents messages from being corrupted while waiting in the send buffer. It also serves as a check against data being corrupted during transport.

## 15.4.4.10. SCI Transport Connections

[**SCI**] sections in the `config.ini` file explicitly define SCI (Scalable Coherent Interface) connections between cluster nodes. Using SCI transporters in MySQL Cluster is supported only when the MySQL-Max binaries are built using `--with-ndb-sci=/your/path/to/SCI`. The `path` should point to a directory that contains at a minimum `lib` and `include` directories containing SISI libraries and header files. (See [Section 15.10, “Using High-Speed Interconnects with MySQL Cluster”](#) for more information about SCI.)

In addition, SCI requires specialized hardware.

It is strongly recommended to use SCI Transporters only for communication between `ndbd` processes. Note also that using SCI Transporters means that the `ndbd` processes never sleep. For this reason, SCI Transporters should be used only on machines having at least two CPUs dedicated for use by `ndbd` processes. There should be at least one CPU per `ndbd` process, with at least one CPU left in reserve to handle operating system activities.

•

### NodeId1, NodeId2

To identify a connection between two nodes it is necessary to provide node identifiers for each of them, as `NodeId1` and `NodeId2`.

•

### Host1SciId0

This identifies the SCI node ID on the first Cluster node (identified by `NodeId1`).

•

### Host1SciId1

It is possible to set up SCI Transporters for failover between two SCI cards which then should use separate networks between the nodes. This identifies the node ID and the second SCI card to be used on the first node.

•

### Host2SciId0

This identifies the SCI node ID on the second Cluster node (identified by `NodeId2`).

•

### Host2SciId1

When using two SCI cards to provide failover, this parameter identifies the second SCI card to be

used on the second node.

- **SharedBufferSize**  
Each SCI transporter has a shared memory segment used for communication between the two nodes. Setting the size of this segment to the default value of 1MB should be sufficient for most applications. Using a smaller value can lead to problems when performing many parallel inserts; if the shared buffer is too small, this can also result in a crash of the `ndbd` process.
- **SendLimit**  
A small buffer in front of the SCI media stores messages before transmitting them over the SCI network. By default, this is set to 8KB. Our benchmarks show that performance is best at 64KB but 16KB reaches within a few percent of this, and there was little if any advantage to increasing it beyond 8KB.
- **SendSignalId**  
To trace a distributed message it is necessary to identify each message uniquely. When this parameter is set to `Y`, message IDs are transported over the network. This feature is disabled by default.
- **Checksum**  
This parameter is a boolean value, and is disabled by default. When `Checksum` is enabled, checksums are calculated for all messages before they are placed in the send buffer. This feature prevents messages from being corrupted while waiting in the send buffer. It also serves as a check against data being corrupted during transport.

## 15.4.5. Overview of Cluster Configuration Parameters

The next three sections provide summary tables of MySQL Cluster configuration parameters used in the `config.ini` file to govern the cluster's functioning. Each table lists the parameters for one of the Cluster node process types (`ndbd`, `ndb_mgmd`, and `mysqld`), and includes the parameter's type as well as its default, minimum, and maximum values as applicable.

It is also stated what type of restart is required (node restart or system restart) — and whether the restart must be done with `--initial` — to change the value of a given configuration parameter. This information is provided in each table's **Restart Type** column, which contains one of the values shown in this list:

- **N**: Node Restart
- **IN**: Initial Node Restart
- **S**: System Restart
- **IS**: Initial System Restart

When performing a node restart or an initial node restart, all of the cluster's data nodes must be restarted in turn (also referred to as a *rolling restart*). It is possible to update cluster configuration parameters marked **N** or **IN** online — that is, without shutting down the cluster — in this fashion. An initial node re-

start requires restarting each `ndbd` process with the `--initial` option.

A system restart requires a complete shutdown and restart of the entire cluster. An initial system restart requires taking a backup of the cluster, wiping the cluster filesystem after shutdown, and then restoring from the backup following the restart.

In any cluster restart, all of the cluster's management servers must be restarted in order for them to read the updated configuration parameter values.

**Important:** Values for numeric cluster parameters can generally be increased without any problems, although it is advisable to do so progressively, making such adjustments in relatively small increments. However, decreasing the values of such parameters — particularly those relating to memory usage and disk space — is not to be undertaken lightly, and it is recommended that you do so only following careful planning and testing. In addition, it is generally the case that parameters relating to memory and disk usage which can be raised using a simple node restart require an initial node restart to be lowered.

Because some of these parameters can be used for configuring more than one type of cluster node, they may appear in more than one of the tables.

(Note that `4294967039` — which often appears as a maximum value in these tables — is equal to  $2^{32} - 2^8 - 1$ .)

### 15.4.5.1. Data Node Configuration Parameters

The following table provides information about parameters used in the `[NDBD]` or `[NDB_DEFAULT]` sections of a `config.ini` file for configuring MySQL Cluster data nodes. For detailed descriptions and other additional information about each of these parameters, see [Section 15.4.4.5, “Defining Data Nodes”](#).

#### *Restart Type Column Values*

- **N:** Node Restart
- **IN:** Initial Node Restart
- **S:** System Restart
- **IS:** Initial System Restart

See [Section 15.4.5, “Overview of Cluster Configuration Parameters”](#), for additional explanations of these abbreviations.

Parameter Name	Type/Units	Default Value	Minimum Value	Maximum Value	Restart Type
<code>ArbitrationTimeout</code>	milli-seconds	1000	10	4294967039	N
<code>BackupDataBufferSize</code>	bytes	2M	0	4294967039	N
<code>BackupDataDir</code>	string	FileSystem-Path/BACKUP	N/A	N/A	IN

MySQL Cluster

BackupLogBufferSize	bytes	2M	0	4294967039	N
BackupMemory	bytes	4M	0	4294967039	N
BackupWriteSize	bytes	32K	2K	4294967039	N
BackupMaxWriteSize	bytes	256K	2K	4294967039	N
BatchSizePerLocalScan	integer	64	1	992	N
DataDir	string	/var/lib/mysql-cluster	N/A	N/A	IN
DataMemory	bytes	80M	1M	1024G (subject to available system RAM and size of IndexMemory)	N
Diskless	true false	0	0	1	IS



	0)				
<code>ExecuteOnComputer</code>	integer				
<code>FileSystemPath</code>	string	value specified for <code>DataDir</code>	N/A	N/A	IN
<code>HeartbeatIntervalDbApi</code>	milliseconds	1500	100	4294967039	N
<code>HeartbeatIntervalDbDb</code>	milliseconds	1500	10	4294967039	N
<code>HostName</code>	string	<code>localhost</code>	N/A	N/A	S
<code>Id</code>	integer	<i>None</i>	1	63	N
<code>IndexMemory</code>	bytes	18M	1M	1024G (subject to available system RAM and size of <code>DataMemory</code> )	N
<code>LockPagesInMainMemory</code>	true false	0	0	1	N

	0)				
LogLevelCheckpoint	integer	0	0	15	IN
LogLevelConnection	integer	0	0	15	N
LogLevelError	integer	0	0	15	N
LogLevelInfo	integer	0	0	15	N
LogLevelNodeRestart	integer	0	0	15	N
LogLevelShutdown	integer	0	0	15	N
LogLevelStartup	integer	1	0	15	N
LogLevelStatistic	integer	0	0	15	N
LongMessageBuffer	bytes	1M	512K	4294967039	N
MaxNoOfAttributes	integer	1000	32	4294967039	N
MaxNoOfConcurrentIndexOperations	integer	8K	0	4294967039	N
MaxNoOfConcurrent-Operations	integer	32768	32	4294967039	N
MaxNoOfConcurrentScans	integer	256	2	500	N
MaxNoOfConcurrent-Transactions	integer	4096	32	4294967039	N
MaxNoOfFiredTriggers	integer	4000	0	4294967039	N
MaxNoOfIndexes (DEPRECATED — use MaxNoOfOrderedIndexes or MaxNoOfUniqueHashIndexes instead)	integer	128	0	4294967039	N

MaxNoOfLocalOperations	integer	UNDEFINED	32	4294967039	N
MaxNoOfLocalScans	integer	UNDEFINED	32	4294967039	N
MaxNoOfOrderedIndexes	integer	128	0	4294967039	N
MaxNoOfSavedMessages	integer	25	0	4294967039	N
MaxNoOfTables	integer	128	8	4294967039	N
MaxNoOfTriggers	integer	768	0	4294967039	N
MaxNoOfUnique-HashIndexes	integer	64	0	4294967039	N
NoOf-DiskPagesToDiskAfterRestartACC	integer (number of 8KB pages per 100 milliseconds)	20 (= 20 * 80KB = 1.6MB/second)	1	4294967039	N
NoOf-DiskPagesToDiskAfterRestartTUP	integer (number of 8KB pages per 100 milliseconds)	40 (= 40 * 80KB = 3.2MB/second)	1	4294967039	N

	seconds)				
NoOfDiskPagesToDisk- DuringRestartACC	integer (number of 8KB pages per 100 mil- li- sec onds)	20 (= 20 * 80KB = 1.6MB/second)	1	4294967039	N
NoOfDiskPagesToDisk- DuringRestartTUP	integer (number of 8KB pages per 100 mil- li- sec onds)	40 (= 40 * 80KB = 3.2MB/second)	1	4294967039	N
NoOfFragmentLogFiles	integer	8	1	4294967039	IN
NoOfReplicas	integer	<i>None</i>	1	4	IS
RedoBuffer	bytes	8M	1M	4294967039	N
RestartOnErrorInsert ( <i>DEBUG BUILDS ONLY</i> )	true   false	0	0	1	N

	0)				
<code>ServerPort</code> ( <i>OBSOLETE</i> )	integer	1186	0	4294967039	N
<code>StartFailureTimeout</code>	milliseconds	0	0	4294967039	N
<code>StartPartialTimeout</code>	milliseconds	30000	0	4294967039	N
<code>StartPartitioned-Timeout</code>	milliseconds	60000	0	4294967039	N
<code>StopOnError</code>	true false	1	0	1	N

	0)				
<a href="#">TimeBetweenGlobal-Checkpoints</a>	mil- li- sec ond s	2000	10	32000	N
<a href="#">TimeBetweenInactiveTransactionAbort-Check</a>	mil- li- sec ond s	1000	1000	4294967039	N
<a href="#">TimeBetweenLocal-Checkpoints</a>	in- tege r (nu mb er of 4-b yte wor ds as a bas e-2 log- arit hm)	20 (= $4 * 2^{20}$ = 4MB write opera- tions)	0	31	N
<a href="#">TimeBetweenWatchDogCheck</a>	mil- li- sec ond s	4000	70	4294967039	N
<a href="#">TransactionBuffer-Memory</a>	byt es	1M	1K	4294967039	N
<a href="#">TransactionDeadlock-DetectionTimeout</a>	mil- li- sec ond s	1200	50	4294967039	N
<a href="#">TransactionInactive-Timeout</a>	mil- li- sec ond s	0	0	4294967039	N
<a href="#">UndoDataBuffer</a>	byt es	16M	1M	4294967039	N
<a href="#">UndoIndexBuffer</a>	byt es	2M	1M	4294967039	N

### 15.4.5.2. Management Node Configuration Parameters

The following table provides information about parameters used in the [NDB\_MGMD] or [MGM] sections of a `config.ini` file for configuring MySQL Cluster management nodes. For detailed descriptions and other additional information about each of these parameters, see [Section 15.4.4.4, “Defining the Management Server”](#).

**Restart Type Column Values**

- **N**: Node Restart
- **IN**: Initial Node Restart
- **S**: System Restart
- **IS**: Initial System Restart

See [Section 15.4.5, “Overview of Cluster Configuration Parameters”](#), for additional explanations of these abbreviations.

Parameter Name	Type/Units	Default Value	Minimum Value	Maximum Value	Restart Type
<a href="#">ArbitrationDelay</a>	mil- li- sec onds	0	0	4294967039	N
<a href="#">ArbitrationRank</a>	in- tege r	1	0	2	N
<a href="#">DataDir</a>	stri ng	N/A	N/A	N/A	IN
<a href="#">ExecuteOnComputer</a>	in- tege r				
<a href="#">HostName</a>	stri ng	localhost	N/A	N/A	IN
<a href="#">Id</a>	in- tege r	None	1	63	IN
<a href="#">LogDestination</a>	CO NS OL E, SY SL OG, or FI LE	CONSOLE	N/A	N/A	N

### 15.4.5.3. SQL Node and API Node Configuration Parameters

The following table provides information about parameters used in the [SQL] and [API] sections of a `config.ini` file for configuring MySQL Cluster SQL nodes and API nodes. For detailed descriptions and other additional information about each of these parameters, see [Section 15.4.4.6, “Defining SQL and Other API Nodes”](#).

#### *Restart Type Column Values*

- **N**: Node Restart
- **IN**: Initial Node Restart
- **S**: System Restart
- **IS**: Initial System Restart

See [Section 15.4.5, “Overview of Cluster Configuration Parameters”](#), for additional explanations of these abbreviations.

Parameter Name	Type/Units	Default Value	Minimum Value	Maximum Value	Restart Type
<code>ArbitrationDelay</code>	milliseconds	0	0	4294967039	N
<code>ArbitrationRank</code>	integer	1	0	2	N
<code>BatchByteSize</code>	bytes	32K	1K	1M	N
<code>BatchSize</code>	integer	64	1	992	N
<code>ExecuteOnComputer</code>	integer				
<code>HostName</code>	string	<code>localhost</code>	N/A	N/A	IN
<code>Id</code>	integer	<i>None</i>	1	63	IN
<code>MaxScanBatchSize</code>	bytes	256K	32K	16M	N

### 15.4.6. Configuring Parameters for Local Checkpoints



The parameters discussed in [Logging and Checkpointing](#) and in [Data Memory, Index Memory, and String Memory](#) that are used to configure local checkpoints for a MySQL Cluster do not exist in isolation, but rather are very much interdependent on each other. In this section, we illustrate how these parameters — including `DataMemory`, `IndexMemory`, `NoOfDiskPagesToDiskAfterRestartTUP`, `NoOfDiskPagesToDiskAfterRestartACC`, and `NoOfFragmentLogFiles` — relate to one another in a working Cluster.

In this example, we assume that our application performs the following numbers of types of operations per hour:

- 50000 selects
- 15000 inserts
- 15000 updates
- 15000 deletes

We also make the following assumptions about the data used in the application:

- We are working with a single table having 40 columns.
- Each column can hold up to 32 bytes of data.
- A typical `UPDATE` run by the application affects the values of 5 columns.
- No `NULL` values are inserted by the application.

A good starting point is to determine the amount of time that should elapse between local checkpoints (LCPs). It worth noting that, in the event of a system restart, it takes 40-60 percent of this interval to execute the REDO log — for example, if the time between LCPs is 5 minutes (300 seconds), then it should take 2 to 3 minutes (120 to 180 seconds) for the REDO log to be read.

The maximum amount of data per node can be assumed to be the size of the `DataMemory` parameter. In this example, we assume that this is 2 GB. The `NoOfDiskPagesToDiskAfterRestartTUP` parameter represents the amount of data to be checkpointed per unit time — however, this parameter is actually expressed as the number of 8K memory pages to be checkpointed per 100 milliseconds. 2 GB per 300 seconds is approximately 6.8 MB per second, or 700 KB per 100 milliseconds, which works out to roughly 85 pages per 100 milliseconds.

Similarly, we can calculate `NoOfDiskPagesToDiskAfterRestartACC` in terms of the time for local checkpoints and the amount of memory required for indexes — that is, the `IndexMemory`. Assuming that we allow 512 MB for indexes, this works out to approximately 20 8-KB pages per 100 milliseconds for this parameter.

Next, we need to determine the number of REDO logfiles required — that is, fragment log files — the corresponding parameter being `NoOfFragmentLogFiles`. We need to make sure that there are sufficient REDO logfiles for keeping records for at least 3 local checkpoints. In a production setting, there are always uncertainties — for instance, we cannot be sure that disks always operate at top speed or with maximum throughput. For this reason, it is best to err on the side of caution, so we double our requirement and calculate a number of fragment logfiles which should be enough to keep records covering 6 local checkpoints.

It is also important to remember that the disk also handles writes to the REDO log and UNDO log, so if you find that the amount of data being written to disk as determined by the values of `NoOf-`

`DiskPagesToDiskAfterRestartACC` and `NoOfDiskPagesToDiskAfterRestartTUP` is approaching the amount of disk bandwidth available, you may wish to increase the time between local checkpoints.

Given 5 minutes (300 seconds) per local checkpoint, this means that we need to support writing log records at maximum speed for  $6 * 300 = 1800$  seconds. The size of a REDO log record is 72 bytes plus 4 bytes per updated column value plus the maximum size of the updated column, and there is one REDO log record for each table record updated in a transaction, on each node where the data reside. Using the numbers of operations set out previously in this section, we derive the following:

- 50000 select operations per hour yields 0 log records (and thus 0 bytes), since `SELECT` statements are not recorded in the REDO log.
- 15000 `DELETE` statements per hour is approximately 5 delete operations per second. (Since we wish to be conservative in our estimate, we round up here and in the following calculations.) No columns are updated by deletes, so these statements consume only 5 operations \* 72 bytes per operation = 360 bytes per second.
- 15000 `UPDATE` statements per hour is roughly the same as 5 updates per second. Each update uses 72 bytes, plus 4 bytes per column \* 5 columns updated, plus 32 bytes per column \* 5 columns — this works out to  $72 + 20 + 160 = 252$  bytes per operation, and multiplying this by 5 operation per second yields 1260 bytes per second.
- 15000 `INSERT` statements per hour is equivalent to 5 insert operations per second. Each insert requires REDO log space of 72 bytes, plus 4 bytes per record \* 40 columns, plus 32 bytes per column \* 40 columns, which is  $72 + 160 + 1280 = 1512$  bytes per operation. This times 5 operations per second yields 7560 bytes per second.

So the total number of REDO log bytes being written per second is approximately  $0 + 360 + 1260 + 7560 = 9180$  bytes. Multiplied by 1800 seconds, this yields 16524000 bytes required for REDO logging, or approximately 15.75 MB. The unit used for `NoOfFragmentLogFiles` represents a set of 4 16-MB logfiles — that is, 64 MB. Thus, the minimum value (3) for this parameter is sufficient for the scenario envisioned in this example, since 3 times 64 = 192 MB, or about 12 times what is required; the default value of 8 (or 512 MB) is more than ample in this case.

A copy of each altered table record is kept in the UNDO log. In the scenario discussed above, the UNDO log would not require any more space than what is provided by the default settings. However, given the size of disks, it is sensible to allocate at least 1 GB for it.

## 15.5. Upgrading and Downgrading MySQL Cluster

This portion of the MySQL Cluster chapter covers upgrading and downgrading a MySQL Cluster from one MySQL release to another. It discusses different types of Cluster upgrades and downgrades, and provides a Cluster upgrade/downgrade compatibility matrix (see [Section 15.5.2, “Cluster Upgrade and Downgrade Compatibility”](#)).

**Important:** You are expected already to be familiar with installing and configuring a MySQL Cluster prior to attempting an upgrade or downgrade. See [Section 15.4, “MySQL Cluster Configuration”](#).

This section remains in development, and continues to be updated and expanded.

### 15.5.1. Performing a Rolling Restart of the Cluster

This section discusses how to perform a *rolling restart* of a MySQL Cluster installation, so called because it involves stopping and starting (or restarting) each node in turn, so that the cluster itself remains operational. This is often done as part of a *rolling upgrade* or *rolling downgrade*, where high availability

of the cluster is mandatory and no downtime of the cluster as a whole is permissible. Where we refer to upgrades, the information provided here also generally applies to downgrades as well.

There are a number of reasons why a rolling restart might be desirable:

- **Cluster Configuration Change:** To make a change in the cluster's configuration, such as adding an SQL node to the cluster, or setting a configuration parameter to a new value.
- **Cluster Software Upgrade/Downgrade:** To upgrade the cluster to a newer version of the MySQL Cluster software (or to downgrade it to an older version). This is usually referred to as a “rolling upgrade” (or “rolling downgrade”, when reverting to an older version of MySQL Cluster).
- **Change on Node Host:** To make changes in the hardware or operating system on which one or more cluster nodes are running
- **Cluster Reset:** To reset the cluster because it has reached an undesirable state

The process for performing a rolling restart may be generalised as follows:

1. Stop, reconfigure, then restart each cluster management node (`ndb_mgmd` process) in turn
2. Stop, reconfigure, then restart each cluster data node (`ndbd` process) in turn
3. Stop, reconfigure, then restart each cluster SQL node (`mysqld` process) in turn

The specifics for implementing a particular rolling upgrade depend upon the actual changes being made. A more detailed view of the process is presented in the table shown here:

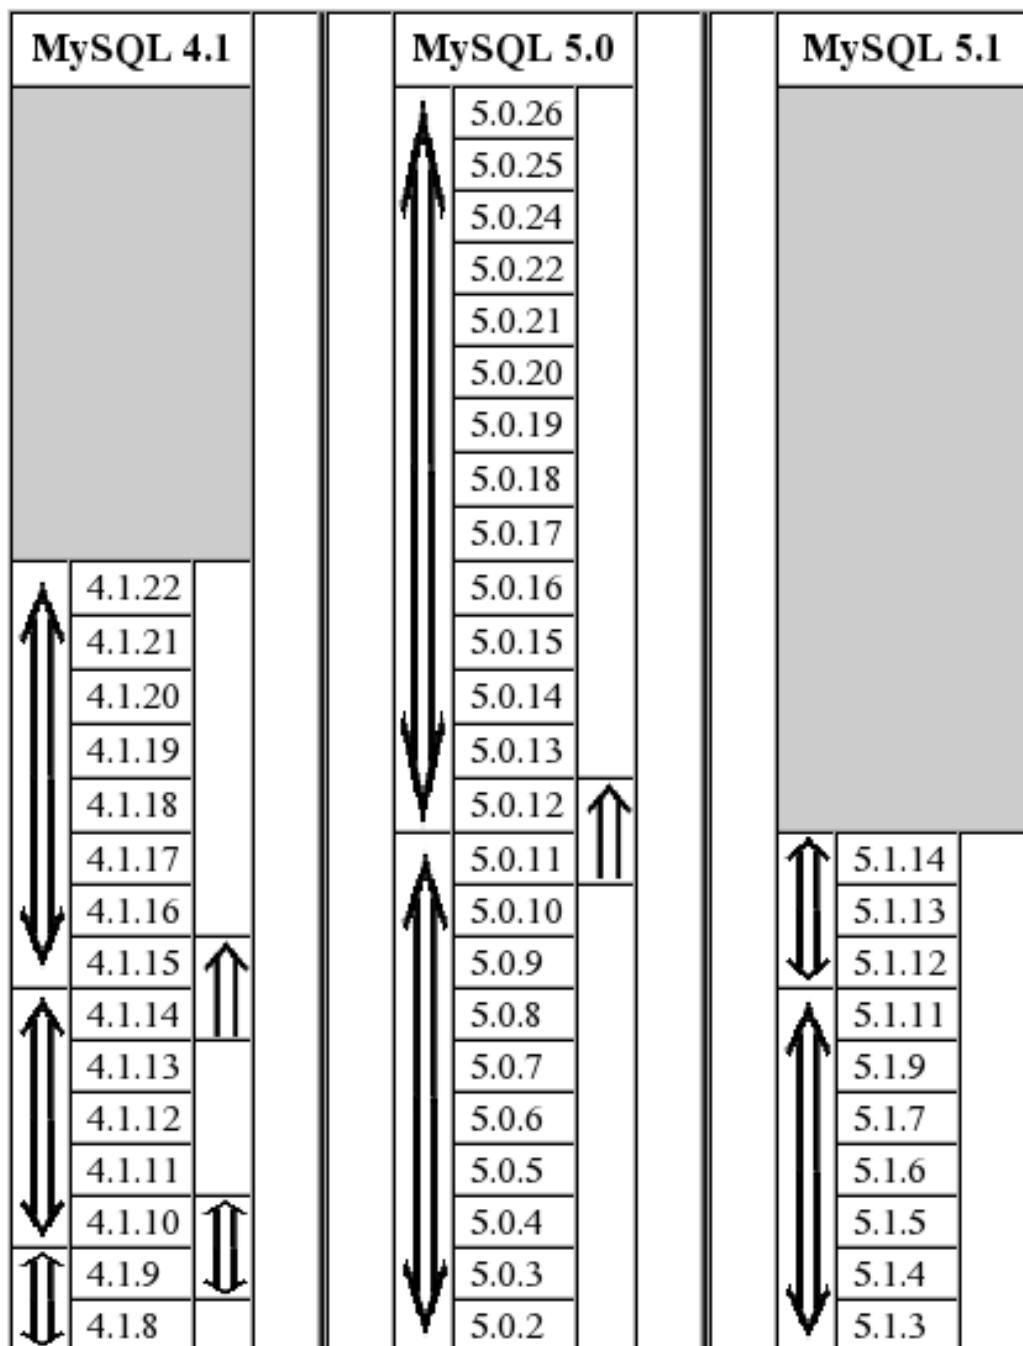
In the previous table, **Stop** and **Start** steps indicate that the process must be stopped completely using a shell command (such as `kill` on most Unix systems) or the management client `STOP` command, then started again from a system shell by invoking the `ndbd` or `ndb_mgmd` executable as appropriate. **Restart** indicates the process may be restarted using the `ndb_mgm` management client `RESTART` command.

## 15.5.2. Cluster Upgrade and Downgrade Compatibility

This section provides information regarding Cluster software and table file compatibility between differing versions of the MySQL Server for purposes of performing upgrades and downgrades.

**Important:** Only compatibility between MySQL versions with regard to `NDB Cluster` is taken into account in this section, and there are likely other issues to be considered. *As with any other MySQL software upgrade or downgrade, you are strongly encouraged to review the relevant portions of the MySQL Manual for the MySQL versions from which and to which you intend to migrate, before attempting an upgrade or downgrade of the MySQL Cluster software. See Section 2.11, “Upgrading MySQL”.*

The following table shows Cluster upgrade and downgrade compatibility between different versions of the MySQL Server.



**KEY:**

↑↓ Upgrades and downgrades supported

↑ Upgrade only

**Notes:**

- **4.1 Series:**

You cannot upgrade directly from 4.1.8 to 4.1.10 (or newer); you must first upgrade from 4.1.8 to 4.1.9, then upgrade to 4.1.10. Similarly, you cannot downgrade directly from 4.1.10 (or newer) to 4.1.8; you must first downgrade from 4.1.10 to 4.1.9, then downgrade from 4.1.9 to 4.1.8.

If you wish to upgrade a MySQL Cluster to 4.1.15, you must upgrade to 4.1.14 first, and you must upgrade to 4.1.15 before upgrading to 4.1.16 or newer.

Cluster downgrades from 4.1.15 to 4.1.14 (or earlier versions) are not supported.

Cluster upgrades from MySQL Server versions previous to 4.1.8 are not supported; when upgrading from these, you must dump all **NDB** tables, install the new version of the software, and then reload the tables from the dump.

- **5.0 Series:**

MySQL 5.0.2 was the first public release in this series.

Cluster downgrades from MySQL 5.0 to MySQL 4.1 are not supported.

Cluster downgrades from 5.0.12 to 5.0.11 (or earlier) are not supported.

You cannot restore with `ndb_restore` to a MySQL 5.0 Cluster using a backup made from a Cluster running MySQL 5.1. You must use `mysqldump` in such cases.

There was no public release for MySQL 5.0.23.

- **5.1 Series:**

MySQL 5.1.3 was the first public release in this series.

You cannot downgrade a MySQL 5.1.6 or later Cluster using Disk Data tables to MySQL 5.1.5 or earlier unless you convert all such tables to in-memory Cluster tables first.

MySQL 5.1.8 and MySQL 5.1.10 were not released.

Online cluster upgrades and downgrades between MySQL 5.1.11 (or an earlier version) and 5.1.12 (or a later version) are not possible due to major changes in the cluster filesystem. In such cases, you must perform a backup or dump, upgrade (or downgrade) the software, start each data node with `-initial`, and then restore from the backup or dump. You can use **NDB** backup/restore or `mysqldump` for this purpose.

## 15.6. Process Management in MySQL Cluster

Understanding how to manage MySQL Cluster requires a knowledge of four essential processes. In the next few sections of this chapter, we cover the roles played by these processes in a cluster, how to use them, and what startup options are available for each of them:

- [Section 15.6.1, “MySQL Server Process Usage for MySQL Cluster”](#)
- [Section 15.6.2, “ndbd, the Storage Engine Node Process”](#)
- [Section 15.6.3, “ndb\\_mgmd, the Management Server Process”](#)

- [Section 15.6.4, “ndb\\_mgm, the Management Client Process”](#)

## 15.6.1. MySQL Server Process Usage for MySQL Cluster

`mysqld` is the traditional MySQL server process. To be used with MySQL Cluster, `mysqld` needs to be built with support for the `NDB Cluster` storage engine, as it is in the precompiled `-max` binaries available from <http://dev.mysql.com/downloads/> for MySQL versions 4.1.3 and newer. If you build MySQL from source, you must invoke `configure` with the `--with-ndbcluster` option to enable `NDB Cluster` storage engine support.

If the `mysqld` binary has been built with Cluster support, the `NDB Cluster` storage engine is still disabled by default. You can use either of two possible options to enable this engine:

- Use `--ndbcluster` as a startup option on the command line when starting `mysqld`.
- Insert a line containing `ndbcluster` in the `[mysqld]` section of your `my.cnf` file.

An easy way to verify that your server is running with the `NDB Cluster` storage engine enabled is to issue the `SHOW ENGINES` statement in the MySQL Monitor (`mysql`). You should see the value `YES` as the `Support` value in the row for `NDBCLUSTER`. If you see `NO` in this row or if there is no such row displayed in the output, you are not running an `NDB-enabled` version of MySQL. If you see `DISABLED` in this row, you need to enable it in either one of the two ways just described.

To read cluster configuration data, the MySQL server requires at a minimum three pieces of information:

- The MySQL server's own cluster node ID
- The hostname or IP address for the management server (MGM node)
- The number of the TCP/IP port on which it can connect to the management server

Beginning with MySQL 4.1.5, node IDs can be dynamically allocated, in which case there is no need to specify them explicitly.

The `mysqld` parameter `ndb-connectstring` is used to specify the connectstring either on the command line when starting `mysqld` or in `my.cnf`. The connectstring contains the hostname or IP address where the management server can be found, as well as the TCP/IP port it uses.

In the following example, `ndb_mgmd.mysql.com` is the host where the management server resides, and the management server listens for cluster messages on port 1186:

```
shell> mysqld --ndb-connectstring=ndb_mgmd.mysql.com:1186
```

See [Section 15.4.4.2, “The Cluster connectstring”](#), for more information on connectstrings.

Given this information, the MySQL server will be a full participant in the cluster. (We sometimes refer to a `mysqld` process running in this manner as an SQL node.) It will be fully aware of all cluster data nodes as well as their status, and will establish connections to all data nodes. In this case, it is able to use any data node as a transaction coordinator and to read and update node data.

You can see in the `mysql` client whether a MySQL server is connected to the cluster using `SHOW PROCESSLIST`. If the MySQL server is connected to the cluster, and you have the `PROCESS` privilege,

then the first row of the output is as shown here:

```
mysql> SHOW PROCESSLIST \G
***** 1. row *****
  Id: 1
  User: system user
  Host:
  db:
  Command: Daemon
  Time: 1
  State: Waiting for event from ndbcluster
  Info: NULL
```

## 15.6.2. `ndbd`, the Storage Engine Node Process

`ndbd` is the process that is used to handle all the data in tables using the NDB Cluster storage engine. This is the process that empowers a data node to accomplish distributed transaction handling, node recovery, checkpointing to disk, online backup, and related tasks.

In a MySQL Cluster, a set of `ndbd` processes cooperate in handling data. These processes can execute on the same computer (host) or on different computers. The correspondences between data nodes and Cluster hosts is completely configurable.

In MySQL versions prior to 4.1.5, each `ndbd` process should be started in a separate directory, the reason for this being that `ndbd` generated a set of log files in its starting directory. In MySQL 4.1.5, this behavior was changed such that these files are placed in the directory specified by `DataDir` in the configuration file. Thus `ndbd` can be started from anywhere.

These log files are listed below. `node_id` is the node's unique identifier. Note that `node_id` represents the node's unique identifier. For example, `ndb_2_error.log` is the error log generated by the data node whose node ID is 2.

- `ndb_node_id_error.log` (was `error.log` in version 4.1.3) is a file containing records of all crashes which the referenced `ndbd` process has encountered. Each record in this file contains a brief error string and a reference to a trace file for this crash. A typical entry in this file might appear as shown here:

```
Date/Time: Saturday 30 July 2004 - 00:20:01
Type of error: error
Message: Internal program error (failed ndbrequire)
Fault ID: 2341
Problem data: DbtupFixAlloc.cpp
Object of reference: DBTUP (Line: 173)
ProgramName: NDB Kernel
ProcessID: 14909
TraceFile: ndb_2_trace.log.2
***EOM***
```

**Note:** *It is very important to be aware that the last entry in the error log file is not necessarily the newest one* (nor is it likely to be). Entries in the error log are *not* listed in chronological order; rather, they correspond to the order of the trace files as determined in the `ndb_node_id_trace.log.next` file (see below). Error log entries are thus overwritten in a cyclical and not sequential fashion.

- `ndb_node_id_trace.log.trace_id` (was `NDB_TraceFile_trace_id.trace` in version 4.1.3) is a trace file describing exactly what happened just before the error occurred. This information is useful for analysis by the MySQL Cluster development team.

It is possible to configure the number of these trace files that will be created before old files are overwritten. `trace_id` is a number which is incremented for each successive trace file.

- `ndb_node_id_trace.log.next` (was `NextTraceFileNo.log` in version 4.1.3) is the file that keeps track of the next trace file number to be assigned.
- `ndb_node_id_out.log` is a file containing any data output by the `ndbd` process. This file is created only if `ndbd` is started as a daemon, which is the default behavior beginning with MySQL 4.1.5. This file was named `nodenode_id.out` in versions 4.1.3 and 4.1.4.
- `ndb_node_id.pid` is a file containing the process ID of the `ndbd` process when started as a daemon. It also functions as a lock file to avoid the starting of nodes with the same identifier.
- `ndb_node_id_signal.log` (was `Signal.log` in version 4.1.3) is a file used only in debug versions of `ndbd`, where it is possible to trace all incoming, outgoing, and internal messages with their data in the `ndbd` process.

It is recommended not to use a directory mounted through NFS because in some environments this can cause problems whereby the lock on the `.pid` file remains in effect even after the process has terminated.

To start `ndbd`, it may also be necessary to specify the hostname of the management server and the port on which it is listening. Optionally, one may also specify the node ID that the process is to use.

```
shell> ndbd --connect-string="nodeid=2;host=ndb_mgmd.mysql.com:1186"
```

See [Section 15.4.4.2, “The Cluster `connectstring`”](#), for additional information about this issue. [Section 15.6.5, “Command Options for MySQL Cluster Processes”](#), describes other options for `ndbd`.

When `ndbd` starts, it actually initiates two processes. The first of these is called the “angel process”; its only job is to discover when the execution process has been completed, and then to restart the `ndbd` process if it is configured to do so. Thus, if you attempt to kill `ndbd` via the Unix `kill` command, it is necessary to kill both processes, beginning with the angel process. The preferred method of terminating an `ndbd` process is to use the management client and stop the process from there.

The execution process uses one thread for reading, writing, and scanning data, as well as all other activities. This thread is implemented asynchronously so that it can easily handle thousands of concurrent activities. In addition, a watch-dog thread supervises the execution thread to make sure that it does not hang in an endless loop. A pool of threads handles file I/O, with each thread able to handle one open file. Threads can also be used for transporter connections by the transporters in the `ndbd` process. In a multi-processor system performing a large number of operations (including updates), the `ndbd` process can consume up to 2 CPUs if permitted to do so.

For a machine with many CPUs it is possible to use several `ndbd` processes which belong to different node groups; however, such a configuration is still considered experimental and is not supported for MySQL 4.1 in a production setting. See [Section 15.11, “Known Limitations of MySQL Cluster”](#).

### 15.6.3. `ndb_mgmd`, the Management Server Process

The management server is the process that reads the cluster configuration file and distributes this information to all nodes in the cluster that request it. It also maintains a log of cluster activities. Management clients can connect to the management server and check the cluster's status.

As of MySQL 4.1.5, it is no longer necessary to specify a connectstring when starting the management server. However, if you are using more than one management server, a connectstring should be provided and each node in the cluster should specify its node ID explicitly.

See [Section 15.4.4.2, “The Cluster `connectstring`”](#), for information about using connectstrings. [Section 15.6.5, “Command Options for MySQL Cluster Processes”](#), describes other options for `ndb_mgmd`.



The following files are created or used by `ndb_mgmd` in its starting directory. From MySQL 4.1.5, the log and PID files are placed in the `DataDir` as specified in the `config.ini` configuration file. In the list that follows, `node_id` is the unique node identifier.

- `config.ini` is the configuration file for the cluster as a whole. This file is created by the user and read by the management server. [Section 15.4, “MySQL Cluster Configuration”](#), discusses how to set up this file.
- `ndb_node_id_cluster.log` (was `cluster.log` in version 4.1.3) is the cluster events log file. Examples of such events include checkpoint startup and completion, node startup events, node failures, and levels of memory usage. A complete listing of cluster events with descriptions may be found in [Section 15.7, “Management of MySQL Cluster”](#).

When the size of the cluster log reaches one million bytes, the file is renamed to `ndb_node_id_cluster.log.seq_id` (was `cluster.log.seq_id` in version 4.1.3) where `seq_id` is the sequence number of the cluster log file. (For example: If files with the sequence numbers 1, 2, and 3 already exist, the next log file is named using the number 4.)

- `ndb_node_id_out.log` (was `node1.out` in version 4.1.3) is the file used for `stdout` and `stderr` when running the management server as a daemon.
- `ndb_node_id.pid` (was `nodenode_id.pid` in version 4.1.3) is the process ID file used when running the management server as a daemon.

## 15.6.4. `ndb_mgm`, the Management Client Process

The management client process is actually not needed to run the cluster. Its value lies in providing a set of commands for checking the cluster's status, starting backups, and performing other administrative functions. The management client accesses the management server using a C API. Advanced users can also employ this API for programming dedicated management processes to perform tasks similar to those performed by `ndb_mgm`.

To start the management client, it is necessary to supply the hostname and port number of the management server:

```
shell> ndb_mgm [host_name [port_num]]
```

For example:

```
shell> ndb_mgm ndb_mgmd.mysql.com 1186
```

The default hostname and port number are `localhost` and 1186, respectively. (Prior to MySQL 4.1.8, the default Cluster port was 2200.)

Additional information about using `ndb_mgm` can be found in [Section 15.6.5.4, “Command Options for `ndb\_mgm`”](#), and [Section 15.7.2, “Commands in the Management Client”](#).

## 15.6.5. Command Options for MySQL Cluster Processes

All MySQL Cluster executables (except for `mysqld`) take the options described in this section as of MySQL 4.1.8. Users of earlier MySQL Cluster versions should note that some of these options have been changed to make consistent with one another as well as with `mysqld`. You can use the `--help` option with any MySQL Cluster executable to view a list of the options which it supports.

The following options are common to all MySQL Cluster executables:

- `--help --usage, -?`

Prints a short list with descriptions of the available command options.

- `--connect-string=connect_string`

`connect_string` sets the connectstring to the management server as a command option. Available with `ndb_mgm` from MySQL 4.1.8 onward.

```
shell> ndbd --connect-string="nodeid=2;host=ndb_mgmd.mysql.com:1186"
```

- `--debug[=options]`

This option can be used only for versions compiled with debugging enabled. It is used to enable output from debug calls in the same manner as for the `mysqld` process.

- `--execute=command, -e command`

Can be used to send a command to a Cluster executable from the system shell. For example, either of the following:

```
shell> ndb_mgm -e "SHOW"
```

or

```
shell> ndb_mgm --execute="SHOW"
```

is equivalent to

```
NDB> SHOW;
```

This is analogous to how the `--execute` or `-e` option works with the `mysql` command-line client. See [Section 4.3.1, “Using Options on the Command Line”](#).

- `--version, -V`

Prints the MySQL Cluster version number of the executable. The version number is relevant because not all versions can be used together, and the MySQL Cluster startup process verifies that the versions of the binaries being used can co-exist in the same cluster. This is also important when performing an online (rolling) software upgrade or downgrade of MySQL Cluster. (See [Section 15.5.1, “Performing a Rolling Restart of the Cluster”](#)).

The next few sections describe options specific to individual NDB programs.

### 15.6.5.1. MySQL Cluster-Related Command Options for `mysqld`

- `--ndb-connectstring=connect_string`

When using the `NDB Cluster` storage engine, this option specifies the management server that distributes cluster configuration data.

- `--ndbcluster`

The `NDB Cluster` storage engine is necessary for using MySQL Cluster. If a `mysqld` binary includes support for the `NDB Cluster` storage engine, the engine is disabled by default. Use the `--ndbcluster` option to enable it. Use `--skip-ndbcluster` to explicitly disable the engine.

### 15.6.5.2. Command Options for `ndbd`

For options common to all `NDB` programs, see [Section 15.6.5, “Command Options for MySQL Cluster Processes”](#).

- `--daemon, -d`

Instructs `ndbd` to execute as a daemon process. From MySQL 4.1.5 on, this is the default behavior, and `--nodaemon` can be used to prevent the process from running as a daemon.

- `--initial`

Instructs `ndbd` to perform an initial start. An initial start erases any files created for recovery purposes by earlier instances of `ndbd`. It also re-creates recovery log files. Note that on some operating systems this process can take a substantial amount of time.

An `--initial` start is to be used only the very first time that the `ndbd` process is started because it removes all files from the Cluster filesystem and re-creates all REDO log files. The exceptions to this rule are:

- When performing a software upgrade which has changed the contents of any files.
- When restarting the node with a new version of `ndbd`.
- As a measure of last resort when for some reason the node restart or system restart repeatedly fails. In this case, be aware that this node can no longer be used to restore data due to the destruction of the datafiles.

This option does not affect any backup files that have already been created by the affected node.

It is possible to achieve the same effect by deleting by other means (such as using `rm -r -f`) all files and directories in the data node's `DataDir` — with the possible exception of the `BACKUP` directory in `DataDir`, should you wish to retain any backups that have been created on that data node — and then starting `ndbd` without having to use the `--initial` option. This may be useful when scripting Cluster administrative tasks.

In older versions of MySQL Cluster, it was possible to use `-i` for this option. This shortcut was removed to prevent this option from being used by mistake.

- `--nodaemon`

Instructs `ndbd` not to start as a daemon process. This is useful when `ndbd` is being debugged and you want output to be redirected to the screen.

- `--nostart, -n`

Instructs `ndbd` not to start automatically. When this option is used, `ndbd` connects to the management server, obtains configuration data from it, and initializes communication objects. However, it does not actually start the execution engine until specifically requested to do so by the management server. This can be accomplished by issuing the proper `START` command in the management client (see [Section 15.7.2, “Commands in the Management Client”](#)).

### 15.6.5.3. Command Options for `ndb_mgmd`

For options common to NDB programs, see [Section 15.6.5, “Command Options for MySQL Cluster Processes”](#).

- `--config-file=filename, -c filename, -f filename`

Instructs the management server as to which file it should use for its configuration file. This option must be specified. The filename defaults to `config.ini`. The `-f` shortcut is available beginning with MySQL Cluster 4.1.8.

- `--daemon, -d`

Instructs `ndb_mgmd` to start as a daemon process. This is the default behavior.

- `--nodaemon`

Instructs `ndb_mgmd` not to start as a daemon process.

### 15.6.5.4. Command Options for `ndb_mgm`

For options common to NDB programs, see [Section 15.6.5, “Command Options for MySQL Cluster Processes”](#).

- `--try-reconnect=number`

If the connection to the management server is broken, the node tries to reconnect to it every 5 seconds until it succeeds. By using this option, it is possible to limit the number of attempts to `number` before giving up and reporting an error instead.

## 15.7. Management of MySQL Cluster

Managing a MySQL Cluster involves a number of tasks, the first of which is to configure and start MySQL Cluster. This is covered in [Section 15.4, “MySQL Cluster Configuration”](#), and [Section 15.6, “Process Management in MySQL Cluster”](#).

The following sections cover the management of a running MySQL Cluster.

There are essentially two methods of actively managing a running MySQL Cluster. The first of these is through the use of commands entered into the management client whereby cluster status can be checked, log levels changed, backups started and stopped, and nodes stopped and started. The second method involves studying the contents of the cluster log `ndb_node_id_cluster.log`; this is usually found in the management server's `DataDir` directory, but this location can be overridden using the `LogDestination` option — see [Section 15.4.4.4, “Defining the Management Server”](#), for details. (Recall that `node_id` represents the unique identifier of the node whose activity is being logged.) The cluster log contains event reports generated by `ndbd`. It is also possible to send cluster log entries to a Unix system log.

### 15.7.1. MySQL Cluster Startup Phases

This section describes the steps involved when the cluster is started.

There are several different startup types and modes, as shown here:

- **Initial Start:** The cluster starts with a clean filesystem on all data nodes. This occurs either when the cluster started for the very first time, or when it is restarted using the `--initial` option.
- **System Restart:** The cluster starts and reads data stored in the data nodes. This occurs when the cluster has been shut down after having been in use, when it is desired for the cluster to resume operations from the point where it left off.
- **Node Restart:** This is the online restart of a cluster node while the cluster itself is running.
- **Initial Node Restart:** This is the same as a node restart, except that the node is reinitialized and started with a clean filesystem.

Prior to startup, each data node (`ndbd` process) must be initialized. Initialization consists of the following steps:

1. Obtain a Node ID.
2. Fetch configuration data.
3. Allocate ports to be used for inter-node communications.
4. Allocate memory according to settings obtained from the configuration file.

When a data node or SQL node first connects to the management node, it reserves a cluster node ID. To make sure that no other node allocates the same node ID, this ID is retained until the node has managed to connect to the cluster and at least one `ndbd` reports that this node is connected. This retention of the node ID is guarded by the connection between the node in question and `ndb_mgmd`.

Normally, in the event of a problem with the node, the node disconnects from the management server, the socket used for the connection is closed, and the reserved node ID is freed. However, if a node is disconnected abruptly — for example, due to a hardware failure in one of the cluster hosts, or because of network issues — the normal closing of the socket by the operating system may not take place. In this case, the node ID continues to be reserved and not released until a TCP timeout occurs 10 or so minutes later.

To take care of this problem, you can use `PURGE STALE SESSIONS`. Running this statement forces all reserved node IDs to be checked; any that are not being used by nodes actually connected to the cluster are then freed.

After each data node has been initialized, the cluster startup process can proceed. The stages which the cluster goes through during this process are listed here:

- **Stage 0**

Clear the cluster filesystem. This stage occurs *only* if the cluster was started with the `--initial` option.

- **Stage 1**

This stage sets up Cluster connections, establishes inter-node communications, and starts Cluster heartbeats.

**Note:** When one or more nodes hang in Phase 1 while the remaining node or nodes hang in Phase 2, this often indicates network problems. One possible cause of such issues is one or more cluster hosts having multiple network interfaces. Another common source of problems causing this condition is

the blocking of TCP/IP ports needed for communications between cluster nodes. In the latter case, this is often due to a misconfigured firewall.

- **Stage 2**

The arbitrator node is elected. If this is a system restart, the cluster determines the latest restorable global checkpoint.

- **Stage 3**

This stage initializes a number of internal cluster variables.

- **Stage 4**

For an initial start or initial node restart, the redo log files are created. The number of these files is equal to `NoOfFragmentLogFiles`.

For a system restart:

- Read schema or schemas.
- Read data from the local checkpoint and undo logs.
- Apply all redo information until the latest restorable global checkpoint has been reached.

For a node restart, find the tail of the redo log.

- **Stage 5**

If this is an initial start, create the `SYSTAB_0` and `NDB$EVENTS` internal system tables.

For a node restart or an initial node restart:

1. The node is included in transaction handling operations.
2. The node schema is compared with that of the master and synchronized with it.
3. Synchronize data received in the form of `INSERT` from the other data nodes in this node's node group.
4. In all cases, wait for complete local checkpoint as determined by the arbitrator.

- **Stage 6**

Update internal variables.

- **Stage 7**

Update internal variables.

- **Stage 8**

In a system restart, rebuild all indexes.

- **Stage 9**

Update internal variables.

- **Stage 10**

At this point in a node restart or initial node restart, APIs may connect to the node and begin to receive events.

- **Stage 11**

At this point in a node restart or initial node restart, event delivery is handed over to the node joining the cluster. The newly-joined node takes over responsibility for delivering its primary data to subscribers.

After this process is completed for an initial start or system restart, transaction handling is enabled. For a node restart or initial node restart, completion of the startup process means that the node may now act as a transaction coordinator.

## 15.7.2. Commands in the Management Client

In addition to the central configuration file, a cluster may also be controlled through a command-line interface available through the management client `ndb_mgm`. This is the primary administrative interface to a running cluster.

Commands for the event logs are given in [Section 15.7.3, “Event Reports Generated in MySQL Cluster”](#); commands for creating backups and restoring from backup are provided in [Section 15.8, “On-line Backup of MySQL Cluster”](#).

The management client has the following basic commands. In the listing that follows, *node\_id* denotes either a database node ID or the keyword `ALL`, which indicates that the command should be applied to all of the cluster's data nodes.

- `HELP`

Displays information on all available commands.

- `SHOW`

Displays information on the cluster's status.

**Note:** In a cluster where multiple management nodes are in use, this command displays information only for data nodes that are actually connected to the current management server.

- `node_id START`

Brings online the data node identified by *node\_id* (or all data nodes).

**Important:** To use this command to bring a data node online, the data node must have been started using `ndbd --nostart` or `ndbd -n`.

- `node_id STOP`

Stops the data node identified by *node\_id* (or all data nodes). A data node affected by this command disconnects from the cluster, and its associated `ndbd` process terminates.

- `node_id RESTART [-n] [-i]`

Restarts the data node identified by `node_id` (or all data nodes).

Using the `-i` option with `RESTART` causes the data node to perform an initial restart; that is, the node's filesystem is deleted and recreated. The effect is the same as that obtained from stopping the data node process and then starting it again using `ndbd --initial` from the system shell.

Using the `-n` option causes the data node process to be restarted, but the data node is not actually brought online until the appropriate `START` command is issued. The effect of this option is the same as that obtained from stopping the data node and then starting it again using `ndbd --nostart` or `ndbd -n` from the system shell.

- `node_id STATUS`  
Displays status information for the data node identified by `node_id` (or for all data nodes).
- `ENTER SINGLE USER MODE node_id`  
Enters single-user mode, whereby only the MySQL server identified by the node ID `node_id` is allowed to access the database.
- `EXIT SINGLE USER MODE`  
Exits single-user mode, allowing all SQL nodes (that is, all running `mysqld` processes) to access the database.
- `QUIT, EXIT`  
Terminates the management client.  
This command does not affect any nodes connected to the cluster.
- `SHUTDOWN`  
Shuts down all cluster data nodes and management nodes. To exit the management client after this has been done, use `EXIT` or `QUIT`.  
This command does *not* shut down any SQL nodes or API nodes that are connected to the cluster.

### 15.7.3. Event Reports Generated in MySQL Cluster

In this section, we discuss the types of event logs provided by MySQL Cluster, and the types of events that are logged.

MySQL Cluster provides two types of event log. These are the *cluster log*, which includes events generated by all cluster nodes, and the *node logs*, which are local to each data node.

Output generated by cluster event logging can have multiple destinations including a file, the management server console window, or `syslog`. Output generated by node event logging is written to the data node's console window.

Both types of event logs can be set to log different subsets of events.



**Note:** The cluster log is the log recommended for most uses because it provides logging information for an entire cluster in a single file. Node logs are intended to be used only during application development, or for debugging application code.

Each reportable event can be distinguished according to three different criteria:

- *Category:* This can be any one of the following values: `STARTUP`, `SHUTDOWN`, `STATISTICS`, `CHECKPOINT`, `NODERESTART`, `CONNECTION`, `ERROR`, or `INFO`.
- *Priority:* This is represented by one of the numbers from 1 to 15 inclusive, where 1 indicates “most important” and 15 “least important.”
- *Severity Level:* This can be any one of the following values: `ALERT`, `CRITICAL`, `ERROR`, `WARNING`, `INFO`, or `DEBUG`.

Both the cluster log and the node log can be filtered on these properties.

### 15.7.3.1. Logging Management Commands

The following management commands are related to the cluster log:

- `CLUSTERLOG ON`  
Turns the cluster log on.
- `CLUSTERLOG OFF`  
Turns the cluster log off.
- `CLUSTERLOG INFO`  
Provides information about cluster log settings.
- `node_id CLUSTERLOG category=threshold`  
Logs *category* events with priority less than or equal to *threshold* in the cluster log.
- `CLUSTERLOG FILTER severity_level`  
Toggles cluster logging of events of the specified *severity\_level*.

The following table describes the default setting (for all data nodes) of the cluster log category threshold. If an event has a priority with a value lower than or equal to the priority threshold, it is reported in the cluster log.

Note that events are reported per data node, and that the threshold can be set to different values on different nodes.

Category	Default threshold (All data nodes)
<code>STARTUP</code>	7
<code>SHUTDOWN</code>	7
<code>STATISTICS</code>	7
<code>CHECKPOINT</code>	7

NODERESTART	7
CONNECTION	7
ERROR	15
INFO	7

Thresholds are used to filter events within each category. For example, a `STARTUP` event with a priority of 3 is not logged unless the threshold for `STARTUP` is changed to 3 or lower. Only events with priority 3 or lower are sent if the threshold is 3.

The following table shows the event severity levels. (**Note:** These correspond to Unix `syslog` levels, except for `LOG_EMERG` and `LOG_NOTICE`, which are not used or mapped.)

1	ALERT	A condition that should be corrected immediately, such as a corrupted system database
2	CRITICAL	Critical conditions, such as device errors or insufficient resources
3	ERROR	Conditions that should be corrected, such as configuration errors
4	WARNING	Conditions that are not errors, but that might require special handling
5	INFO	Informational messages
6	DEBUG	Debugging messages used for NDB Cluster development

Event severity levels can be turned on or off (using `CLUSTERLOG FILTER` — see above). If a severity level is turned on, then all events with a priority less than or equal to the category thresholds are logged. If the severity level is turned off then no events belonging to that severity level are logged.

### 15.7.3.2. Log Events

An event report reported in the event logs has the following format:

```
datetime [string] severity -- message
```

For example:

```
09:19:30 2005-07-24 [NDB] INFO -- Node 4 Start phase 4 completed
```

This section discusses all reportable events, ordered by category and severity level within each category.

In the event descriptions, GCP and LCP mean “Global Checkpoint” and “Local Checkpoint,” respectively.

#### CONNECTION Events

These events are associated with connections between Cluster nodes.

Event	Priority	Severity Level	Description
data nodes connected	8	INFO	Data nodes connected
data nodes disconnected	8	INFO	Data nodes disconnected
Communication closed	8	INFO	SQL node or data node connection closed
Communication opened	8	INFO	SQL node or data node connection opened

**CHECKPOINT Events**

The logging messages shown here are associated with checkpoints.

Event	Priority	Severity Level	Description
LCP stopped in calc keep GCI	0	ALERT	LCP stopped
Local checkpoint fragment completed	11	INFO	LCP on a fragment has been completed
Global checkpoint completed	10	INFO	GCP finished
Global checkpoint started	9	INFO	Start of GCP: REDO log is written to disk
Local checkpoint completed	8	INFO	LCP completed normally
Local checkpoint started	7	INFO	Start of LCP: data written to disk
Report undo log blocked	7	INFO	UNDO logging blocked; buffer near overflow

**STARTUP Events**

The following events are generated in response to the startup of a node or of the cluster and of its success or failure. They also provide information relating to the progress of the startup process, including information concerning logging activities.

Event	Priority	Severity Level	Description
Internal start signal received ST-TORRY	15	INFO	Blocks received after completion of restart
Undo records executed	15	INFO	
New REDO log started	10	INFO	GCI keep <i>X</i> , newest restorable GCI <i>Y</i>
New log started	10	INFO	Log part <i>X</i> , start MB <i>Y</i> , stop MB <i>Z</i>
Node has been refused for inclusion in the cluster	8	INFO	Node cannot be included in cluster due to misconfiguration, inability to establish communication, or other problem
data node neighbors	8	INFO	Shows neighboring data nodes
data node start phase <i>X</i> completed	4	INFO	A data node start phase has been completed
Node has been successfully included into the cluster	3	INFO	Displays the node, managing node, and dynamic ID
data node start phases initiated	1	INFO	NDB Cluster nodes starting
data node all start phases completed	1	INFO	NDB Cluster nodes started
data node shutdown initiated	1	INFO	Shutdown of data node has commenced
data node shutdown aborted	1	INFO	Unable to shut down data node normally

**NODERESTART Events**

The following events are generated when restarting a node and relate to the success or failure of the node restart process.

Event	Priority	Severity Level	Description
-------	----------	----------------	-------------

Node failure phase completed	8	ALERT	Reports completion of node failure phases
Node has failed, node state was <i>X</i>	8	ALERT	Reports that a node has failed
Report arbitrator results	2	ALERT	<p>There are eight different possible results for arbitration attempts:</p> <ul style="list-style-type: none"> <li>• Arbitration check failed — less than 1/2 nodes left</li> <li>• Arbitration check succeeded — node group majority</li> <li>• Arbitration check failed — missing node group</li> <li>• Network partitioning — arbitration required</li> <li>• Arbitration succeeded — affirmative response from node <i>X</i></li> <li>• Arbitration failed - negative response from node <i>X</i></li> <li>• Network partitioning - no arbitrator available</li> <li>• Network partitioning - no arbitrator configured</li> </ul>
Completed copying a fragment	10	INFO	
Completed copying of dictionary information	8	INFO	
Completed copying distribution information	8	INFO	
Starting to copy fragments	8	INFO	
Completed copying all fragments	8	INFO	
GCP takeover started	7	INFO	
GCP takeover completed	7	INFO	
LCP takeover started	7	INFO	
LCP takeover completed (state = <i>X</i> )	7	INFO	
Report whether an arbitrator is found or not	6	INFO	<p>There are seven different possible outcomes when seeking an arbitrator:</p> <ul style="list-style-type: none"> <li>• Management server restarts arbitration thread [state=<i>X</i>]</li> <li>• Prepare arbitrator node <i>X</i> [ticket=<i>Y</i>]</li> <li>• Receive arbitrator node <i>X</i> [ticket=<i>Y</i>]</li> <li>• Started arbitrator node <i>X</i> [ticket=<i>Y</i>]</li> </ul>

			<ul style="list-style-type: none"> <li>• Lost arbitrator node <i>X</i> - process failure [state=<i>Y</i>]</li> <li>• Lost arbitrator node <i>X</i> - process exit [state=<i>Y</i>]</li> <li>• Lost arbitrator node <i>X</i> &lt;error msg&gt; [state=<i>Y</i>]</li> </ul>
--	--	--	---

**STATISTICS Events**

The following events are of a statistical nature. They provide information such as numbers of transactions and other operations, amount of data sent or received by individual nodes, and memory usage.

Event	Priority	Severity Level	Description
Report job scheduling statistics	9	INFO	Mean internal job scheduling statistics
Sent number of bytes	9	INFO	Mean number of bytes sent to node <i>X</i>
Received # of bytes	9	INFO	Mean number of bytes received from node <i>X</i>
Report transaction statistics	8	INFO	Numbers of: transactions, commits, reads, simple reads, writes, concurrent operations, attribute information, and aborts
Report operations	8	INFO	Number of operations
Report table create	7	INFO	
Memory usage	5	INFO	Data and index memory usage (80%, 90%, and 100%)

**ERROR Events**

These events relate to Cluster errors and warnings. The presence of one or more of these generally indicates that a major malfunction or failure has occurred.

Event	Priority	Severity	Description
Dead due to missed heartbeat	8	ALERT	Node <i>X</i> declared “dead” due to missed heartbeat
Transporter errors	2	ERROR	
Transporter warnings	8	WARNING	
Missed heartbeats	8	WARNING	Node <i>X</i> missed heartbeat # <i>Y</i>
General warning events	2	WARNING	

**INFO Events**

These events provide general information about the state of the cluster and activities associated with Cluster maintenance, such as logging and heartbeat transmission.

Event	Priority	Severity	Description
Sent heartbeat	12	INFO	Heartbeat sent to node X
Create log bytes	11	INFO	Log part, log file, MB
General information events	2	INFO	

### 15.7.3.3. Using CLUSTERLOG STATISTICS

The NDB management client's `CLUSTERLOG STATISTICS` command can provide a number of useful statistics in its output. The following statistics are reported by the transaction coordinator:

Statistic	Description (Number of...)
<code>Trans. Count</code>	Transactions attempted with this node as coordinator
<code>Commit Count</code>	Transactions committed with this node as coordinator
<code>Read Count</code>	Primary key reads (all)
<code>Simple Read Count</code>	Primary key reads reading the latest committed value
<code>Write Count</code>	Primary key writes (includes all <code>INSERT</code> , <code>UPDATE</code> , and <code>DELETE</code> operations)
<code>AttrInfoCount</code>	Data words used to describe all reads and writes received
<code>Concurrent Operations</code>	All concurrent operations ongoing at the moment the report is taken
<code>Abort Count</code>	Transactions with this node as coordinator that were aborted
<code>Scans</code>	Scans (all)
<code>Range Scans</code>	Index scans

The `ndbd` process has a scheduler that runs in an infinite loop. During each loop scheduler performs the following tasks:

1. Read any incoming messages from sockets into a job buffer.
2. Check whether there are any timed messages to be executed; if so, put these into the job buffer as well.
3. Execute (in a loop) any messages in the job buffer.
4. Send any distributed messages that were generated by executing the messages in the job buffer.
5. Wait for any new incoming messages.

The number of loops executed in the third step is reported as the `Mean Loop Counter`. This statistic increases in size as the utilisation of the TCP/IP buffer improves. You can use this to monitor performance as you add new processes to the cluster.

The `Mean send size` and `Mean receive size` statistics allow you to gauge the efficiency of writes and reads (respectively) between nodes. These values are given in bytes. Higher values mean a lower cost per byte sent or received; the maximum is 64k.

To generate a report of all cluster log statistics, you can use the following command in the NDB management client:

```
ndb_mgm> ALL CLUSTERLOG STATISTICS=15
```

## 15.7.4. Single-User Mode

Single-user mode allows the database administrator to restrict access to the database system to a single MySQL server (SQL node). When entering single-user mode, all connections to all other MySQL servers are closed gracefully and all running transactions are aborted. No new transactions are allowed to be started.

Once the cluster has entered single-user mode, only the designated SQL node is granted access to the database.

You can use the `ALL STATUS` command to see when the cluster has entered single-user mode.

Example:

```
NDB> ENTER SINGLE USER MODE 5
```

After this command has executed and the cluster has entered single-user mode, the SQL node whose node ID is 5 becomes the cluster's only permitted user.

The node specified in the preceding command must be a MySQL Server node; An attempt to specify any other type of node will be rejected.

**Note:** When the preceding command is invoked, all transactions running on the designated node are aborted, the connection is closed, and the server must be restarted.

The command `EXIT SINGLE USER MODE` changes the state of the cluster's data nodes from single-user mode to normal mode. MySQL Servers waiting for a connection (that is, for the cluster to become ready and available), are again permitted to connect. The MySQL Server denoted as the single-user SQL node continues to run (if still connected) during and after the state change.

Example:

```
NDB> EXIT SINGLE USER MODE
```

There are two recommended ways to handle a node failure when running in single-user mode:

- Method 1:
  1. Finish all single-user mode transactions
  2. Issue the `EXIT SINGLE USER MODE` command
  3. Restart the cluster's data nodes

- Method 2:

Restart database nodes prior to entering single-user mode.

## 15.8. On-line Backup of MySQL Cluster

This section describes how to create a backup and how to restore the database from a backup at a later time.

## 15.8.1. Cluster Backup Concepts

A backup is a snapshot of the database at a given time. The backup consists of three main parts:

- **Metadata:** the names and definitions of all database tables
- **Table records:** the data actually stored in the database tables at the time that the backup was made
- **Transaction log:** a sequential record telling how and when data was stored in the database

Each of these parts is saved on all nodes participating in the backup. During backup, each node saves these three parts into three files on disk:

- `BACKUP-backup_id.node_id.ctl`

A control file containing control information and metadata. Each node saves the same table definitions (for all tables in the cluster) to its own version of this file.

- `BACKUP-backup_id-0.node_id.data`

A data file containing the table records, which are saved on a per-fragment basis. That is, different nodes save different fragments during the backup. The file saved by each node starts with a header that states the tables to which the records belong. Following the list of records there is a footer containing a checksum for all records.

- `BACKUP-backup_id.node_id.log`

A log file containing records of committed transactions. Only transactions on tables stored in the backup are stored in the log. Nodes involved in the backup save different records because different nodes host different database fragments.

In the listing above, *backup\_id* stands for the backup identifier and *node\_id* is the unique identifier for the node creating the file.

## 15.8.2. Using The Management Client to Create a Backup

Before starting a backup, make sure that the cluster is properly configured for performing one. (See [Section 15.8.4, “Configuration for Cluster Backup”](#).)

Creating a backup using the management client involves the following steps:

1. Start the management client (`ndb_mgm`).
2. Execute the command `START BACKUP`.
3. The management client will reply with the message `Start of backup ordered`. This means that the management client has submitted the request to the cluster, but has not yet received any response.
4. The management client will reply `Backup backup_id started`, where *backup\_id* is the unique identifier for this particular backup. (This identifier will also be saved in the cluster log, if it has not been configured otherwise.) This means that the cluster has received and processed the backup request. It does *not* mean that the backup has finished.



5. The management client will signal that the backup is finished with the message `Backup backup_id completed`.

Cluster backups are created by default in the `BACKUP` subdirectory of the `DataDir` on each data node. This can be overridden for one or more data nodes individually, or for all cluster data nodes in the `config.ini` file using the `BackupDataDir` configuration parameter as discussed in [Identifying Data Nodes](#). The backup files created for a backup with a given `backup_id` are stored in a subdirectory named `BACKUP-backup_id` in the backup directory.

To abort a backup that is already in progress:

1. Start the management client.
2. Execute the command `ABORT BACKUP backup_id`. The number `backup_id` is the identifier of the backup that was included in the response of the management client when the backup was started (in the message `Backup backup_id started`).
3. The management client will acknowledge the abort request with `Abort of backup backup_id ordered`; note that it has received no actual response to this request yet.
4. After the backup has been aborted, the management client will report `Backup backup_id has been aborted for reason XYZ`. This means that the cluster has terminated the backup and that all files related to this backup have been removed from the cluster filesystem.

It is also possible to abort a backup in progress from the system shell using this command:

```
shell> ndb_mgm -e "ABORT BACKUP backup_id"
```

**Note:** If there is no backup with ID `backup_id` running when it is aborted, the management client makes no explicit response. However, the fact that an invalid abort command was sent is indicated in the cluster log.

### 15.8.3. How to Restore a Cluster Backup

The cluster restoration program is implemented as a separate command-line utility `ndb_restore`, which can normally be found in the MySQL `bin` directory. This program reads the files created as a result of the backup and inserts the stored information into the database. `ndb_restore` must be executed once for each set of backup files — that is, once for each data node in the cluster at the time that the backup was created.

**Important:** In order to use `ndb_restore`, the cluster must be running in single user mode. See [Section 15.7.4, “Single-User Mode”](#).

Typical options for this utility are shown here:

```
ndb_restore [-c connectstring] -n node_id [-m] -b backup_id -r /path/to/backup/files
```

The `-c` option is used to specify a connectstring which tells `ndb_restore` where to locate the cluster management server. (See [Section 15.4.4.2, “The Cluster connectstring”](#), for information on connectstrings.) If this option is not used, then `ndb_restore` attempts to connect to a management server on `localhost:1186`. This utility acts as a cluster API node, and so requires a free connection “slot” to connect to the cluster management server. This means that there must be at least one `[API]` section that can be used by it in the cluster `config.ini` file. It is a good idea to keep at least one empty `[API]` section in `config.ini` that is not being used for a MySQL server or for an application for this reason (see [Section 15.4.4.6, “Defining SQL and Other API Nodes”](#)). `ndb_restore` can also employ

an unused `[MYSQLD]` section in the `config.ini` file for this purpose.

You can verify that `ndb_restore` is connected to the cluster by using the `SHOW` command in the `ndb_mgm` management client. You can also accomplish this from a system shell, as shown here:

```
shell> ndb_mgm -e "SHOW"
```

`-n` is used to specify the node ID of the data node on which the backups were taken.

The first time you run the `ndb_restore` restoration program, you also need to restore the metadata. In other words, you must re-create the database tables — this can be done by running it with the `-m` option. Note that the cluster should have an empty database when starting to restore a backup. (In other words, you should start `ndbd` with `--initial` prior to performing the restore.)

The `-b` option is used to specify the ID or sequence number of the backup, and is the same number shown by the management client in the `Backup backup_id completed` message displayed upon completion of a backup. (See [Section 15.8.2, “Using The Management Client to Create a Backup”](#).)

The `-r` option is required, and is used to tell `ndb_restore` the directory where it can find the backup files. **Important:** When restoring cluster backups, you must be sure to restore all data nodes from backups having the same backup ID.

It is possible to restore a backup to a database with a different configuration than it was created from. For example, suppose that a backup with backup ID `12`, created in a cluster with two database nodes having the node IDs `2` and `3`, is to be restored to a cluster with four nodes. Then `ndb_restore` must be run twice — once for each database node in the cluster where the backup was taken. However, `ndb_restore` cannot always restore backups made from a cluster running one version of MySQL to a cluster running a different MySQL version. See [Section 15.5.2, “Cluster Upgrade and Downgrade Compatibility”](#), for more information.

**Note:** For rapid restoration, the data may be restored in parallel, provided that there is a sufficient number of cluster connections available. However, the data files must always be applied before the logs.

A complete listing of options available for this program is shown in the following table:

Long Form	Short Form	Description	Default Value
<code>--backup-id</code>	<code>-b</code>	Backup sequence ID	<code>0</code>
<code>--backup_path</code>	<i>None</i>	Path to backup files	<code>./</code>
<code>-char-acter-sets-dir</code>	<i>None</i>	Specify the directory where character set information can be found	<i>None</i>
<code>--connect, -connectstring, or -ndb-connectstring</code>	<code>-c</code> or <code>-C</code>	Set the connectstring in <i>hos</i>	<i>None</i>

		<code>port</code> ] format	
<code>--core-file</code>	<i>None</i>	Write a core file in the event of an error	<code>TRUE</code>
<code>--debug</code>	<code>-#</code>	Output debug log	<code>d:t:0,/tmp/ndb_restore.trace</code>
<code>--help</code> or <code>--usage</code>	<code>-?</code>	Display help message with available options and current values, then exit	[N/A]
<code>--ndb-mgmd-host</code>	<i>None</i>	Set the host and port in <code>host[:port]</code> format for the management server to connect to; this is the same as <code>-connect</code> , <code>--connectstring</code> , or <code>--ndb-connectstring</code> , but without a way to specify the <code>nodeid</code>	<i>None</i>
<code>-ndb-nodegroup-map</code>	<code>-z</code>	Specifies a nodegroup map — <i>Syntax</i> : list of ( <code>source_nodegroup</code> , <code>destination_nodegroup</code> )	<i>None</i>
<code>--ndb-nodeid</code>	<i>None</i>	Specify a node ID for the <code>ndb_restore</code> process	<code>0</code>
<code>-ndb-optimized-node-selection</code>	<i>None</i>	Optimize selection of nodes for transactions	<code>TRUE</code>
<code>--ndb-shm</code>	<i>None</i>	Use shared memory connections when available	<code>FALSE</code>
<code>--nodeid</code>	<code>-n</code>	Use backup files from node with the specified ID	<code>0</code>
<code>--parallelism</code>	<code>-p</code>	Set from 1 to 1024 parallel transactions to be used during the restoration process	<code>128</code>
<code>--print</code>	<i>None</i>	Print data and log to <code>stdout</code>	<code>FALSE</code>
<code>--print_data</code>	<i>None</i>	Print data to <code>stdout</code>	<code>FALSE</code>
<code>--print_log</code>	<i>None</i>	Print log to <code>stdout</code>	<code>FALSE</code>
<code>--print_meta</code>	<i>None</i>	Print metadata to <code>stdout</code>	<code>FALSE</code>
<code>--restore_data</code>	<code>-r</code>	Restore data and logs	<code>FALSE</code>
<code>--restore_meta</code>	<code>-m</code>	Restore table metadata	<code>FALSE</code>
<code>--version</code>	<code>-V</code>	Output version information and exit	[N/A]

## 15.8.4. Configuration for Cluster Backup

Four configuration parameters are essential for backup:

- `BackupDataBufferSize`

The amount of memory used to buffer data before it is written to disk.

- `BackupLogBufferSize`  
The amount of memory used to buffer log records before these are written to disk.
- `BackupMemory`  
The total memory allocated in a database node for backups. This should be the sum of the memory allocated for the backup data buffer and the backup log buffer.
- `BackupWriteSize`  
The default size of blocks written to disk. This applies for both the backup data buffer and the backup log buffer.
- `BackupMaxWriteSize`  
The maximum size of blocks written to disk. This applies for both the backup data buffer and the backup log buffer.

More detailed information about these parameters can be found in [Backup Parameters](#).

## 15.8.5. Backup Troubleshooting

If an error code is returned when issuing a backup request, the most likely cause is insufficient memory or disk space. You should check that there is enough memory allocated for the backup. **Important:** If you have set `BackupDataBufferSize` and `BackupLogBufferSize` and their sum is greater than 4MB, then you must also set `BackupMemory` as well. See [BackupMemory](#).

You should also make sure that there is sufficient space on the hard drive partition of the backup target.

**NDB** does not support repeatable reads, which can cause problems with the restoration process. Although the backup process is “hot”, restoring a MySQL Cluster from backup is not a 100% “hot” process. This is due to the fact that, for the duration of the restore process, running transactions get non-repeatable reads from the restored data. This means that the state of the data is inconsistent while the restore is in progress.

## 15.9. Cluster Utility Programs

This section discusses the MySQL Cluster utility programs that can be found in the `mysql/bin` directory. Each of these — except for `ndb_size.pl` and `ndb_error_reporter` — is a standalone binary that can be used from a system shell, and that does not need to connect to a MySQL server (nor even requires that a MySQL server be connected to the cluster).

These utilities can also serve as examples for writing your own applications using the **NDB** API. The source code for most of these programs may be found in the `ndb/tools` directory of the MySQL 4.1 tree (see [Section 2.9, “MySQL Installation Using a Source Distribution”](#)). The **NDB** API is not covered in this manual; please refer to the [NDB API Guide](http://dev.mysql.com/doc//ndbapi/en/) [http://dev.mysql.com/doc//ndbapi/en/] for information about this API.

All of the **NDB** utilities are listed here with brief descriptions:

- `ndb_delete_all`: Deletes all rows from a given table.

- `ndb_desc`: Lists all properties of an `NDB` table.
- `ndb_drop_index`: Drops the specified index from an `NDB` table.
- `ndb_drop_table`: Drops an `NDB` table.
- `ndb_error_reporter`: Can be used to gather information useful for diagnosing problems with the cluster.
- `ndb_mgm`: This is the MySQL Cluster management client, which is discussed in [Section 15.7.2, “Commands in the Management Client”](#).
- `ndb_print_backup_file`: Prints diagnostic information obtained from cluster backup files.
- `ndb_print_schema_file`: Prints diagnostic information obtained from cluster schema files.
- `ndb_print_sys_file`: Prints diagnostic information obtained from cluster system files.
- `ndb_restore`: This utility is used to restore a cluster from backup. See [Section 15.8.3, “How to Restore a Cluster Backup”](#), for more information.
- `ndb_select_all`: Prints all rows from an `NDB` table.
- `ndb_select_count`: Gets the number of rows in one or more `NDB` tables.
- `ndb_show_tables`: Shows all `NDB` tables anywhere in the cluster.
- `ndb_size.pl`: Examines all the tables in a given non-Cluster database and calculates the amount of storage each would require if it were converted to use the `NDB` storage engine.
- `ndb_waiter`: Reports on the status of cluster data nodes in a manner similar to that of the management client command `ALL STATUS`.

Most of these utilities need to connect to a Cluster management server in order to function. The exceptions are `ndb_size.pl` (see below), and the following utilities which access a cluster data node filesystem and so need to be run on a data node host:

- `ndb_print_backup_file`
- `ndb_print_schema_file`
- `ndb_print_sys_file`

`ndb_size.pl` is a Perl script which is also intended to be used from the shell; however it is a MySQL application and must be able to connect to a MySQL server. See [Section 15.9.12, “ndb\\_size.pl”](#), for additional requirements for using this script.

`ndb_error_reporter` is also a Perl script. It is used to gather cluster data node and management node logs together into a tarball to submit along with a bug report. It can use `ssh` or `scp` to access the node filesystems remotely.

Additional information about each of these utilities (except for `ndb_mgm` and `ndb_restore`) can be found in the sections that follow.

**Note:** All of these utilities (except for `ndb_size.pl`) can use the options discussed in [Section 15.6.5, “Command Options for MySQL Cluster Processes”](#). Additional options specific to each utility program are discussed in the individual program listings.

The order in which these options are used is generally not important. For example, all of these commands produce exactly the same output:

- `ndb_desc -c localhost fish -d test`
- `ndb_desc fish -c localhost -d test`
- `ndb_desc -d test fish -c localhost`

## 15.9.1. `ndb_delete_all`

**Description:** Deletes all rows from the given NDB table. In some cases, this can be much faster than `DELETE` or even `TRUNCATE`.

**Usage:**

```
ndb_delete_all -c connect_string tbl_name -d db_name
```

This deletes all rows from the table named `tbl_name` in the database named `db_name`. It is exactly equivalent to executing `TRUNCATE db_name.tbl_name` in MySQL.

**Additional Options:**

- `-t, --transactional`

Use of this option causes the delete operation to be performed as a single transaction.

**Warning:** With very large tables, this using this option may cause the number of operations available to the cluster to be exceeded.

## 15.9.2. `ndb_desc`

**Description:** Provides a detailed description of one or more NDB tables.

**Usage:**

```
ndb_desc -c connect_string tbl_name -d db_name
```

**Sample Output:**

MySQL table creation and population statements:

```
USE test;

CREATE TABLE fish (
  id INT(11) NOT NULL AUTO_INCREMENT,
  name VARCHAR(20),

  PRIMARY KEY pk (id),
  UNIQUE KEY uk (name)
) ENGINE=NDBCLUSTER;

INSERT INTO fish VALUES
  ('','guppy'), ('','tuna'), ('','shark'),
  ('','manta ray'), ('','grouper'), ('','puffer');
```

Output from `ndb_desc`:

```
shell> ./ndb_desc -c localhost fish -d test -p
-- fish --
Version: 16777221
Fragment type: 5
K Value: 6
Min load factor: 78
Max load factor: 80
Temporary table: no
Number of attributes: 2
Number of primary keys: 1
Length of frm data: 268
Row Checksum: 1
Row GCI: 1
TableStatus: Retrieved
-- Attributes --
id Int PRIMARY KEY DISTRIBUTION KEY AT=FIXED ST=MEMORY
name Varchar(20;latin1_swedish_ci) NULL AT=SHORT_VAR ST=MEMORY

-- Indexes --
PRIMARY KEY(id) - UniqueHashIndex
uk(name) - OrderedIndex
PRIMARY(id) - OrderedIndex
uk$unique(name) - UniqueHashIndex

-- Per partition info --
Partition  Row count  Commit count  Frag fixed memory  Frag var sized memory
2          2          2             65536              327680
1          2          2             65536              327680
3          2          2             65536              327680

NDBT_ProgramExit: 0 - OK
```

**Additional Options:**

- `-p, --extra-partition-info`  
Prints additional information about the table's partitions.
- Information about multiple tables can be obtained in a single invocation of `ndb_desc` by using their names, separated by spaces. All of the tables must be in the same database.

### 15.9.3. `ndb_drop_index`

**Description:** Drops the specified index from an NDB table. *It is recommended that you use this utility only as an example for writing NDB API applications* — see the Warning later in this section for details.

**Usage:**

```
ndb_drop_index -c connect_string table_name index -d db_name
```

The statement shown above drops the index named `index` from the `table` in the `database`.

**Additional Options:** None that are specific to this application.

**Warning:** Operations performed on Cluster table indexes using the NDB API are not visible to MySQL and make the table unusable by a MySQL server. If you use this program to drop an index, then try to access the table from an SQL node, an error results, as shown here:

```
shell> ./ndb_drop_index -c localhost dogs ix -d ctest1
Dropping index dogs/idx...OK

NDBT_ProgramExit: 0 - OK

shell> ./mysql -u jon -p ctest1
```

```

Enter password: *****
Reading table information for completion of table and column names
You can turn off this feature to get a quicker startup with -A

Welcome to the MySQL monitor.  Commands end with ; or \g.
Your MySQL connection id is 7 to server version: 5.1.12-beta-20060817

Type 'help;' or '\h' for help. Type '\c' to clear the buffer.

mysql> SHOW TABLES;
+-----+
| Tables_in_ctest1 |
+-----+
| a                 |
| bt1               |
| bt2               |
| dogs              |
| employees         |
| fish              |
+-----+
6 rows in set (0.00 sec)

mysql> SELECT * FROM dogs;
ERROR 1296 (HY000): Got error 4243 'Index not found' from NDBCLUSTER

```

In such a case, your *only* option for making the table available to MySQL again is to drop the table and re-create it. You can use either the SQL statement `DROP TABLE` or the `ndb_drop_table` utility (see [Section 15.9.4, “ndb\\_drop\\_table”](#)) to drop the table.

## 15.9.4. `ndb_drop_table`

**Description:** Drops the specified NDB table. (If you try to use this on a table created with a storage engine other than NDB, it fails with the error 723: No such table exists.) This operation is extremely fast — in some cases, it can be an order of magnitude faster than using `DROP TABLE` on an NDB table from MySQL.

**Usage:**

```
ndb_drop_table -c connect_string tbl_name -d db_name
```

**Additional Options:** None.

## 15.9.5. `ndb_error_reporter`

**Description:** Creates an archive from data node and management node logfiles that can be used to help diagnose bugs or other problems with a cluster. *It is highly recommended that you make use of this utility when filing reports of bugs in MySQL Cluster.*

**Usage:**

```
ndb_error_reporter path/to/config-file [username] [--fs]
```

This utility is intended for use on a management node host, and requires the path to the management host configuration file (`config.ini`). Optionally, you can supply the name of a user that is able to access the cluster's data nodes via SSH, in order to copy the data node logfiles. `ndb_error_reporter` then includes all of these files in archive that is created in the same directory in which it is run. The archive is named `ndb_error_report_YYYYMMDDHHMMSS.tar.bz2`, where `YYYYMMDDHHMMSS` is a date-time string.

If the `--fs` is used, then the data node filesystems are also copied to the management host and included in the archive that is produced by this script. As data node filesystems can be extremely large even after being compressed, we ask that you please do *not* send archives created using this option to MySQL AB unless you are specifically requested to do so.



## 15.9.6. `ndb_print_backup_file`

**Description:** Obtains diagnostic information from a cluster backup file.

**Usage:**

```
ndb_print_backup_file file_name
```

*file\_name* is the name of a cluster backup file. This can be any of the files (`.Data`, `.ctl`, or `.log` file) found in a cluster backup directory. These files are found in the data node's backup directory under the subdirectory `BACKUP-#`, where `#` is the sequence number for the backup. For more information about cluster backup files and their contents, see [Section 15.8.1, “Cluster Backup Concepts”](#).

Like `ndb_print_schema_file` and `ndb_print_sys_file` (and unlike most of the other NDB utilities that are intended to be run on a management server host or to connect to a management server) `ndb_print_backup_file` must be run on a cluster data node, since it accesses the data node filesystem directly. Because it does not make use of the management server, this utility can be used when the management server is not running, and even when the cluster has been completely shut down.

**Additional Options:** None.

## 15.9.7. `ndb_print_schema_file`

**Description:** Obtains diagnostic information from a cluster schema file.

**Usage:**

```
ndb_print_schema_file file_name
```

*file\_name* is the name of a cluster schema file.

Like `ndb_print_backup_file` and `ndb_print_sys_file` (and unlike most of the other NDB utilities that are intended to be run on a management server host or to connect to a management server) `ndb_print_schema_file` must be run on a cluster data node, since it accesses the data node filesystem directly. Because it does not make use of the management server, this utility can be used when the management server is not running, and even when the cluster has been completely shut down.

**Additional Options:** None.

## 15.9.8. `ndb_print_sys_file`

**Description:** Obtains diagnostic information from a cluster system file.

**Usage:**

```
ndb_print_sys_file file_name
```

*file\_name* is the name of a cluster system file (sysfile). Cluster system files are located in a data node's data directory (`DataDir`); the path under this directory to system files matches the pattern `ndb_#_fs/D#/DBDIH/P#.sysfile`. In each case, the `#` represents a number (not necessarily the same number).

Like `ndb_print_backup_file` and `ndb_print_schema_file` (and unlike most of the other NDB utilities that are intended to be run on a management server host or to connect to a management server) `ndb_print_sys_file` must be run on a cluster data node, since it accesses the data node filesystem directly. Because it does not make use of the management server, this utility can be used when the management server is not running, and even when the cluster has been completely shut down.

**Additional Options:** None.

### 15.9.9. `ndb_select_all`

**Description:** Prints all rows from an NDB table to `stdout`.

**Usage:**

```
ndb_select_all -c connect_string tbl_name -d db_name [> file_name]
```

**Additional Options:**

- `-l lock_type, --lock=lock_type`

Employs a lock when reading the table. Possible values for `lock_type` are:

- 0: Read lock
- 1: Read lock with hold
- 2: Exclusive read lock

There is no default value for this option.

- `-o index_name, --order=index_name`

Orders the output according to the index named `index_name`. Note that this is the name of an index, not of a column, and that the index must have been explicitly named when created.

- `-z, --descending`

Sorts the output in descending order. This option can be used only in conjunction with the `-o` (`--order`) option.

- `--header=FALSE`

Excludes column headers from the output.

- `-x, --useHexFormat`

Causes all numeric values to be displayed in hexadecimal format. This does not affect the output of numerals contained in strings or datetime values.

`-D character, --delimiter=character`

Causes the `character` to be used as a column delimiter. Only table data columns are separated by this delimiter.

The default delimiter is the tab character.

- `--rowid`

Adds a `ROWID` column providing information about the fragments in which rows are stored.

- `--gci`

Adds a column to the output showing the global checkpoint at which each row was last updated. See [Section 15.13, “MySQL Cluster Glossary”](#), and [Section 15.7.3.2, “Log Events”](#), for more informa-

tion about checkpoints.

- `-t, --tupscan`

Scan the table in the order of the tuples.

- `--nodata`

Causes any table data to be omitted.

### Sample Output:

Output from a MySQL `SELECT` statement:

```
mysql> SELECT * FROM ctest1.fish;
+----+-----+
| id | name  |
+----+-----+
| 3  | shark |
| 6  | puffer|
| 2  | tuna  |
| 4  | manta ray|
| 5  | grouper|
| 1  | guppy |
+----+-----+
6 rows in set (0.04 sec)
```

Output from the equivalent invocation of `ndb_select_all`:

```
shell> ./ndb_select_all -c localhost fish -d ctest1
id      name
3       [shark]
6       [puffer]
2       [tuna]
4       [manta ray]
5       [grouper]
1       [guppy]
6 rows returned

NDBT_ProgramExit: 0 - OK
```

Note that all string values are enclosed by square brackets (“[...]”) in the output of `ndb_select_all`. For a further example, consider the table created and populated as shown here:

```
CREATE TABLE dogs (
  id INT(11) NOT NULL AUTO_INCREMENT,
  name VARCHAR(25) NOT NULL,
  breed VARCHAR(50) NOT NULL,
  PRIMARY KEY pk (id),
  KEY ix (name)
)
ENGINE=NDB;

INSERT INTO dogs VALUES
  ('', 'Lassie', 'collie'),
  ('', 'Scooby-Doo', 'Great Dane'),
  ('', 'Rin-Tin-Tin', 'German Shepherd'),
  ('', 'Rosscoe', 'Mutt');
```

This demonstrates the use of several additional `ndb_select_all` options:

```
shell> ./ndb_select_all -d ctest1 dogs -o ix -z --gci
GCI      id name      breed
834461  2 [Scooby-Doo] [Great Dane]
834878  4 [Rosscoe]    [Mutt]
834463  3 [Rin-Tin-Tin] [German Shepherd]
835657  1 [Lassie]     [Collie]
4 rows returned
```

```
NDBT_ProgramExit: 0 - OK
```

## 15.9.10. `ndb_select_count`

**Description:** Prints the number of rows in one or more NDB table. With a single table, the result is equivalent to that obtained by using the MySQL statement `SELECT COUNT(*) FROM tbl_name`.

**Usage:**

```
ndb_select_count [-c connect_string] -ddb_name tbl_name[, tbl_name2[, ...]]
```

**Additional Options:** None that are specific to this application. However, you can obtain row counts from multiple tables in the same database by listing the table names separated by spaces when invoking this command, as shown under **Sample Output**.

**Sample Output:**

```
shell> ./ndb_select_count -c localhost -d ctest1 fish dogs
6 records in table fish
4 records in table dogs
NDBT_ProgramExit: 0 - OK
```

## 15.9.11. `ndb_show_tables`

**Description:** Displays a list of all NDB database objects in the cluster. By default, this includes not only both user-created tables and NDB system tables, but NDB-specific indexes, and internal triggers, as well.

**Usage:**

```
ndb_show_tables [-c connect_string]
```

**Additional Options:**

- `-l, --loops`

Specifies the number of times the utility should execute. This is 1 when this option is not specified, but if you do use the option, you must supply an integer argument for it.

- `-p, --parsable`

Using this option causes the output to be in a format suitable for use with `LOAD DATA INFILE`.

- `-t, --type`

Can be used to restrict the output to one type of object, specified by an integer type code as shown here:

- **1:** System table
- **2:** User-created table
- **3:** Unique hash index

Any other value causes all NDB database objects to be listed (the default).

- `-u, --unqualified`

If specified, this causes unqualified object names to be displayed.

**Note:** Only user-created Cluster tables may be accessed from MySQL; system tables such as `SYSTAB_0` are not visible to `mysqld`. However, you can examine the contents of system tables using NDB API applications such as `ndb_select_all` (see [Section 15.9.9, “ndb\\_select\\_all”](#)).

## 15.9.12. `ndb_size.pl`

**Description:** This is a Perl script that can be used to estimate the amount of space that would be required by a MySQL database if it were converted to use the `NDBCluster` storage engine. Unlike the other utilities discussed in this section, it does not require access to a MySQL Cluster (in fact, there is no reason for it to do so). However, it does need to access the MySQL server on which the database to be tested resides.

### Requirements:

- A running MySQL server. The server instance does not have to provide support for MySQL Cluster.
- A working installation of Perl.
- The `DBI` and `HTML::Template` modules, both of which can be obtained from CPAN if they are not already part of your Perl installation. (Many Linux and other operating system distribution provide their own packages for one or both of these libraries.)
- The `ndb_size.tpl` template file, which you should be able to find in the `share/mysql` directory of your MySQL installation. This file should be copied or moved into the same directory as `ndb_size.pl` — if it is not there already — before running the script.
- A MySQL user account having the necessary privileges. If you do not wish to use an existing account, then creating one using `GRANT USAGE ON db_name.*` — where `db_name` is the name of the database to be examined — is sufficient for this purpose.

`ndb_size.pl` and `ndb_size.tpl` can also be found in the MySQL sources in `storage/ndb/tools`. If these files are not present in your MySQL installation, you can obtain them from the [MySQLForge project page](http://forge.mysql.com/projects/view.php?id=88) [http://forge.mysql.com/projects/view.php?id=88].

### Usage:

```
perl ndb_size.pl db_name hostname username password > file_name.html
```

The command shown connects to the MySQL server at `hostname` using the account of the user `username` having the password `password`, analyses all of the tables in database `db_name`, and generates a report in HTML format which is directed to the file `file_name.html`. (Without the redirection, the output is sent to `stdout`.) This figure shows partial sample output as viewed in a Web browser:

## MySQL Cluster analysis for world

This is an automated analysis of the DBI:mysql:database=world;host=192.168.0.176 database for migration into MySQL Cluster. No warranty is made to the accuracy of the information.

This information should be valid for MySQL 4.1 and 5.0. Since 5.1 is not a final release yet, the numbers should be used as a guide only.

### Parameter Settings

**NOTE** the configuration parameters below do not take into account system tables and other requirements.

Parameter	4.1	5.0	5.1
DataMemory (kb)	544	544	544
IndexMemory (kb)	192	136	136
MaxNoOfTables	3	3	3
MaxNoOfAttributes	24	24	24
MaxNoOfOrderedIndexes	3	3	3
MaxNoOfUniqueHashIndexes	3	3	3
MaxNoOfTriggers	12	12	12

### Memory usage because of parameters

Usage is in kilobytes. Actual usage will vary as you should set the parameters larger than those listed in the table above.

Parameter	4.1	5.0	5.1
Attributes	5	5	5
Tables	60	60	60
OrderedIndexes	30	30	30
UniqueHashIndexes	45	45	45

### Table List

- [City](#)
- [Country](#)
- [CountryLanguage](#)

#### City

Column	Type	Size	Key	4.1 NDB Size	5.0 NDB Size	5.1 NDB Size
ID	int	11	PR1	4	4	4
District	char	20		20	20	20
Name	char	35		36	36	36
Population	int	11		4	4	4
CountryCode	char	3		4	4	4

### Indexes

We assume that indexes are ORDERED (not created USING HASH). If order is not required, 10 bytes of data memory can be saved per row if the index is created USING HASH

Index	Type	Columns	4.1 IdxMem	5.0 IdxMem	5.1 IdxMem	4.1 DatMem	5.0 DatMem	5.1 DatMem
PRIMARY	BTREE	ID	29	25	25	10	10	10

### DataMemory Usage

	4.1	5.0	5.1
Row Overhead	16	16	16
Column DataMemory/Row	68	68	68
Index DataMemory/Row	10	10	10
Total DataMemory/Row	94	94	94
Rows per 32 kb page	347	347	347
Current number of rows	4079	4079	4079
Total DataMemory (kb)	384	384	384

### IndexMemory Usage

	4.1	5.0	5.1
IndexMemory/Row	29	25	25
Rows per 8kb page	282	327	327
Current number of rows	4079	4079	4079
Total IndexMemory (kb)	120	104	104

The output from this script includes:

- Minimum values for the `DataMemory`, `IndexMemory`, `MaxNoOfTables`, `MaxNoOfAttributes`, `MaxNoOfOrderedIndexes`, `MaxNoOfUniqueHashIndexes`, and `MaxNoOfTriggers` configuration parameters required to accommodate the tables analysed.
- Memory requirements for all of the tables, attributes, ordered indexes, and unique hash indexes defined in the database.
- The `IndexMemory` and `DataMemory` required per table and table row.

### 15.9.13. `ndb_waiter`

**Description:** Repeatedly (each 100 milliseconds) prints out the status of all cluster data nodes until either the cluster reaches a given status or the `--timeout` limit is exceeded, then exits. By default, it waits for the cluster to achieve `STARTED` status, in which all nodes have started and connected to the cluster. This can be overridden using the `--no-contact` and `--not-started` options (see [Additional Options](#)).

The node states reported by this utility are as follows:

- `NO_CONTACT`: The node cannot be contacted.
- `UNKNOWN`: The node can be contacted, but its status is not yet known. Usually, this means that the node has received a `START` or `RESTART` command from the management server, but has not yet acted on it.
- `NOT_STARTED`: The node has stopped, but remains in contact with the cluster. This is seen when restarting the node using the management client's `RESTART` command.
- `STARTING`: The node's `ndbd` process has started, but the node has not yet joined the cluster.
- `STARTED`: The node is operational, and has joined the cluster.
- `SHUTTING_DOWN`: The node is shutting down.
- `SINGLE USER MODE`: This is shown for all cluster data nodes when the cluster is in single user mode.

**Usage:**

```
ndb_waiter [-c connect_string]
```

**Additional Options:**

- `-n, --no-contact`

Instead of waiting for the `STARTED` state, `ndb_waiter` continues running until the cluster reaches `NO_CONTACT` status before exiting.

- `--not-started`

Instead of waiting for the `STARTED` state, `ndb_waiter` continues running until the cluster reaches `NOT_STARTED` status before exiting.

- `-t, --timeout=#`

Time to wait. The program exits if the desired state is not achieved within this number of seconds. The default is 120 seconds (1200 reporting cycles).

### Sample Output:

Shown here is the output from `ndb_waiter` when run against a 4-node cluster in which two nodes have been shut down and then started again manually. Duplicate reports (indicated by “...”) are omitted.

```
shell> ./ndb_waiter -c localhost

Connecting to mgmsrv at (localhost)
State node 1 STARTED
State node 2 NO_CONTACT
State node 3 STARTED
State node 4 NO_CONTACT
Waiting for cluster enter state STARTED

...

State node 1 STARTED
State node 2 UNKNOWN
State node 3 STARTED
State node 4 NO_CONTACT
Waiting for cluster enter state STARTED

...

State node 1 STARTED
State node 2 STARTING
State node 3 STARTED
State node 4 NO_CONTACT
Waiting for cluster enter state STARTED

...

State node 1 STARTED
State node 2 STARTING
State node 3 STARTED
State node 4 UNKNOWN
Waiting for cluster enter state STARTED

...

State node 1 STARTED
State node 2 STARTING
State node 3 STARTED
State node 4 STARTING
Waiting for cluster enter state STARTED

...

State node 1 STARTED
State node 2 STARTED
State node 3 STARTED
State node 4 STARTING
Waiting for cluster enter state STARTED

...

State node 1 STARTED
State node 2 STARTED
State node 3 STARTED
State node 4 STARTED
Waiting for cluster enter state STARTED

NDBT_ProgramExit: 0 - OK
```

**Note:** If no connectstring is specified, then `ndb_waiter` tries to connect to a management on `localhost`, and reports `Connecting to mgmsrv at (null)`.



## 15.10. Using High-Speed Interconnects with MySQL Cluster

Even before design of [NDB Cluster](#) began in 1996, it was evident that one of the major problems to be encountered in building parallel databases would be communication between the nodes in the network. For this reason, [NDB Cluster](#) was designed from the very beginning to allow for the use of a number of different data transport mechanisms. In this Manual, we use the term *transporter* for these.

The MySQL Cluster codebase includes support for four different transporters:

- *TCP/IP using 100 Mbps or gigabit Ethernet*, as discussed in [Section 15.4.4.7, “Cluster TCP/IP Connections”](#).
- *Direct (machine-to-machine) TCP/IP*; although this transporter uses the same TCP/IP protocol as mentioned in the previous item, it requires setting up the hardware differently and is configured differently as well. For this reason, it is considered a separate transport mechanism for MySQL Cluster. See [Section 15.4.4.8, “TCP/IP Connections Using Direct Connections”](#), for details.
- *Shared memory (SHM)*. For more information about SHM, see [Section 15.4.4.9, “Shared-Memory Connections”](#).
- *Scalable Coherent Interface (SCI)*, as described in the next section of this chapter, [Section 15.4.4.10, “SCI Transport Connections”](#).

Most users today employ TCP/IP over Ethernet because it is ubiquitous. TCP/IP is also by far the best-tested transporter for use with MySQL Cluster.

We are working to make sure that communication with the `ndbd` process is made in “chunks” that are as large as possible because this benefits all types of data transmission.

For users who desire it, it is also possible to use cluster interconnects to enhance performance even further. There are two ways to achieve this: Either a custom transporter can be designed to handle this case, or you can use socket implementations that bypass the TCP/IP stack to one extent or another. We have experimented with both of these techniques using the SCI (Scalable Coherent Interface) technology developed by [Dolphin](http://www.dolphinics.com/) [http://www.dolphinics.com/].

### 15.10.1. Configuring MySQL Cluster to use SCI Sockets

In this section, we show how to adapt a cluster configured for normal TCP/IP communication to use SCI Sockets instead. This documentation is based on SCI Sockets version 2.3.0 as of 01 October 2004.

#### Prerequisites

Any machines with which you wish to use SCI Sockets must be equipped with SCI cards.

It is possible to use SCI Sockets with any version of MySQL Cluster. No special builds are needed because it uses normal socket calls which are already available in MySQL Cluster. However, SCI Sockets are currently supported only on the Linux 2.4 and 2.6 kernels. SCI Transporters have been tested successfully on additional operating systems although we have verified these only with Linux 2.4 to date.

There are essentially four requirements for SCI Sockets:

- Building the SCI Socket libraries.
- Installation of the SCI Socket kernel libraries.

- Installation of one or two configuration files.
- The SCI Socket kernel library must be enabled either for the entire machine or for the shell where the MySQL Cluster processes are started.

This process needs to be repeated for each machine in the cluster where you plan to use SCI Sockets for inter-node communication.

Two packages need to be retrieved to get SCI Sockets working:

- The source code package containing the DIS support libraries for the SCI Sockets libraries.
- The source code package for the SCI Socket libraries themselves.

Currently, these are available only in source code format. The latest versions of these packages at the time of this writing were available as (respectively) `DIS_GPL_2_5_0_SEP_10_2004.tar.gz` and `SCI_SOCKET_2_3_0_OKT_01_2004.tar.gz`. You should be able to find these (or possibly newer versions) at <http://www.dolphinics.no/support/downloads.html>.

### Package Installation

Once you have obtained the library packages, the next step is to unpack them into appropriate directories, with the SCI Sockets library unpacked into a directory below the DIS code. Next, you need to build the libraries. This example shows the commands used on Linux/x86 to perform this task:

```
shell> tar xzf DIS_GPL_2_5_0_SEP_10_2004.tar.gz
shell> cd DIS_GPL_2_5_0_SEP_10_2004/src/
shell> tar xzf ../../SCI_SOCKET_2_3_0_OKT_01_2004.tar.gz
shell> cd ../adm/bin/Linux_pkgs
shell> ./make_PSB_66_release
```

It is possible to build these libraries for some 64-bit processors. To build the libraries for Opteron CPUs using the 64-bit extensions, run `make_PSB_66_X86_64_release` rather than `make_PSB_66_release`. If the build is made on an Itanium machine, you should use `make_PSB_66_IA64_release`. The X86-64 variant should work for Intel EM64T architectures but this has not yet (to our knowledge) been tested.

Once the build process is complete, the compiled libraries will be found in a zipped tar file with a name along the lines of `DIS-<operating-system>-time-date`. It is now time to install the package in the proper place. In this example we will place the installation in `/opt/DIS`. (**Note:** You will most likely need to run the following as the system `root` user.)

```
shell> cp DIS_Linux_2.4.20-8_181004.tar.gz /opt/
shell> cd /opt
shell> tar xzf DIS_Linux_2.4.20-8_181004.tar.gz
shell> mv DIS_Linux_2.4.20-8_181004 DIS
```

### Network Configuration

Now that all the libraries and binaries are in their proper place, we need to ensure that the SCI cards have proper node IDs within the SCI address space.

It is also necessary to decide on the network structure before proceeding. There are three types of network structures which can be used in this context:

- A simple one-dimensional ring

- One or more SCI switches with one ring per switch port
- A two- or three-dimensional torus.

Each of these topologies has its own method for providing node IDs. We discuss each of them in brief.

A simple ring uses node IDs which are non-zero multiples of 4: 4, 8, 12,...

The next possibility uses SCI switches. An SCI switch has 8 ports, each of which can support a ring. It is necessary to make sure that different rings use different node ID spaces. In a typical configuration, the first port uses node IDs below 64 (4 – 60), the next 64 node IDs (68 – 124) are assigned to the next port, and so on, with node IDs 452 – 508 being assigned to the eighth port.

Two- and three-dimensional torus network structures take into account where each node is located in each dimension, incrementing by 4 for each node in the first dimension, by 64 in the second dimension, and (where applicable) by 1024 in the third dimension. See [Dolphin's Web site](http://www.dolphinics.com/support/index.html) [http://www.dolphinics.com/support/index.html] for more thorough documentation.

In our testing we have used switches, although most large cluster installations use 2- or 3-dimensional torus structures. The advantage provided by switches is that, with dual SCI cards and dual switches, it is possible to build with relative ease a redundant network where the average failover time on the SCI network is on the order of 100 microseconds. This is supported by the SCI transporter in MySQL Cluster and is also under development for the SCI Socket implementation.

Failover for the 2D/3D torus is also possible but requires sending out new routing indexes to all nodes. However, this requires only 100 milliseconds or so to complete and should be acceptable for most high-availability cases.

By placing cluster data nodes properly within the switched architecture, it is possible to use 2 switches to build a structure whereby 16 computers can be interconnected and no single failure can hinder more than one of them. With 32 computers and 2 switches it is possible to configure the cluster in such a manner that no single failure can cause the loss of more than two nodes; in this case, it is also possible to know which pair of nodes is affected. Thus, by placing the two nodes in separate node groups, it is possible to build a “safe” MySQL Cluster installation.

To set the node ID for an SCI card use the following command in the `/opt/DIS/sbin` directory. In this example, `-c 1` refers to the number of the SCI card (this is always 1 if there is only 1 card in the machine); `-a 0` refers to adapter 0; and `68` is the node ID:

```
shell> ./sciconfig -c 1 -a 0 -n 68
```

If you have multiple SCI cards in the same machine, you can determine which card has which slot by issuing the following command (again we assume that the current working directory is `/opt/DIS/sbin`):

```
shell> ./sciconfig -c 1 -gsn
```

This will give you the SCI card's serial number. Then repeat this procedure with `-c 2`, and so on, for each card in the machine. Once you have matched each card with a slot, you can set node IDs for all cards.

After the necessary libraries and binaries are installed, and the SCI node IDs are set, the next step is to set up the mapping from hostnames (or IP addresses) to SCI node IDs. This is done in the SCI sockets configuration file, which should be saved as `/etc/sci/scisock.conf`. In this file, each SCI node ID is mapped through the proper SCI card to the hostname or IP address that it is to communicate with. Here is a very simple example of such a configuration file:

```
#host      #nodeId
alpha      8
beta       12
192.168.10.20 16
```

It is also possible to limit the configuration so that it applies only to a subset of the available ports for these hosts. An additional configuration file `/etc/sci/scisock_opt.conf` can be used to accomplish this, as shown here:

```
#-key          -type      -values
EnablePortsByDefault  yes
EnablePort          tcp         2200
DisablePort         tcp         2201
EnablePortRange      tcp         2202 2219
DisablePortRange     tcp         2220 2231
```

## Driver Installation

With the configuration files in place, the drivers can be installed.

First, the low-level drivers and then the SCI socket driver need to be installed:

```
shell> cd DIS/sbin/
shell> ./drv-install add PSB66
shell> ./scisocket-install add
```

If desired, the installation can be checked by invoking a script which verifies that all nodes in the SCI socket configuration files are accessible:

```
shell> cd /opt/DIS/sbin/
shell> ./status.sh
```

If you discover an error and need to change the SCI socket configuration, it is necessary to use `ksocket-etconfig` to accomplish this task:

```
shell> cd /opt/DIS/util
shell> ./ksocketconfig -f
```

## Testing the Setup

To ensure that SCI sockets are actually being used, you can employ the `latency_bench` test program. Using this utility's server component, clients can connect to the server to test the latency of the connection. Determining whether SCI is enabled should be fairly simple from observing the latency. (**Note:** Before using `latency_bench`, it is necessary to set the `LD_PRELOAD` environment variable as shown later in this section.)

To set up a server, use the following:

```
shell> cd /opt/DIS/bin/socket
shell> ./latency_bench -server
```

To run a client, use `latency_bench` again, except this time with the `-client` option:

```
shell> cd /opt/DIS/bin/socket
shell> ./latency_bench -client server_hostname
```

SCI socket configuration should now be complete and MySQL Cluster ready to use both SCI Sockets and the SCI transporter (see [Section 15.4.4.10, "SCI Transport Connections"](#)).

## Starting the Cluster

The next step in the process is to start MySQL Cluster. To enable usage of SCI Sockets it is necessary to set the environment variable `LD_PRELOAD` before starting `ndbd`, `mysqld`, and `ndb_mgmd`. This variable should point to the kernel library for SCI Sockets.

To start `ndbd` in a bash shell, do the following:

```
bash-shell> export LD_PRELOAD=/opt/DIS/lib/libkscisock.so
bash-shell> ndbd
```

In a `tcsh` environment the same thing can be accomplished with:

```
tcsh-shell> setenv LD_PRELOAD=/opt/DIS/lib/libkscisock.so
tcsh-shell> ndbd
```

**Note:** MySQL Cluster can use only the kernel variant of SCI Sockets.

## 15.10.2. Understanding the Impact of Cluster Interconnects

The `ndbd` process has a number of simple constructs which are used to access the data in a MySQL Cluster. We have created a very simple benchmark to check the performance of each of these and the effects which various interconnects have on their performance.

There are four access methods:

- **Primary key access**

This is access of a record through its primary key. In the simplest case, only one record is accessed at a time, which means that the full cost of setting up a number of TCP/IP messages and a number of costs for context switching are borne by this single request. In the case where multiple primary key accesses are sent in one batch, those accesses share the cost of setting up the necessary TCP/IP messages and context switches. If the TCP/IP messages are for different destinations, additional TCP/IP messages need to be set up.

- **Unique key access**

Unique key accesses are similar to primary key accesses, except that a unique key access is executed as a read on an index table followed by a primary key access on the table. However, only one request is sent from the MySQL Server, and the read of the index table is handled by `ndbd`. Such requests also benefit from batching.

- **Full table scan**

When no indexes exist for a lookup on a table, a full table scan is performed. This is sent as a single request to the `ndbd` process, which then divides the table scan into a set of parallel scans on all cluster `ndbd` processes. In future versions of MySQL Cluster, an SQL node will be able to filter some of these scans.

- **Range scan using ordered index**

When an ordered index is used, it performs a scan in the same manner as the full table scan, except that it scans only those records which are in the range used by the query transmitted by the MySQL server (SQL node). All partitions are scanned in parallel when all bound index attributes include all attributes in the partitioning key.

To check the base performance of these access methods, we have developed a set of benchmarks. One such benchmark, `testReadPerf`, tests simple and batched primary and unique key accesses. This

benchmark also measures the setup cost of range scans by issuing scans returning a single record. There is also a variant of this benchmark which uses a range scan to fetch a batch of records.

In this way, we can determine the cost of both a single key access and a single record scan access, as well as measure the impact of the communication media used, on base access methods.

In our tests, we ran the base benchmarks for both a normal transporter using TCP/IP sockets and a similar setup using SCI sockets. The figures reported in the following table are for small accesses of 20 records per access. The difference between serial and batched access decreases by a factor of 3 to 4 when using 2KB records instead. SCI Sockets were not tested with 2KB records. Tests were performed on a cluster with 2 data nodes running on 2 dual-CPU machines equipped with AMD MP1900+ processors.

Access Type	TCP/IP Sockets	SCI Socket
Serial pk access	400 microseconds	160 microseconds
Batched pk access	28 microseconds	22 microseconds
Serial uk access	500 microseconds	250 microseconds
Batched uk access	70 microseconds	36 microseconds
Indexed eq-bound	1250 microseconds	750 microseconds
Index range	24 microseconds	12 microseconds

We also performed another set of tests to check the performance of SCI Sockets *vis-à-vis* that of the SCI transporter, and both of these as compared with the TCP/IP transporter. All these tests used primary key accesses either serially and multi-threaded, or multi-threaded and batched.

The tests showed that SCI sockets were about 100% faster than TCP/IP. The SCI transporter was faster in most cases compared to SCI sockets. One notable case occurred with many threads in the test program, which showed that the SCI transporter did not perform very well when used for the `mysqld` process.

Our overall conclusion was that, for most benchmarks, using SCI sockets improves performance by approximately 100% over TCP/IP, except in rare instances when communication performance is not an issue. This can occur when scan filters make up most of processing time or when very large batches of primary key accesses are achieved. In that case, the CPU processing in the `ndbd` processes becomes a fairly large part of the overhead.

Using the SCI transporter instead of SCI Sockets is only of interest in communicating between `ndbd` processes. Using the SCI transporter is also only of interest if a CPU can be dedicated to the `ndbd` process because the SCI transporter ensures that this process will never go to sleep. It is also important to ensure that the `ndbd` process priority is set in such a way that the process does not lose priority due to running for an extended period of time, as can be done by locking processes to CPUs in Linux 2.6. If such a configuration is possible, the `ndbd` process will benefit by 10–70% as compared with using SCI sockets. (The larger figures will be seen when performing updates and probably on parallel scan operations as well.)

There are several other optimized socket implementations for computer clusters, including Myrinet, Gigabit Ethernet, Infiniband and the VIA interface. We have tested MySQL Cluster so far only with SCI sockets. See [Section 15.10.1, “Configuring MySQL Cluster to use SCI Sockets”](#) for information on how to set up SCI sockets using ordinary TCP/IP for MySQL Cluster.

## 15.11. Known Limitations of MySQL Cluster

In this section, we provide a list of known limitations in MySQL Cluster releases in the 4.1.x series compared to features available when using the `MyISAM` and `InnoDB` storage engines. Currently, there are no plans to address these in coming releases of MySQL 4.1; however, we will attempt to supply fixes for these issues in subsequent release series. If you check the “Cluster” category in the MySQL

bugs database at <http://bugs.mysql.com>, you can find known bugs which (if marked “4.1”) we intend to correct in upcoming releases of MySQL 4.1.

The list here is intended to be complete with respect to the conditions just set forth. You can report any discrepancies that you encounter to the MySQL bugs database using the instructions given in [Section 1.8, “How to Report Bugs or Problems”](#). If we do not plan to fix the problem in MySQL 4.1, we will add it to the list.

- **Noncompliance in syntax** (resulting in errors when running existing applications):
  - Temporary tables are not supported.
  - Not all charsets and collations are supported; see [Section D.5.14, “Changes in MySQL Cluster-4.1.6 \(10 October 2004\)”](#), for a list of those that are supported.
  - There are no prefix indexes; only entire fields can be indexed.
  - Text indexes are not supported. That is, you cannot create indexes on columns of any of the [TEXT](#) datatypes, nor does the [NDB](#) storage engine support [FULLTEXT](#) indexes (these are supported by [MyISAM](#) only). However, you can index [CHAR](#) or [VARCHAR](#) columns of [NDB](#) tables.
  - A [BIT](#) column cannot be a primary key, unique key, or index, nor can it be part of a composite primary key, unique key, or index.
  - Geometry datatypes ([WKT](#) and [WKB](#)) are not supported.
  - [INSERT IGNORE](#) is supported only for primary keys, but not for unique keys. One possible workaround is to remove the constraint by dropping the unique index, perform any inserts, and then add the unique index again. ([Bug#17431](#) [<http://bugs.mysql.com/17431>])
- **Non-compliance in limits or behavior** (may result in errors when running existing applications):
  - **Error Reporting:**
    - A duplicate key error returns the error message ERROR 23000: Can't write; duplicate key in table '*tbl\_name*'.
    - Like other MySQL storage engines, the [NDB](#) storage engine can handle a maximum of one [AUTO\\_INCREMENT](#) column per table. However, in the case of a Cluster table with no explicit primary key, an [AUTO\\_INCREMENT](#) column is automatically defined and used as a “hidden” primary key. For this reason, you cannot define a table that has an explicit [AUTO\\_INCREMENT](#) column unless that column is also declared using the [PRIMARY KEY](#) option.  
  
Attempting to create a table with an [AUTO\\_INCREMENT](#) column that is not the table's primary key, and using the [NDB](#) storage engine, fails with an error.
  - **Transaction Handling:**
    - [NDB Cluster](#) supports only the [READ COMMITTED](#) transaction isolation level.
    - There is no partial rollback of transactions. A duplicate key or similar error results in a rollback of the entire transaction.
    - **Important:** If a [SELECT](#) from a Cluster table includes a [BLOB](#), [TEXT](#), or [VARCHAR](#)

column, the `READ COMMITTED` transaction isolation level is converted to a read with read lock. This is done to guarantee consistency, due to the fact that parts of the values stored in columns of these types are actually read from a separate table.

- As noted elsewhere in this chapter, MySQL Cluster does not handle large transactions well; it is better to perform a number of small transactions with a few operations each than to attempt a single large transaction containing a great many operations.

Among other considerations, large transactions require very large amounts of memory. Because of this, the transactional behaviour of a number of MySQL statements is effected as described in the following list:

- `TRUNCATE` is not transactional when used on `NDB` tables. If a `TRUNCATE` fails to empty the table, then it must be re-run until it is successful.
- `DELETE FROM` (even with no `WHERE` clause) *is* transactional. For tables containing a great many rows, you may find that performance is improved by using several `DELETE FROM ... LIMIT ...` statements to “chunk” the delete operation. If the objective is to empty the table, then you may wish to use `TRUNCATE` instead.
- `LOAD DATA INFILE` is not transactional. During such an operation the `NDB` engine can and does commit at will.

`LOAD DATA FROM MASTER` is not supported in MySQL Cluster.

- When copying a table as part of an `ALTER TABLE`, the creation of the copy is non-transactional. (In any case, this operation is rolled back when the copy is deleted.)
- **Node Start, Stop, or Restart:** Starting, stopping, or restarting a node may give rise to temporary errors causing some transactions to fail. These include the following cases:
  - When first starting a node, it is possible that you may see Error 1204 Temporary failure, distribution changed and similar temporary errors.
  - The stopping or failure of any data node can result in a number of different node failure errors. (However, there should be no aborted transactions when performing a planned shutdown of the cluster.)

In either of these cases, any errors that are generated must be handled within the application. This should be done by retrying the transaction.

- A number of hard limits exist which are configurable, but available main memory in the cluster sets limits. See the complete list of configuration parameters in [Section 15.4.4, “Configuration File”](#). Most configuration parameters can be upgraded online. These hard limits include:
  - Database memory size and index memory size (`DataMemory` and `IndexMemory`, respectively).

`DataMemory` is allocated as 32KB pages. As each `DataMemory` page is used, it is assigned to a specific table; once allocated, this memory cannot be freed except by dropping the table.

See [Section 15.4.4.5, “Defining Data Nodes”](#), for further information about `DataMemory` and `IndexMemory`.

- The maximum number of operations that can be performed per transaction is set using the



configuration parameters `MaxNoOfConcurrentOperations` and `MaxNoOfLocalOperations`. Note that bulk loading, `TRUNCATE TABLE`, and `ALTER TABLE` are handled as special cases by running multiple transactions, and so are not subject to this limitation.

- Different limits related to tables and indexes. For example, the maximum number of ordered indexes per table is determined by `MaxNoOfOrderedIndexes`.

- Database names, table names and attribute names cannot be as long in `NDB` tables as with other table handlers. Attribute names are truncated to 31 characters, and if not unique after truncation give rise to errors. Database names and table names can total a maximum of 122 characters. (That is, the maximum length for an `NDB Cluster` table name is 122 characters less the number of characters in the name of the database of which that table is a part.)

- All Cluster table rows are of fixed length. This means (for example) that if a table has one or more `VARCHAR` fields containing only relatively small values, more memory and disk space is required when using the `NDB` storage engine than would be the case for the same table and data using the `MyISAM` engine. (In other words, in the case of a `VARCHAR` column, the column requires the same amount of storage as a `CHAR` column of the same size.)
- The maximum number of metadata objects is limited to 1600, including database tables, system tables, indexes and `BLOBs`.
- The maximum number of attributes per table is limited to 128.
- The maximum permitted size of any one row is 8KB. Note that each `BLOB` or `TEXT` column contributes a maximum of 256 bytes towards this total.
- The maximum number of attributes per key is 32.

- **Unsupported features** (do not cause errors, but are not supported or enforced):

- The foreign key construct is ignored, just as it is in `MyISAM` tables.
- Savepoints and rollbacks to savepoints are ignored as in `MyISAM`.
- `OPTIMIZE` operations are not supported.
- `LOAD TABLE ... FROM MASTER` is not supported.

- **Performance and limitation-related issues:**

- The query cache is disabled, since it is not invalidated if an update occurs on a different MySQL server.
- There are query performance issues due to sequential access to the `NDB` storage engine; it is also relatively more expensive to do many range scans than it is with either `MyISAM` or `InnoDB`.
- The `Records in range` statistic is not supported, resulting in non-optimal query plans in some cases. Employ `USE INDEX` or `FORCE INDEX` as a workaround.
- Unique hash indexes created with `USING HASH` cannot be used for accessing a table if `NULL` is given as part of the key.
- MySQL Cluster does not support durable commits on disk. Commits are replicated, but there is

no guarantee that logs are flushed to disk on commit.

- `SQL_LOG_BIN` has no effect on data operations; however, it is supported for schema operations.

MySQL Cluster cannot produce a binlog for tables having `BLOB` columns but no primary key.

Only the following schema operations are logged in a cluster binlog which is *not* on the `mysqld` executing the statement:

- `CREATE TABLE`
- `ALTER TABLE`
- `DROP TABLE`
- `CREATE DATABASE`
- `DROP DATABASE`
- `ndbd` searches only the default path (typically `/usr/local/mysql/share/mysql/charsets`) for character sets. Thus, it is not possible to install MySQL with Cluster support in a different path (in the case of the `.tar.gz` archives, other than `/usr/local/mysql`) if character sets that are not compiled into the MySQL Server need to be used.

• **Missing features:**

- The only supported isolation level is `READ COMMITTED`. (InnoDB supports `READ COMMITTED`, `READ UNCOMMITTED`, `REPEATABLE READ`, and `SERIALIZABLE`.) See [Section 15.8.5, “Backup Troubleshooting”](#), for information on how this can affect backup and restore of Cluster databases.
- No durable commits on disk. Commits are replicated, but there is no guarantee that logs are flushed to disk on commit.

• **Problems relating to multiple MySQL servers** (not relating to `MyISAM` or `InnoDB`):

- `ALTER TABLE` is not fully locking when running multiple MySQL servers (no distributed table lock).
- MySQL replication will not work correctly if updates are done on multiple MySQL servers. However, if the database partitioning scheme is done at the application level and no transactions take place across these partitions, replication can be made to work.
- Autodiscovery of databases is not supported for multiple MySQL servers accessing the same MySQL Cluster. However, autodiscovery of tables is supported in such cases. What this means is that after a database named `db_name` is created or imported using one MySQL server, you should issue a `CREATE DATABASE db_name` statement on each additional MySQL server that accesses the same MySQL Cluster. Once this has been done for a given MySQL server, that server should be able to detect the database tables without error.
- DDL operations are not node failure safe. If a node fails while trying to perform one of these (such as `CREATE TABLE` or `ALTER TABLE`), the data dictionary is locked and no further DDL statements can be executed without restarting the cluster.

---

• **Issues exclusive to MySQL Cluster** (not related to [MyISAM](#) or [InnoDB](#)):

- All machines used in the cluster must have the same architecture. That is, all machines hosting nodes must be either big-endian or little-endian, and you cannot use a mixture of both. For example, you cannot have a management node running on a PowerPC which directs a data node that is running on an x86 machine. This restriction does not apply to machines simply running `mysql` or other clients that may be accessing the cluster's SQL nodes.
- It is also not possible to perform a Cluster backup and restore between different architectures. For example, you cannot back up a cluster running on a big-endian platform and then restore from that backup to a cluster running on a little-endian system. ([Bug#19255](#) [<http://bugs.mysql.com/19255>])
- It is not possible to make online schema changes such as those accomplished using `ALTER TABLE` or `CREATE INDEX`, as the `NDB Cluster` engine does not support autodiscovery of such changes. (However, you can import or create a table that uses a different storage engine, and then convert it to `NDB` using `ALTER TABLE tbl_name ENGINE=NDBCLUSTER`. In such a case, you must issue a `FLUSH TABLES` statement to force the cluster to pick up the change.)
- Online adding or dropping of nodes is not possible (the cluster must be restarted in such cases).
- While it is possible to run multiple cluster processes concurrently on a single host, it is not always advisable to do so for reasons of performance and high availability, as well as other considerations. In particular, we do not in MySQL 4.1 support for production use any MySQL Cluster deployment in which more than one `ndbd` process is run on a single physical machine.

We may support multiple data nodes per host in a future MySQL release, following additional testing. However, in MySQL 4.1, such configurations can be considered experimental only.

• When using multiple management servers:

- You must give nodes explicit IDs in connectstrings because automatic allocation of node IDs does not work across multiple management servers.
- You must take extreme care to have the same configurations for all management servers. No special checks for this are performed by the cluster.
- Prior to MySQL 4.1.15, all data nodes had to be restarted after bringing up the cluster in order for the management nodes to be able to see one another.

(See [Bug#12307](#) [<http://bugs.mysql.com/12307>] and [Bug#13070](#) [<http://bugs.mysql.com/13070>] for more information.)

- Multiple network addresses per data node are not supported. Use of these is liable to cause problems: In the event of a data node failure, an SQL node waits for confirmation that the data node went down but never receives it because another route to that data node remains open. This can effectively make the cluster inoperable.

It is possible to use multiple network hardware interfaces (such as Ethernet cards) for a single data node, but these must be bound to the same address. This also means that it not possible to use more than one `[TCP]` section per connection in the `config.ini` file. See [Section 15.4.4.7, “Cluster TCP/IP Connections”](#), for more information.

- The maximum number of data nodes is 48.
- The total maximum number of nodes in a MySQL Cluster is 63. This number includes all SQL nodes (MySQL Servers), API nodes (applications accessing the cluster other than MySQL serv-

ers), data nodes, and management servers.

## 15.12. MySQL Cluster FAQ

This section answers questions that are often asked about MySQL Cluster.

- *What does “NDB” mean?*

This stands for “**N**etwork **D**atabase.”

- *What's the difference in using Cluster vs using replication?*

In a replication setup, a master MySQL server updates one or more slaves. Transactions are committed sequentially, and a slow transaction can cause the slave to lag behind the master. This means that if the master fails, it is possible that the slave might not have recorded the last few transactions. If a transaction-safe engine such as [InnoDB](#) is being used, a transaction will either be complete on the slave or not applied at all, but replication does not guarantee that all data on the master and the slave will be consistent at all times. In MySQL Cluster, all data nodes are kept in synchrony, and a transaction committed by any one data node is committed for all data nodes. In the event of a data node failure, all remaining data nodes remain in a consistent state.

In short, whereas standard MySQL replication is asynchronous, MySQL Cluster is synchronous.

We have implemented (asynchronous) replication for Cluster in MySQL 5.1. This includes the capability to replicate both between two clusters, and from a MySQL cluster to a non-Cluster MySQL server. However, we do not plan to backport this functionality to MySQL 4.1.

- *Do I need to do any special networking to run Cluster? (How do computers in a cluster communicate?)*

MySQL Cluster is intended to be used in a high-bandwidth environment, with computers connecting via TCP/IP. Its performance depends directly upon the connection speed between the cluster's computers. The minimum connectivity requirements for Cluster include a typical 100-megabit Ethernet network or the equivalent. We recommend you use gigabit Ethernet whenever available.

The faster SCI protocol is also supported, but requires special hardware. See [Section 15.10, “Using High-Speed Interconnects with MySQL Cluster”](#), for more information about SCI.

- *How many computers do I need to run a cluster, and why?*

A minimum of three computers is required to run a viable cluster. However, the minimum **recommended** number of computers in a MySQL Cluster is four: one each to run the management and SQL nodes, and two computers to serve as data nodes. The purpose of the two data nodes is to provide redundancy; the management node must run on a separate machine to guarantee continued arbitration services in the event that one of the data nodes fails.

- *What do the different computers do in a cluster?*

A MySQL Cluster has both a physical and logical organization, with computers being the physical elements. The logical or functional elements of a cluster are referred to as *nodes*, and a computer housing a cluster node is sometimes referred to as a *cluster host*. There are three types of nodes, each

corresponding to a specific role within the cluster. These are:

- **Management node (MGM node):** Provides management services for the cluster as a whole, including startup, shutdown, backups, and configuration data for the other nodes. The management node server is implemented as the application `ndb_mgmd`; the management client used to control MySQL Cluster via the MGM node is `ndb_mgm`.
- **Data node:** Stores and replicates data. Data node functionality is handled by an instance of the NDB data node process `ndbd`.
- **SQL node:** This is simply an instance of MySQL Server (`mysqld`) that is built with support for the `NDB Cluster` storage engine and started with the `--ndb-cluster` option to enable the engine.

- *With which operating systems can I use Cluster?*

MySQL Cluster is supported on most Unix-like operating systems, including Linux, Mac OS X, Solaris, BSD, HP-UX, AIX, and IRIX, among others, as well as Novell Netware. Cluster is not supported for Windows at this time. However, we are working to add Cluster support for other platforms, including Windows, and our goal is to offer MySQL Cluster on all platforms for which MySQL itself is supported.

For more detailed information concerning the level of support which is offered for MySQL Cluster on various operating system versions, OS distributions, and hardware platforms, please refer to <http://www.mysql.com/support/supportedplatforms.html>.

- *What are the hardware requirements for running MySQL Cluster?*

Cluster should run on any platform for which NDB-enabled binaries are available. Naturally, faster CPUs and more memory will improve performance, and 64-bit CPUs will likely be more effective than 32-bit processors. There must be sufficient memory on machines used for data nodes to hold each node's share of the database (see *How much RAM do I Need?* for more information). Nodes can communicate via a standard TCP/IP network and hardware. For SCI support, special networking hardware is required.

- *How much RAM do I need? Is it possible to use disk memory at all?*

In MySQL-4.1, Cluster is in-memory only. This means that all table data (including indexes) is stored in RAM. Therefore, if your data takes up 1GB of space and you want to replicate it once in the cluster, you need 2GB of memory to do so (1 GB per replica). This is in addition to the memory required by the operating system and any applications running on the cluster computers.

If a data node's memory usage exceeds what is available in RAM, then the system will attempt to use swap space up to the limit set for `DataMemory`. However, this will at best result in severely degraded performance, and may cause the node to be dropped due to slow response time (missed heartbeats). We do not recommend on relying on disk swapping in a production environment for this reason. In any case, once the `DataMemory` limit is reached, any operations requiring additional memory (such as inserts) will fail.

(We have implemented disk data storage for MySQL Cluster in MySQL 5.1, but we have no plans to add this capability in MySQL 4.1.)

You can use the following formula for obtaining a rough estimate of how much RAM is needed for each data node in the cluster:

```
(SizeofDatabase × NumberOfReplicas × 1.1 ) / NumberOfDataNodes
```

To calculate the memory requirements more exactly requires determining, for each table in the cluster database, the storage space required per row (see [Section 11.5, “Data Type Storage Requirements”](#), for details), and multiplying this by the number of rows. You must also remember to account for any column indexes as follows:

- Each primary key or hash index created for an `NDBCluster` table requires 21–25 bytes per record. These indexes use `IndexMemory`.
- Each ordered index requires 10 bytes storage per record, using `DataMemory`.
- Creating a primary key or unique index also creates an ordered index, unless this index is created with `USING HASH`. In other words:
  - A primary key or unique index on a Cluster table normally takes up 31 to 35 bytes per record.
  - However, if the primary key or unique index is created with `USING HASH`, then it requires only 21 to 25 bytes per record.

Note that creating MySQL Cluster tables with `USING HASH` for all primary keys and unique indexes will generally cause table updates to run more quickly — in some cases by as much as 20 to 30 percent faster than updates on tables where `USING HASH` was not used in creating primary and unique keys. This is due to the fact that less memory is required (because no ordered indexes are created), and that less CPU must be utilized (because fewer indexes must be read and possibly updated). However, it also means that queries that could otherwise use range scans must be satisfied by other means, which can result in slower selects.

When calculating Cluster memory requirements, you may find useful the `ndb_size.pl` utility which is available on [MySQLForge](http://forge.mysql.com/projects/view.php?id=88) [http://forge.mysql.com/projects/view.php?id=88]. This Perl script connects to a current MySQL (non-Cluster) database and creates a report on how much space that database would require if it used the `NDBCluster` storage engine. For more information, see [Section 15.9.12, “ndb\\_size.pl”](#).

It is especially important to keep in mind that *every MySQL Cluster table must have a primary key*. The `NDB` storage engine creates a primary key automatically if none is defined, and this primary key is created without `USING HASH`.

There is no easy way to determine exactly how much memory is being used for storage of Cluster indexes at any given time; however, warnings are written to the Cluster log when 80% of available `DataMemory` or `IndexMemory` is in use, and again when use reaches 85%, 90%, and so on.

We often see questions from users who report that, when they are trying to populate a Cluster database, the loading process terminates prematurely and an error message like this one is observed:

```
ERROR 1114: The table 'my_cluster_table' is full
```

When this occurs, the cause is very likely to be that your setup does not provide sufficient RAM for all table data and all indexes, *including the primary key required by the `NDB` storage engine and automatically created in the event that the table definition does not include the definition of a primary key*.

It is also worth noting that all data nodes should have the same amount of RAM, as no data node in a cluster can use more memory than the least amount available to any individual data node. In other words, if there are three computers hosting Cluster data nodes, with two of these having 3GB of RAM available to store Cluster data, and one having only 1GB RAM, then each data node can devote only 1GB to clustering.

- *Because MySQL Cluster uses TCP/IP, does that mean I can run it over the Internet, with one or more nodes in a remote location?*

It is very doubtful in any case that a cluster would perform reliably under such conditions, as MySQL Cluster was designed and implemented with the assumption that it would be run under conditions guaranteeing dedicated high-speed connectivity such as that found in a LAN setting using 100 Mbps or gigabit Ethernet (preferably the latter). We neither test nor warrant its performance using anything slower than this.

Also, it is extremely important to keep in mind that communications between the nodes in a MySQL Cluster are not secure; they are neither encrypted nor safeguarded by any other protective mechanism. The most secure configuration for a cluster is in a private network behind a firewall, with no direct access to any Cluster data or management nodes from outside. (For SQL nodes, you should take the same precautions as you would with any other instance of the MySQL server.)

- *Do I have to learn a new programming or query language to use Cluster?*

No. Although some specialized commands are used to manage and configure the cluster itself, only standard (My)SQL queries and commands are required for the following operations:

- Creating, altering, and dropping tables
- Inserting, updating, and deleting table data
- Creating, changing, and dropping primary and unique indexes
- Configuring and managing SQL nodes (MySQL servers)

- *How do I find out what an error or warning message means when using Cluster?*

There are two ways in which this can be done:

- From within the `mysql` client, use `SHOW ERRORS` or `SHOW WARNINGS` immediately upon being notified of the error or warning condition. Errors and warnings also be displayed in MySQL Query Browser.
- From a system shell prompt, use `pererror --ndb error_code`.

- *Is MySQL Cluster transaction-safe? What isolation levels are supported?*

*Yes:* For tables created with the `NDB` storage engine, transactions are supported. In MySQL 4.1, Cluster supports only the `READ COMMITTED` transaction isolation level.

- *What storage engines are supported by MySQL Cluster?*

Clustering in MySQL is supported only by the `NDB` storage engine. That is, in order for a table to be shared between nodes in a cluster, it must be created using `ENGINE=NDB` (or `ENGINE=NDBCLUSTER`, which is equivalent).

It is possible to create tables using other storage engines (such as `MyISAM` or `InnoDB`) on a MySQL server being used for clustering, but these non-`NDB` tables will **not** participate in the cluster; they are local to the individual MySQL server instance on which they are created.

- *Which versions of the MySQL software support Cluster? Do I have to compile from source?*

Cluster is supported in the MySQL-max binaries from version 4.1.3 onward. You can determine whether your server has NDB support using either the `SHOW VARIABLES LIKE 'have_%%'` or `SHOW ENGINES` statement. (See [Section 5.3](#), “The `mysqld-max` Extended MySQL Server”, for more information.)

Linux users, please note that NDB is *not* included in the 4.1.x RPMs; you should use the binaries supplied as `.tar.gz` archives from <http://dev.mysql.com/downloads/> instead. You can also obtain NDB support by compiling the `-max` binaries from source, but it is not necessary to do so simply to use MySQL Cluster. To download the latest binary, RPM, or source distribution in the MySQL 4.1 series, visit <http://dev.mysql.com/downloads/mysql/4.1.html>.

- *In the event of a catastrophic failure — say, for instance, the whole city loses power **and** my UPS fails — would I lose all my data?*

All committed transactions are logged. Therefore, although it is possible that some data could be lost in the event of a catastrophe, this should be quite limited. Data loss can be further reduced by minimizing the number of operations per transaction. (It is not a good idea to perform large numbers of operations per transaction in any case.)

- *Is it possible to use `FULLTEXT` indexes with Cluster?*

`FULLTEXT` indexing is not currently supported by the NDB storage engine, or by any storage engine other than `MyISAM`. We are working to add this capability in a future release.

- *Can I run multiple nodes on a single computer?*

It is possible but not advisable. One of the chief reasons to run a cluster is to provide redundancy. To enjoy the full benefits of this redundancy, each node should reside on a separate machine. If you place multiple nodes on a single machine and that machine fails, you lose all of those nodes. Given that MySQL Cluster can be run on commodity hardware loaded with a low-cost (or even no-cost) operating system, the expense of an extra machine or two is well worth it to safeguard mission-critical data. It also worth noting that the requirements for a cluster host running a management node are minimal. This task can be accomplished with a 200 MHz Pentium CPU and sufficient RAM for the operating system plus a small amount of overhead for the `ndb_mgmd` and `ndb_mgm` processes.

It is acceptable to run multiple cluster data nodes on a single host for learning about MySQL Cluster, or for testing purposes; however, this is not supported for production use.

- *Can I add nodes to a cluster without restarting it?*

Not at present. A simple restart is all that is required for adding new MGM or SQL nodes to a Cluster. When adding data nodes the process is more complex, and requires the following steps:

1. Make a complete backup of all Cluster data.
2. Completely shut down the cluster and all cluster node processes.
3. Restart the cluster, using the `--initial` startup option.
4. Restore all cluster data from the backup.

In a future MySQL Cluster release series, we hope to implement a “hot” reconfiguration capability for MySQL Cluster to minimize (if not eliminate) requirements for restarting the cluster when adding new nodes. However, this is not planned for the MySQL 4.1 release series.

- *Are there any limitations that I should be aware of when using Cluster?*

NDB tables in MySQL 4.1 are subject to the following limitations:



- Temporary tables are not supported; a `CREATE TEMPORARY TABLE` statement using `ENGINE=NDB` or `ENGINE=NDBCLUSTER` fails with an error.
- Not all character sets and collations are supported. (For a complete listing of those that are supported, see [Section D.5.14, “Changes in MySQL Cluster-4.1.6 \(10 October 2004\)”](#)).
- `FULLTEXT` indexes and index prefixes are not supported. Only complete columns may be indexed.
- Spatial data types are not supported. See [Chapter 16, \*Spatial Extensions\*](#).
- Only complete rollbacks for transactions are supported. Partial rollbacks and rollbacks to savepoints are not supported.
- The maximum number of attributes allowed per table is 128, and attribute names cannot be any longer than 31 characters. For each table, the maximum combined length of the table and database names is 122 characters.
- The maximum size for a table row is 8 kilobytes, not counting `BLOB` values. There is no set limit for the number of rows per table. Table size limits depend on a number of factors, in particular on the amount of RAM available to each data node.
- The `NDB` engine does not support foreign key constraints. As with `MyISAM` tables, these are ignored.
- Query caching is not supported.

For additional information on Cluster limitations, see [Section 15.11, “Known Limitations of MySQL Cluster”](#).

- *How do I import an existing MySQL database into a cluster?*

You can import databases into MySQL Cluster much as you would with any other version of MySQL. Other than the limitation mentioned in the previous question, the only other special requirement is that any tables to be included in the cluster must use the `NDB` storage engine. This means that the tables must be created with `ENGINE=NDB` or `ENGINE=NDBCLUSTER`. It is also possible to convert existing tables using other storage engines to `NDB Cluster` using `ALTER TABLE`, but requires an additional workaround. See [Section 15.11, “Known Limitations of MySQL Cluster”](#), for details.

- *How do cluster nodes communicate with one another?*

Cluster nodes can communicate via any of three different protocols: TCP/IP, SHM (shared memory), and SCI (Scalable Coherent Interface). Where available, SHM is used by default between nodes residing on the same cluster host. SCI is a high-speed (1 gigabit per second and higher), high-availability protocol used in building scalable multi-processor systems; it requires special hardware and drivers. See [Section 15.10, “Using High-Speed Interconnects with MySQL Cluster”](#), for more about using SCI as a transport mechanism in MySQL Cluster.

- *What is an “arbiter”?*

If one or more nodes in a cluster fail, it is possible that not all cluster nodes will be able to “see” one another. In fact, it is possible that two sets of nodes might become isolated from one another in a network partitioning, also known as a “split brain” scenario. This type of situation is undesirable because each set of nodes tries to behave as though it is the entire cluster.

When cluster nodes go down, there are two possibilities. If more than 50% of the remaining nodes can communicate with each other, we have what is sometimes called a “majority rules” situation, and this set of nodes is considered to be the cluster. The arbitrator comes into play when there is an even number of nodes: in such cases, the set of nodes to which the arbitrator belongs is considered to be the cluster, and nodes not belonging to this set are shut down.

The preceding information is somewhat simplified. A more complete explanation taking into account node groups follows:

When all nodes in at least one node group are alive, network partitioning is not an issue, because no one portion of the cluster can form a functional cluster. The real problem arises when no single node group has all its nodes alive, in which case network partitioning (the “split-brain” scenario) becomes possible. Then an arbitrator is required. All cluster nodes recognize the same node as the arbitrator, which is normally the management server; however, it is possible to configure any of the MySQL Servers in the cluster to act as the arbitrator instead. The arbitrator accepts the first set of cluster nodes to contact it, and tells the remaining set to shut down. Arbitrator selection is controlled by the `ArbitrationRank` configuration parameter for MySQL Server and management server nodes. (See [Section 15.4.4.4, “Defining the Management Server”](#), for details.) It should also be noted that the role of arbitrator does not in and of itself impose any heavy demands upon the host so designated, and thus the arbitrator host does not need to be particularly fast or to have extra memory especially for this purpose.

- *What data types are supported by MySQL Cluster?*

MySQL Cluster supports all of the usual MySQL data types, with the exception of those associated with MySQL's spatial extensions. (See [Chapter 16, \*Spatial Extensions\*](#).) In addition, there are some differences with regard to indexes when used with NDB tables. **Note:** MySQL Cluster tables (that is, tables created with `ENGINE=NDBCLUSTER`) have only fixed-width rows. This means that (for example) each record containing a `VARCHAR(255)` column will require space for 255 characters (as required for the character set and collation being used for the table), regardless of the actual number of characters stored therein. This issue is expected to be fixed in a future MySQL release series.

See [Section 15.11, “Known Limitations of MySQL Cluster”](#), for more information about these issues.

- *How do I start and stop MySQL Cluster?*

It is necessary to start each node in the cluster separately, in the following order:

1. Start the management node with the `ndb_mgmd` command.
2. Start each data node with the `ndbd` command.
3. Start each MySQL server (SQL node) using `mysqld_safe --user=mysql &`.

Each of these commands must be run from a system shell on the machine housing the affected node. You can verify the cluster is running by starting the MGM management client `ndb_mgm` on the machine housing the MGM node.

- *What happens to cluster data when the cluster is shut down?*

The data held in memory by the cluster's data nodes is written to disk, and is reloaded in memory the next time that the cluster is started.

To shut down the cluster, enter the following command in a shell on the machine hosting the MGM node:

```
shell> ndb_mgm -e shutdown
```

This causes the `ndb_mgm`, `ndb_mgm`, and any `ndbd` processes to terminate gracefully. MySQL servers running as Cluster SQL nodes can be stopped using `mysqladmin shutdown`.

For more information, see [Section 15.7.2, “Commands in the Management Client”](#), and [Section 15.3.6, “Safe Shutdown and Restart”](#).

- *Is it helpful to have more than one management node for a cluster?*

It can be helpful as a fail-safe. Only one MGM node controls the cluster at any given time, but it is possible to configure one MGM as primary, and one or more additional management nodes to take over in the event that the primary MGM node fails.

- *Can I mix different kinds of hardware and operating systems in a Cluster?*

Yes, so long as all machines and operating systems have the same endianness (all big-endian or all little-endian). It is also possible to use different MySQL Cluster releases on different nodes. However, we recommend this be done only as part of a rolling upgrade procedure.

- *Can I run two data nodes on a single host? Two SQL nodes?*

Yes, it is possible to do this. In the case of multiple data nodes, each node must use a different data directory. If you want to run multiple SQL nodes on one machine, each instance of `mysqld` must use a different TCP/IP port. However, running more than one node of a given type per machine is not supported for production use.

- *Can I use hostnames with MySQL Cluster?*

Yes, it is possible to use DNS and DHCP for cluster hosts. However, if your application requires “five nines” availability, we recommend using fixed IP addresses. Making communication between Cluster hosts dependent on services such as DNS and DHCP introduces additional points of failure, and the fewer of these, the better.

## 15.13. MySQL Cluster Glossary

The following terms are useful to an understanding of MySQL Cluster or have specialized meanings when used in relation to it.

- **Cluster:**

In its generic sense, a cluster is a set of computers functioning as a unit and working together to accomplish a single task.

### **NDB Cluster:**

This is the storage engine used in MySQL to implement data storage, retrieval, and management distributed among several computers.

### **MySQL Cluster:**

This refers to a group of computers working together using the **NDB** storage engine to support a distributed MySQL database in a *shared-nothing architecture* using *in-memory storage*.

- **Configuration files:**

Text files containing directives and information regarding the cluster, its hosts, and its nodes. These are read by the cluster's management nodes when the cluster is started. See [Section 15.4.4, “Configuration File”](#), for details.

- **Backup:**

A complete copy of all cluster data, transactions and logs, saved to disk or other long-term storage.

- **Restore:**

Returning the cluster to a previous state, as stored in a backup.

- **Checkpoint:**

Generally speaking, when data is saved to disk, it is said that a checkpoint has been reached. More specific to Cluster, it is a point in time where all committed transactions are stored on disk. With regard to the **NDB** storage engine, there are two types of checkpoints which work together to ensure that a consistent view of the cluster's data is maintained:

- **Local Checkpoint (LCP):**

This is a checkpoint that is specific to a single node; however, LCP's take place for all nodes in the cluster more or less concurrently. An LCP involves saving all of a node's data to disk, and so usually occurs every few minutes. The precise interval varies, and depends upon the amount of data stored by the node, the level of cluster activity, and other factors.

- **Global Checkpoint (GCP):**

A GCP occurs every few seconds, when transactions for all nodes are synchronized and the redo-log is flushed to disk.

- **Cluster host:**

A computer making up part of a MySQL Cluster. A cluster has both a *physical* structure and a *logical* structure. Physically, the cluster consists of a number of computers, known as *cluster hosts* (or more simply as *hosts*). See also **Node** and **Node group** below.

- **Node:**

This refers to a logical or functional unit of MySQL Cluster, and is sometimes also referred to as a *cluster node*. In the context of MySQL Cluster, we use the term “node” to indicate a *process* rather than a physical component of the cluster. There are three node types required to implement a working MySQL Cluster:

- **Management (MGM) nodes:**

Manages the other nodes within the MySQL Cluster. It provides configuration data to the other nodes; starts and stops nodes; handles network partitioning; creates backups and restores from them, and so forth.

- **SQL (MySQL server) nodes:**

Instances of MySQL Server which serve as front ends to data kept in the cluster's **data nodes**. Clients desiring to store, retrieve, or update data can access an SQL node just as they would any other MySQL Server, employing the usual authentication methods and API's; the underlying distribution of data between node groups is transparent to users and applications. SQL nodes access the cluster's databases as a whole without regard to the data's distribution across different data nodes or cluster hosts.

- **Data nodes:**

These nodes store the actual data. Table data fragments are stored in a set of node groups; each node group stores a different subset of the table data. Each of the nodes making up a node group stores a replica of the fragment for which that node group is responsible. Currently, a single cluster can support up to 48 data nodes total.

It is possible for more than one node to co-exist on a single machine. (In fact, it is even possible to set up a complete cluster on one machine, although one would almost certainly *not* want to do this in a production environment.) It may be helpful to remember that, when working with MySQL Cluster, the term *host* refers to a physical component of the cluster whereas a *node* is a logical or functional component (that is, a process).

**Note Regarding Terms:** In older versions of the MySQL Cluster documentation, data nodes were sometimes referred to as “database nodes”. The term “storage nodes” has also been used. In addition, SQL nodes were sometimes known as “client nodes”. This older terminology has been deprecated to minimize confusion, and for this reason should be avoided. They are also often referred to as “API nodes” — an SQL node is actually an API node that provides an SQL interface to the cluster.

- **Node group:**

A set of data nodes. All data nodes in a node group contain the same data (fragments), and all nodes in a single group should reside on different hosts. It is possible to control which nodes belong to which node groups.

For more information, see [Section 15.2.1, “MySQL Cluster Nodes, Node Groups, Replicas, and Partitions”](#).

- **Node failure:**

MySQL Cluster is not solely dependent upon the functioning of any single node making up the cluster; the cluster can continue to run if one or more nodes fail. The precise number of node failures that a given cluster can tolerate depends upon the number of nodes and the cluster's configuration.

- **Node restart:**

The process of restarting a failed cluster node.

- **Initial node restart:**

The process of starting a cluster node with its filesystem removed. This is sometimes used in the course of software upgrades and in other special circumstances.

- **System crash (or system failure):**

This can occur when so many cluster nodes have failed that the cluster's state can no longer be guaranteed.

- **System restart:**

The process of restarting the cluster and reinitializing its state from disk logs and checkpoints. This is required after either a planned or an unplanned shutdown of the cluster.

- **Fragment:**

A portion of a database table; in the [NDB](#) storage engine, a table is broken up into and stored as a number of fragments. A fragment is sometimes also called a “partition”; however, “fragment” is the preferred term. Tables are fragmented in MySQL Cluster in order to facilitate load balancing

between machines and nodes.

- **Replica:**

Under the [NDB](#) storage engine, each table fragment has number of replicas stored on other data nodes in order to provide redundancy. Currently, there may be up four replicas per fragment.

- **Transporter:**

A protocol providing data transfer between nodes. MySQL Cluster currently supports four different types of transporter connections:

- **TCP/IP**

This is, of course, the familiar network protocol that underlies HTTP, FTP (and so on) on the Internet. TCP/IP can be used for both local and remote connections.

- **SCI**

**Scalable Coherent Interface** is a high-speed protocol used in building multiprocessor systems and parallel-processing applications. Use of SCI with MySQL Cluster requires specialized hardware, as discussed in [Section 15.10.1, “Configuring MySQL Cluster to use SCI Sockets”](#). For a basic introduction to SCI, see [this essay at dolphinics.com](http://www.dolphinics.com) [<http://www.dolphinics.com/corporate/scitech.html>].

- **SHM**

Unix-style **shared memory** segments. Where supported, SHM is used automatically to connect nodes running on the same host. The [Unix man page for shmop\(2\)](#) [<http://www.scit.wlv.ac.uk/cgi-bin/mansec?2+shmop>] is a good place to begin obtaining additional information about this topic.

**Note:** The cluster transporter is internal to the cluster. Applications using MySQL Cluster communicate with SQL nodes just as they do with any other version of MySQL Server (via TCP/IP, or through the use of Unix socket files or Windows named pipes). Queries can be sent and results retrieved using the standard MySQL client APIs.

- **NDB:**

This stands for **Network Database**, and refers to the storage engine used to enable MySQL Cluster. The [NDB](#) storage engine supports all the usual MySQL data types and SQL statements, and is ACID-compliant. This engine also provides full support for transactions (commits and rollbacks).

- **shared-nothing architecture:**

The ideal architecture for a MySQL Cluster. In a true shared-nothing setup, each node runs on a separate host. The advantage such an arrangement is that there no single host or node can act as single point of failure or as a performance bottle neck for the system as a whole.

- **In-memory storage:**

All data stored in each data node is kept in memory on the node's host computer. For each data node in the cluster, you must have available an amount of RAM equal to the size of the database times the number of replicas, divided by the number of data nodes. Thus, if the database takes up 1GB of memory, and you want to set up the cluster with four replicas and eight data nodes, a minimum of 500MB memory will be required per node. Note that this is in addition to any requirements for the operating system and any other applications that might be running on the host.

- **Table:**

As is usual in the context of a relational database, the term “table” denotes a set of identically structured records. In MySQL Cluster, a database table is stored in a data node as a set of fragments, each of which is replicated on additional data nodes. The set of data nodes replicating the same fragment or set of fragments is referred to as a *node group*.

- **Cluster programs:**

These are command-line programs used in running, configuring, and administering MySQL Cluster. They include both server daemons:

- `ndbd`:

The data node daemon (runs a data node process)

- `ndb_mgmd`:

The management server daemon (runs a management server process)

and client programs:

- `ndb_mgm`:

The management client (provides an interface for executing management commands)

- `ndb_waiter`:

Used to verify status of all nodes in a cluster

- `ndb_restore`:

Restores cluster data from backup

For more about these programs and their uses, see [Section 15.6, “Process Management in MySQL Cluster”](#).

- **Event log:**

MySQL Cluster logs events by category (startup, shutdown, errors, checkpoints, and so on), priority, and severity. A complete listing of all reportable events may be found in [Section 15.7.3, “Event Reports Generated in MySQL Cluster”](#). Event logs are of two types:

- **Cluster log:**

Keeps a record of all desired reportable events for the cluster as a whole.

- **Node log:**

A separate log which is also kept for each individual node.

Under normal circumstances, it is necessary and sufficient to keep and examine only the cluster log. The node logs need be consulted only for application development and debugging purposes.

---

# Chapter 16. Spatial Extensions

MySQL 4.1 introduces spatial extensions to allow the generation, storage, and analysis of geographic features. Currently, these features are available for [MyISAM](#) tables only.

This chapter covers the following topics:

- The basis of these spatial extensions in the OpenGIS geometry model
- Data formats for representing spatial data
- How to use spatial data in MySQL
- Use of indexing for spatial data
- MySQL differences from the OpenGIS specification

## Additional resources

- The Open Geospatial Consortium publishes the *OpenGIS® Simple Features Specifications For SQL*, a document that proposes several conceptual ways for extending an SQL RDBMS to support spatial data. This specification is available from the OGC Web site at <http://www.opengis.org/docs/99-049.pdf>.
- If you have questions or concerns about the use of the spatial extensions to MySQL, you can discuss them in the GIS forum: <http://forums.mysql.com/list.php?23>.

## 16.1. Introduction to MySQL Spatial Support

MySQL implements spatial extensions following the specification of the Open Geospatial Consortium (OGC). This is an international consortium of more than 250 companies, agencies, and universities participating in the development of publicly available conceptual solutions that can be useful with all kinds of applications that manage spatial data. The OGC maintains a Web site at <http://www.opengis.org/>.

In 1997, the Open Geospatial Consortium published the *OpenGIS® Simple Features Specifications For SQL*, a document that proposes several conceptual ways for extending an SQL RDBMS to support spatial data. This specification is available from the OGC Web site at <http://www.opengis.org/docs/99-049.pdf>. It contains additional information relevant to this chapter.

MySQL implements a subset of the **SQL with Geometry Types** environment proposed by OGC. This term refers to an SQL environment that has been extended with a set of geometry types. A geometry-valued SQL column is implemented as a column that has a geometry type. The specification describe a set of SQL geometry types, as well as functions on those types to create and analyze geometry values.

A **geographic feature** is anything in the world that has a location. A feature can be:

- An entity. For example, a mountain, a pond, a city.
- A space. For example, a postcode area, the tropics.
- A definable location. For example, a crossroad, as a particular place where two streets intersect.



Some documents use the term **geospatial feature** to refer to geographic features.

**Geometry** is another word that denotes a geographic feature. Originally the word **geometry** meant measurement of the earth. Another meaning comes from cartography, referring to the geometric features that cartographers use to map the world.

This chapter uses all of these terms synonymously: **geographic feature**, **geospatial feature**, **feature**, or **geometry**. Here, the term most commonly used is **geometry**, defined as *a point or an aggregate of points representing anything in the world that has a location*.

## 16.2. The OpenGIS Geometry Model

The set of geometry types proposed by OGC's **SQL with Geometry Types** environment is based on the **OpenGIS Geometry Model**. In this model, each geometric object has the following general properties:

- It is associated with a Spatial Reference System, which describes the coordinate space in which the object is defined.
- It belongs to some geometry class.

### 16.2.1. The Geometry Class Hierarchy

The geometry classes define a hierarchy as follows:

- `Geometry` (non-instantiable)
  - `Point` (instantiable)
  - `Curve` (non-instantiable)
    - `LineString` (instantiable)
      - `Line`
      - `LinearRing`
  - `Surface` (non-instantiable)
    - `Polygon` (instantiable)
- `GeometryCollection` (instantiable)
  - `MultiPoint` (instantiable)
  - `MultiCurve` (non-instantiable)
    - `MultiLineString` (instantiable)
  - `MultiSurface` (non-instantiable)
    - `MultiPolygon` (instantiable)

It is not possible to create objects in non-instantiable classes. It is possible to create objects in instantiable classes. All classes have properties, and instantiable classes may also have assertions (rules that

define valid class instances).

`Geometry` is the base class. It is an abstract class. The instantiable subclasses of `Geometry` are restricted to zero-, one-, and two-dimensional geometric objects that exist in two-dimensional coordinate space. All instantiable geometry classes are defined so that valid instances of a geometry class are topologically closed (that is, all defined geometries include their boundary).

The base `Geometry` class has subclasses for `Point`, `Curve`, `Surface`, and `GeometryCollection`:

- `Point` represents zero-dimensional objects.
- `Curve` represents one-dimensional objects, and has subclass `LineString`, with sub-subclasses `Line` and `LinearRing`.
- `Surface` is designed for two-dimensional objects and has subclass `Polygon`.
- `GeometryCollection` has specialized zero-, one-, and two-dimensional collection classes named `MultiPoint`, `MultiLineString`, and `MultiPolygon` for modeling geometries corresponding to collections of `Points`, `LineStrings`, and `Polygons`, respectively. `MultiCurve` and `MultiSurface` are introduced as abstract superclasses that generalize the collection interfaces to handle `Curves` and `Surfaces`.

`Geometry`, `Curve`, `Surface`, `MultiCurve`, and `MultiSurface` are defined as non-instantiable classes. They define a common set of methods for their subclasses and are included for extensibility.

`Point`, `LineString`, `Polygon`, `GeometryCollection`, `MultiPoint`, `MultiLineString`, and `MultiPolygon` are instantiable classes.

## 16.2.2. Class `Geometry`

`Geometry` is the root class of the hierarchy. It is a non-instantiable class but has a number of properties that are common to all geometry values created from any of the `Geometry` subclasses. These properties are described in the following list. Particular subclasses have their own specific properties, described later.

### Geometry Properties

A geometry value has the following properties:

- Its **type**. Each geometry belongs to one of the instantiable classes in the hierarchy.
- Its **SRID**, or Spatial Reference Identifier. This value identifies the geometry's associated Spatial Reference System that describes the coordinate space in which the geometry object is defined.

In MySQL, the SRID value is just an integer associated with the geometry value. All calculations are done assuming Euclidean (planar) geometry.

- Its **coordinates** in its Spatial Reference System, represented as double-precision (eight-byte) numbers. All non-empty geometries include at least one pair of (X,Y) coordinates. Empty geometries contain no coordinates.

Coordinates are related to the SRID. For example, in different coordinate systems, the distance between two objects may differ even when objects have the same coordinates, because the distance on the **planar** coordinate system and the distance on the **geocentric** system (coordinates on the Earth's surface) are different things.

- Its **interior**, **boundary**, and **exterior**.

Every geometry occupies some position in space. The exterior of a geometry is all space not occupied by the geometry. The interior is the space occupied by the geometry. The boundary is the interface between the geometry's interior and exterior.

- Its **MBR** (Minimum Bounding Rectangle), or Envelope. This is the bounding geometry, formed by the minimum and maximum (X,Y) coordinates:

```
((MINX MINY, MAXX MINY, MAXX MAXY, MINX MAXY, MINX MINY))
```

- Whether the value is **simple** or **non-simple**. Geometry values of types ([LineString](#), [MultiPoint](#), [MultiLineString](#)) are either simple or non-simple. Each type determines its own assertions for being simple or non-simple.
- Whether the value is **closed** or **not closed**. Geometry values of types ([LineString](#), [MultiString](#)) are either closed or not closed. Each type determines its own assertions for being closed or not closed.
- Whether the value is **empty** or **non-empty**. A geometry is empty if it does not have any points. Exterior, interior, and boundary of an empty geometry are not defined (that is, they are represented by a [NULL](#) value). An empty geometry is defined to be always simple and has an area of 0.
- Its **dimension**. A geometry can have a dimension of -1, 0, 1, or 2:
  - -1 for an empty geometry.
  - 0 for a geometry with no length and no area.
  - 1 for a geometry with non-zero length and zero area.
  - 2 for a geometry with non-zero area.

[Point](#) objects have a dimension of zero. [LineString](#) objects have a dimension of 1. [Polygon](#) objects have a dimension of 2. The dimensions of [MultiPoint](#), [MultiLineString](#), and [MultiPolygon](#) objects are the same as the dimensions of the elements they consist of.

## 16.2.3. Class [Point](#)

A [Point](#) is a geometry that represents a single location in coordinate space.

### [Point](#) Examples

- Imagine a large-scale map of the world with many cities. A [Point](#) object could represent each city.
- On a city map, a [Point](#) object could represent a bus stop.

### [Point](#) Properties

- X-coordinate value.
- Y-coordinate value.
- [Point](#) is defined as a zero-dimensional geometry.

- The boundary of a `Point` is the empty set.

## 16.2.4. Class `Curve`

A `Curve` is a one-dimensional geometry, usually represented by a sequence of points. Particular subclasses of `Curve` define the type of interpolation between points. `Curve` is a non-instantiable class.

### `Curve` Properties

- A `Curve` has the coordinates of its points.
- A `Curve` is defined as a one-dimensional geometry.
- A `Curve` is simple if it does not pass through the same point twice.
- A `Curve` is closed if its start point is equal to its endpoint.
- The boundary of a closed `Curve` is empty.
- The boundary of a non-closed `Curve` consists of its two endpoints.
- A `Curve` that is simple and closed is a `LinearRing`.

## 16.2.5. Class `LineString`

A `LineString` is a `Curve` with linear interpolation between points.

### `LineString` Examples

- On a world map, `LineString` objects could represent rivers.
- In a city map, `LineString` objects could represent streets.

### `LineString` Properties

- A `LineString` has coordinates of segments, defined by each consecutive pair of points.
- A `LineString` is a `Line` if it consists of exactly two points.
- A `LineString` is a `LinearRing` if it is both closed and simple.

## 16.2.6. Class `Surface`

A `Surface` is a two-dimensional geometry. It is a non-instantiable class. Its only instantiable subclass is `Polygon`.

### `Surface` Properties

- A `Surface` is defined as a two-dimensional geometry.

- The OpenGIS specification defines a simple [Surface](#) as a geometry that consists of a single “patch” that is associated with a single exterior boundary and zero or more interior boundaries.
- The boundary of a simple [Surface](#) is the set of closed curves corresponding to its exterior and interior boundaries.

### 16.2.7. Class [Polygon](#)

A [Polygon](#) is a planar [Surface](#) representing a multisided geometry. It is defined by a single exterior boundary and zero or more interior boundaries, where each interior boundary defines a hole in the [Polygon](#).

#### [Polygon](#) Examples

- On a region map, [Polygon](#) objects could represent forests, districts, and so on.

#### [Polygon](#) Assertions

- The boundary of a [Polygon](#) consists of a set of [LinearRing](#) objects (that is, [LineString](#) objects that are both simple and closed) that make up its exterior and interior boundaries.
- A [Polygon](#) has no rings that cross. The rings in the boundary of a [Polygon](#) may intersect at a [Point](#), but only as a tangent.
- A [Polygon](#) has no lines, spikes, or punctures.
- A [Polygon](#) has an interior that is a connected point set.
- A [Polygon](#) may have holes. The exterior of a [Polygon](#) with holes is not connected. Each hole defines a connected component of the exterior.

The preceding assertions make a [Polygon](#) a simple geometry.

### 16.2.8. Class [GeometryCollection](#)

A [GeometryCollection](#) is a geometry that is a collection of one or more geometries of any class.

All the elements in a [GeometryCollection](#) must be in the same Spatial Reference System (that is, in the same coordinate system). There are no other constraints on the elements of a [GeometryCollection](#), although the subclasses of [GeometryCollection](#) described in the following sections may restrict membership. Restrictions may be based on:

- Element type (for example, a [MultiPoint](#) may contain only [Point](#) elements)
- Dimension
- Constraints on the degree of spatial overlap between elements

### 16.2.9. Class [MultiPoint](#)

A [MultiPoint](#) is a geometry collection composed of [Point](#) elements. The points are not connected

or ordered in any way.

#### **MultiPoint Examples**

- On a world map, a `MultiPoint` could represent a chain of small islands.
- On a city map, a `MultiPoint` could represent the outlets for a ticket office.

#### **MultiPoint Properties**

- A `MultiPoint` is a zero-dimensional geometry.
- A `MultiPoint` is simple if no two of its `Point` values are equal (have identical coordinate values).
- The boundary of a `MultiPoint` is the empty set.

### 16.2.10. Class `MultiCurve`

A `MultiCurve` is a geometry collection composed of `Curve` elements. `MultiCurve` is a non-instantiable class.

#### **MultiCurve Properties**

- A `MultiCurve` is a one-dimensional geometry.
- A `MultiCurve` is simple if and only if all of its elements are simple; the only intersections between any two elements occur at points that are on the boundaries of both elements.
- A `MultiCurve` boundary is obtained by applying the “mod 2 union rule” (also known as the “odd-even rule”): A point is in the boundary of a `MultiCurve` if it is in the boundaries of an odd number of `MultiCurve` elements.
- A `MultiCurve` is closed if all of its elements are closed.
- The boundary of a closed `MultiCurve` is always empty.

### 16.2.11. Class `MultiLineString`

A `MultiLineString` is a `MultiCurve` geometry collection composed of `LineString` elements.

#### **MultiLineString Examples**

- On a region map, a `MultiLineString` could represent a river system or a highway system.

### 16.2.12. Class `MultiSurface`

A `MultiSurface` is a geometry collection composed of surface elements. `MultiSurface` is a non-instantiable class. Its only instantiable subclass is `MultiPolygon`.

#### **MultiSurface Assertions**

- Two `MultiSurface` surfaces have no interiors that intersect.
- Two `MultiSurface` elements have boundaries that intersect at most at a finite number of points.

### 16.2.13. Class `MultiPolygon`

A `MultiPolygon` is a `MultiSurface` object composed of `Polygon` elements.

#### `MultiPolygon` Examples

- On a region map, a `MultiPolygon` could represent a system of lakes.

#### `MultiPolygon` Assertions

- A `MultiPolygon` has no two `Polygon` elements with interiors that intersect.
- A `MultiPolygon` has no two `Polygon` elements that cross (crossing is also forbidden by the previous assertion), or that touch at an infinite number of points.
- A `MultiPolygon` may not have cut lines, spikes, or punctures. A `MultiPolygon` is a regular, closed point set.
- A `MultiPolygon` that has more than one `Polygon` has an interior that is not connected. The number of connected components of the interior of a `MultiPolygon` is equal to the number of `Polygon` values in the `MultiPolygon`.

#### `MultiPolygon` Properties

- A `MultiPolygon` is a two-dimensional geometry.
- A `MultiPolygon` boundary is a set of closed curves (`LineString` values) corresponding to the boundaries of its `Polygon` elements.
- Each `Curve` in the boundary of the `MultiPolygon` is in the boundary of exactly one `Polygon` element.
- Every `Curve` in the boundary of an `Polygon` element is in the boundary of the `MultiPolygon`.

## 16.3. Supported Spatial Data Formats

This section describes the standard spatial data formats that are used to represent geometry objects in queries. They are:

- Well-Known Text (WKT) format
- Well-Known Binary (WKB) format

Internally, MySQL stores geometry values in a format that is not identical to either WKT or WKB format.

### 16.3.1. Well-Known Text (WKT) Format

The Well-Known Text (WKT) representation of Geometry is designed to exchange geometry data in ASCII form.

Examples of WKT representations of geometry objects:

- A [Point](#):

```
POINT(15 20)
```

Note that point coordinates are specified with no separating comma.

- A [LineString](#) with four points:

```
LINESTRING(0 0, 10 10, 20 25, 50 60)
```

Note that point coordinate pairs are separated by commas.

- A [Polygon](#) with one exterior ring and one interior ring:

```
POLYGON((0 0,10 0,10 10,0 10,0 0),(5 5,7 5,7 7,5 7, 5 5))
```

- A [MultiPoint](#) with three [Point](#) values:

```
MULTIPOINT(0 0, 20 20, 60 60)
```

- A [MultiLineString](#) with two [LineString](#) values:

```
MULTILINESTRING((10 10, 20 20), (15 15, 30 15))
```

- A [MultiPolygon](#) with two [Polygon](#) values:

```
MULTIPOLYGON(((0 0,10 0,10 10,0 10,0 0)),((5 5,7 5,7 7,5 7, 5 5)))
```

- A [GeometryCollection](#) consisting of two [Point](#) values and one [LineString](#):

```
GEOMETRYCOLLECTION(POINT(10 10), POINT(30 30), LINESTRING(15 15, 20 20))
```

A Backus-Naur grammar that specifies the formal production rules for writing WKT values can be found in the OpenGIS specification document referenced near the beginning of this chapter.

### 16.3.2. Well-Known Binary (WKB) Format

The Well-Known Binary (WKB) representation for geometric values is defined by the OpenGIS specification. It is also defined in the ISO *SQL/MM Part 3: Spatial* standard.

WKB is used to exchange geometry data as binary streams represented by [BLOB](#) values containing geometric WKB information.

WKB uses one-byte unsigned integers, four-byte unsigned integers, and eight-byte double-precision numbers (IEEE 754 format). A byte is eight bits.

For example, a WKB value that corresponds to [POINT\(1 1\)](#) consists of this sequence of 21 bytes



(each represented here by two hex digits):

```
010100000000000000000000F03F000000000000F03F
```

The sequence may be broken down into these components:

```
Byte order : 01
WKB type   : 01000000
X          : 00000000000000F03F
Y          : 00000000000000F03F
```

Component representation is as follows:

- The byte order may be either 0 or 1 to indicate little-endian or big-endian storage. The little-endian and big-endian byte orders are also known as Network Data Representation (NDR) and External Data Representation (XDR), respectively.
- The WKB type is a code that indicates the geometry type. Values from 1 through 7 indicate [Point](#), [LineString](#), [Polygon](#), [MultiPoint](#), [MultiLineString](#), [MultiPolygon](#), and [GeometryCollection](#).
- A [Point](#) value has X and Y coordinates, each represented as a double-precision value.

WKB values for more complex geometry values are represented by more complex data structures, as detailed in the OpenGIS specification.

## 16.4. Creating a Spatially Enabled MySQL Database

This section describes the data types you can use for representing spatial data in MySQL, and the functions available for creating and retrieving spatial values.

### 16.4.1. MySQL Spatial Data Types

MySQL has data types that correspond to OpenGIS classes. Some of these types hold single geometry values:

- [GEOMETRY](#)
- [POINT](#)
- [LINESTRING](#)
- [POLYGON](#)

[GEOMETRY](#) can store geometry values of any type. The other single-value types ([POINT](#), [LINESTRING](#), and [POLYGON](#)) restrict their values to a particular geometry type.

The other data types hold collections of values:

- [MULTIPOINT](#)
- [MULTILINESTRING](#)
- [MULTIPOLYGON](#)

- `GEOMETRYCOLLECTION`

`GEOMETRYCOLLECTION` can store a collection of objects of any type. The other collection types (`MULTIPOINT`, `MULTILINESTRING`, `MULTIPOLYGON`, and `GEOMETRYCOLLECTION`) restrict collection members to those having a particular geometry type.

## 16.4.2. Creating Spatial Values

This section describes how to create spatial values using Well-Known Text and Well-Known Binary functions that are defined in the OpenGIS standard, and using MySQL-specific functions.

### 16.4.2.1. Creating Geometry Values Using WKT Functions

MySQL provides a number of functions that take as input parameters a Well-Known Text representation and, optionally, a spatial reference system identifier (SRID). They return the corresponding geometry.

`GeomFromText()` accepts a WKT of any geometry type as its first argument. An implementation also provides type-specific construction functions for construction of geometry values of each geometry type.

- `GeomCollFromText(wkt[,srid]), GeometryCollectionFromText(wkt[,srid])`

Constructs a `GEOMETRYCOLLECTION` value using its WKT representation and SRID.

- `GeomFromText(wkt[,srid]), GeometryFromText(wkt[,srid])`

Constructs a geometry value of any type using its WKT representation and SRID.

- `LineFromText(wkt[,srid]), LineStringFromText(wkt[,srid])`

Constructs a `LINestring` value using its WKT representation and SRID.

- `MLineFromText(wkt[,srid]), MultiLineStringFromText(wkt[,srid])`

Constructs a `MULTILINESTRING` value using its WKT representation and SRID.

- `MPointFromText(wkt[,srid]), MultiPointFromText(wkt[,srid])`

Constructs a `MULTIPOINT` value using its WKT representation and SRID.

- `MPolyFromText(wkt[,srid]), MultiPolygonFromText(wkt[,srid])`

Constructs a `MULTIPOLYGON` value using its WKT representation and SRID.

- `PointFromText(wkt[,srid])`

Constructs a `POINT` value using its WKT representation and SRID.

- `PolyFromText(wkt[,srid]), PolygonFromText(wkt[,srid])`

Constructs a `POLYGON` value using its WKT representation and SRID.

The OpenGIS specification also defines the following optional functions, which MySQL does not implement. These functions construct `Polygon` or `MultiPolygon` values based on the WKT representation of a collection of rings or closed `LineString` values. These values may intersect.

- `BdMPolyFromText(wkt, srid)`

Constructs a `MultiPolygon` value from a `MultiLineString` value in WKT format containing an arbitrary collection of closed `LineString` values.

- `BdPolyFromText(wkt, srid)`

Constructs a `Polygon` value from a `MultiLineString` value in WKT format containing an arbitrary collection of closed `LineString` values.

### 16.4.2.2. Creating Geometry Values Using WKB Functions

MySQL provides a number of functions that take as input parameters a `BLOB` containing a Well-Known Binary representation and, optionally, a spatial reference system identifier (SRID). They return the corresponding geometry.

`GeomFromWKB()` accepts a WKB of any geometry type as its first argument. An implementation also provides type-specific construction functions for construction of geometry values of each geometry type.

- `GeomCollFromWKB(wkb[, srid]), GeometryCollectionFromWKB(wkb[, srid])`

Constructs a `GEOMETRYCOLLECTION` value using its WKB representation and SRID.

- `GeomFromWKB(wkb[, srid]), GeometryFromWKB(wkb[, srid])`

Constructs a geometry value of any type using its WKB representation and SRID.

- `LineFromWKB(wkb[, srid]), LineStringFromWKB(wkb[, srid])`

Constructs a `LINESTRING` value using its WKB representation and SRID.

- `MLineFromWKB(wkb[, srid]), MultiLineStringFromWKB(wkb[, srid])`

Constructs a `MULTILINESTRING` value using its WKB representation and SRID.

- `MPointFromWKB(wkb[, srid]), MultiPointFromWKB(wkb[, srid])`

Constructs a `MULTIPOINT` value using its WKB representation and SRID.

- `MPolyFromWKB(wkb[, srid]), MultiPolygonFromWKB(wkb[, srid])`

Constructs a `MULTIPOLYGON` value using its WKB representation and SRID.

- `PointFromWKB(wkb[, srid])`

Constructs a `POINT` value using its WKB representation and SRID.

- `PolyFromWKB(wkb[, srid]), PolygonFromWKB(wkb[, srid])`

Constructs a `POLYGON` value using its WKB representation and SRID.

The OpenGIS specification also describes optional functions for constructing `Polygon` or `MultiPolygon` values based on the WKB representation of a collection of rings or closed `LineString` values. These values may intersect. MySQL does not implement these functions:

- `BdMPolyFromWKB(wkb, srid)`

Constructs a `MultiPolygon` value from a `MultiLineString` value in WKB format containing an arbitrary collection of closed `LineString` values.

- `BdPolyFromWKB(wkb,srid)`

Constructs a `Polygon` value from a `MultiLineString` value in WKB format containing an arbitrary collection of closed `LineString` values.

### 16.4.2.3. Creating Geometry Values Using MySQL-Specific Functions

MySQL provides a set of useful non-standard functions for creating geometry WKB representations. The functions described in this section are MySQL extensions to the OpenGIS specification. The results of these functions are `BLOB` values containing WKB representations of geometry values with no SRID. The results of these functions can be substituted as the first argument for any function in the `GeomFromWKB()` function family.

- `GeometryCollection(g1,g2,...)`

Constructs a WKB `GeometryCollection`. If any argument is not a well-formed WKB representation of a geometry, the return value is `NULL`.

- `LineString(pt1,pt2,...)`

Constructs a WKB `LineString` value from a number of WKB `Point` arguments. If any argument is not a WKB `Point`, the return value is `NULL`. If the number of `Point` arguments is less than two, the return value is `NULL`.

- `MultiLineString(ls1,ls2,...)`

Constructs a WKB `MultiLineString` value using WKB `LineString` arguments. If any argument is not a WKB `LineString`, the return value is `NULL`.

- `MultiPoint(pt1,pt2,...)`

Constructs a WKB `MultiPoint` value using WKB `Point` arguments. If any argument is not a WKB `Point`, the return value is `NULL`.

- `MultiPolygon(poly1,poly2,...)`

Constructs a WKB `MultiPolygon` value from a set of WKB `Polygon` arguments. If any argument is not a WKB `Polygon`, the return value is `NULL`.

- `Point(x,y)`

Constructs a WKB `Point` using its coordinates.

- `Polygon(ls1,ls2,...)`

Constructs a WKB `Polygon` value from a number of WKB `LineString` arguments. If any argument does not represent the WKB of a `LinearRing` (that is, not a closed and simple `LineString`) the return value is `NULL`.

### 16.4.3. Creating Spatial Columns

MySQL provides a standard way of creating spatial columns for geometry types, for example, with

`CREATE TABLE` or `ALTER TABLE`. Currently, spatial columns are supported only for `MyISAM` tables.

- Use the `CREATE TABLE` statement to create a table with a spatial column:

```
CREATE TABLE geom (g GEOMETRY);
```

- Use the `ALTER TABLE` statement to add or drop a spatial column to or from an existing table:

```
ALTER TABLE geom ADD pt POINT;
ALTER TABLE geom DROP pt;
```

## 16.4.4. Populating Spatial Columns

After you have created spatial columns, you can populate them with spatial data.

Values should be stored in internal geometry format, but you can convert them to that format from either Well-Known Text (WKT) or Well-Known Binary (WKB) format. The following examples demonstrate how to insert geometry values into a table by converting WKT values into internal geometry format:

- Perform the conversion directly in the `INSERT` statement:

```
INSERT INTO geom VALUES (GeomFromText('POINT(1 1)'));
SET @g = 'POINT(1 1)';
INSERT INTO geom VALUES (GeomFromText(@g));
```

- Perform the conversion prior to the `INSERT`:

```
SET @g = GeomFromText('POINT(1 1)');
INSERT INTO geom VALUES (@g);
```

The following examples insert more complex geometries into the table:

```
SET @g = 'LINESTRING(0 0,1 1,2 2)';
INSERT INTO geom VALUES (GeomFromText(@g));

SET @g = 'POLYGON((0 0,10 0,10 10,0 10,0 0),(5 5,7 5,7 7,5 7, 5 5))';
INSERT INTO geom VALUES (GeomFromText(@g));

SET @g =
'GEOMETRYCOLLECTION(POINT(1 1),LINESTRING(0 0,1 1,2 2,3 3,4 4))';
INSERT INTO geom VALUES (GeomFromText(@g));
```

The preceding examples all use `GeomFromText()` to create geometry values. You can also use type-specific functions:

```
SET @g = 'POINT(1 1)';
INSERT INTO geom VALUES (PointFromText(@g));

SET @g = 'LINESTRING(0 0,1 1,2 2)';
INSERT INTO geom VALUES (LineStringFromText(@g));

SET @g = 'POLYGON((0 0,10 0,10 10,0 10,0 0),(5 5,7 5,7 7,5 7, 5 5))';
INSERT INTO geom VALUES (PolygonFromText(@g));

SET @g =
'GEOMETRYCOLLECTION(POINT(1 1),LINESTRING(0 0,1 1,2 2,3 3,4 4))';
INSERT INTO geom VALUES (GeomCollFromText(@g));
```

Note that if a client application program wants to use WKB representations of geometry values, it is responsible for sending correctly formed WKB in queries to the server. However, there are several ways of satisfying this requirement. For example:

- Inserting a `POINT(1 1)` value with hex literal syntax:

```
mysql> INSERT INTO geom VALUES
-> (GeomFromWKB(0x010100000000000000000000F03F000000000000F03F));
```

- An ODBC application can send a WKB representation, binding it to a placeholder using an argument of `BLOB` type:

```
INSERT INTO geom VALUES (GeomFromWKB(?))
```

Other programming interfaces may support a similar placeholder mechanism.

- In a C program, you can escape a binary value using `mysql_real_escape_string()` and include the result in a query string that is sent to the server. See [Section 17.2.3.51](#), “`mysql_real_escape_string()`”.

## 16.4.5. Fetching Spatial Data

Geometry values stored in a table can be fetched in internal format. You can also convert them into WKT or WKB format.

- Fetching spatial data in internal format:

Fetching geometry values using internal format can be useful in table-to-table transfers:

```
CREATE TABLE geom2 (g GEOMETRY) SELECT g FROM geom;
```

- Fetching spatial data in WKT format:

The `AsText()` function converts a geometry from internal format into a WKT string.

```
SELECT AsText(g) FROM geom;
```

- Fetching spatial data in WKB format:

The `AsBinary()` function converts a geometry from internal format into a `BLOB` containing the WKB value.

```
SELECT AsBinary(g) FROM geom;
```

## 16.5. Analyzing Spatial Information

After populating spatial columns with values, you are ready to query and analyze them. MySQL provides a set of functions to perform various operations on spatial data. These functions can be grouped into four major categories according to the type of operation they perform:

- Functions that convert geometries between various formats
-

- Functions that provide access to qualitative or quantitative properties of a geometry
- Functions that describe relations between two geometries
- Functions that create new geometries from existing ones

Spatial analysis functions can be used in many contexts, such as:

- Any interactive SQL program, such as `mysql` or MySQL Query Browser
- Application programs written in any language that supports a MySQL client API

## 16.5.1. Geometry Format Conversion Functions

MySQL supports the following functions for converting geometry values between internal format and either WKT or WKB format:

- `AsBinary(g)`

Converts a value in internal geometry format to its WKB representation and returns the binary result.

```
SELECT AsBinary(g) FROM geom;
```

- `AsText(g)`

Converts a value in internal geometry format to its WKT representation and returns the string result.

```
mysql> SET @g = 'LineString(1 1,2 2,3 3)';
mysql> SELECT AsText(GeomFromText(@g));
+-----+
| AsText(GeomFromText(@g)) |
+-----+
| LINESTRING(1 1,2 2,3 3) |
+-----+
```

- `GeomFromText(wkt [, srid])`

Converts a string value from its WKT representation into internal geometry format and returns the result. A number of type-specific functions are also supported, such as `PointFromText()` and `LineFromText()`. See [Section 16.4.2.1, “Creating Geometry Values Using WKT Functions”](#).

- `GeomFromWKB(wkb [, srid])`

Converts a binary value from its WKB representation into internal geometry format and returns the result. A number of type-specific functions are also supported, such as `PointFromWKB()` and `LineFromWKB()`. See [Section 16.4.2.2, “Creating Geometry Values Using WKB Functions”](#).

## 16.5.2. Geometry Functions

Each function that belongs to this group takes a geometry value as its argument and returns some quantitative or qualitative property of the geometry. Some functions restrict their argument type. Such functions return `NULL` if the argument is of an incorrect geometry type. For example, `Area()` returns `NULL` if the object type is neither `Polygon` nor `MultiPolygon`.

### 16.5.2.1. General Geometry Functions

The functions listed in this section do not restrict their argument and accept a geometry value of any type.

- `Dimension(g)`

Returns the inherent dimension of the geometry value *g*. The result can be -1, 0, 1, or 2. The meaning of these values is given in [Section 16.2.2, “Class Geometry”](#).

```
mysql> SELECT Dimension(GeomFromText('LineString(1 1,2 2)'));
+-----+
| Dimension(GeomFromText('LineString(1 1,2 2)')) |
+-----+
| 1 |
+-----+
```

- `Envelope(g)`

Returns the Minimum Bounding Rectangle (MBR) for the geometry value *g*. The result is returned as a `Polygon` value.

The polygon is defined by the corner points of the bounding box:

```
POLYGON((MINX MINY, MAXX MINY, MAXX MAXY, MINX MAXY, MINX MINY))
```

```
mysql> SELECT AsText(Envelope(GeomFromText('LineString(1 1,2 2)')));
+-----+
| AsText(Envelope(GeomFromText('LineString(1 1,2 2)'))) |
+-----+
| POLYGON((1 1,2 1,2 2,1 2,1 1)) |
+-----+
```

- `GeometryType(g)`

Returns as a string the name of the geometry type of which the geometry instance *g* is a member. The name corresponds to one of the instantiable `Geometry` subclasses.

```
mysql> SELECT GeometryType(GeomFromText('POINT(1 1)'));
+-----+
| GeometryType(GeomFromText('POINT(1 1)')) |
+-----+
| POINT |
+-----+
```

- `SRID(g)`

Returns an integer indicating the Spatial Reference System ID for the geometry value *g*.

In MySQL, the SRID value is just an integer associated with the geometry value. All calculations are done assuming Euclidean (planar) geometry.

```
mysql> SELECT SRID(GeomFromText('LineString(1 1,2 2)',101));
+-----+
| SRID(GeomFromText('LineString(1 1,2 2)',101)) |
+-----+
| 101 |
+-----+
```

The OpenGIS specification also defines the following functions, which MySQL does not implement:



- `Boundary(g)`

Returns a geometry that is the closure of the combinatorial boundary of the geometry value *g*.

- `IsEmpty(g)`

Returns 1 if the geometry value *g* is the empty geometry, 0 if it is not empty, and -1 if the argument is `NULL`. If the geometry is empty, it represents the empty point set.

- `IsSimple(g)`

Currently, this function is a placeholder and should not be used. If implemented, its behavior will be as described in the next paragraph.

Returns 1 if the geometry value *g* has no anomalous geometric points, such as self-intersection or self-tangency. `IsSimple()` returns 0 if the argument is not simple, and -1 if it is `NULL`.

The description of each instantiable geometric class given earlier in the chapter includes the specific conditions that cause an instance of that class to be classified as not simple. (See [Section 16.2.1](#), “The Geometry Class Hierarchy”.)

## 16.5.2.2. Point Functions

A `Point` consists of X and Y coordinates, which may be obtained using the following functions:

- `X(p)`

Returns the X-coordinate value for the point *p* as a double-precision number.

```
mysql> SET @pt = 'Point(56.7 53.34)';
mysql> SELECT X(GeomFromText(@pt));
+-----+
| X(GeomFromText(@pt)) |
+-----+
|                    56.7 |
+-----+
```

- `Y(p)`

Returns the Y-coordinate value for the point *p* as a double-precision number.

```
mysql> SET @pt = 'Point(56.7 53.34)';
mysql> SELECT Y(GeomFromText(@pt));
+-----+
| Y(GeomFromText(@pt)) |
+-----+
|                    53.34 |
+-----+
```

## 16.5.2.3. LineString Functions

A `LineString` consists of `Point` values. You can extract particular points of a `LineString`, count the number of points that it contains, or obtain its length.

- `EndPoint(ls)`

Returns the `Point` that is the endpoint of the `LineString` value *ls*.

```
mysql> SET @ls = 'LineString(1 1,2 2,3 3)';
mysql> SELECT AsText(EndPoint(GeomFromText(@ls)));
+-----+
| AsText(EndPoint(GeomFromText(@ls))) |
+-----+
| POINT(3 3) |
+-----+
```

- `GLength(ls)`

Returns as a double-precision number the length of the `LineString` value `ls` in its associated spatial reference.

```
mysql> SET @ls = 'LineString(1 1,2 2,3 3)';
mysql> SELECT GLength(GeomFromText(@ls));
+-----+
| GLength(GeomFromText(@ls)) |
+-----+
| 2.8284271247462 |
+-----+
```

`GLength()` is a non-standard name. It corresponds to the OpenGIS `Length()` function.

- `NumPoints(ls)`

Returns the number of `Point` objects in the `LineString` value `ls`.

```
mysql> SET @ls = 'LineString(1 1,2 2,3 3)';
mysql> SELECT NumPoints(GeomFromText(@ls));
+-----+
| NumPoints(GeomFromText(@ls)) |
+-----+
| 3 |
+-----+
```

- `PointN(ls,N)`

Returns the `N`-th `Point` in the `LineString` value `ls`. Points are numbered beginning with 1.

```
mysql> SET @ls = 'LineString(1 1,2 2,3 3)';
mysql> SELECT AsText(PointN(GeomFromText(@ls),2));
+-----+
| AsText(PointN(GeomFromText(@ls),2)) |
+-----+
| POINT(2 2) |
+-----+
```

- `StartPoint(ls)`

Returns the `Point` that is the start point of the `LineString` value `ls`.

```
mysql> SET @ls = 'LineString(1 1,2 2,3 3)';
mysql> SELECT AsText(StartPoint(GeomFromText(@ls)));
+-----+
| AsText(StartPoint(GeomFromText(@ls))) |
+-----+
| POINT(1 1) |
+-----+
```

The OpenGIS specification also defines the following function, which MySQL does not implement:

- `IsRing(ls)`

Returns 1 if the `LineString` value `ls` is closed (that is, its `StartPoint()` and `EndPoint()` values are the same) and is simple (does not pass through the same point more than once). Returns 0 if `ls` is not a ring, and -1 if it is `NULL`.

### 16.5.2.4. MultiLineString Functions

- `GLength(mls)`

Returns as a double-precision number the length of the `MultiLineString` value `mls`. The length of `mls` is equal to the sum of the lengths of its elements.

```
mysql> SET @mls = 'MultiLineString((1 1,2 2,3 3),(4 4,5 5))';
mysql> SELECT GLength(GeomFromText(@mls));
+-----+
| GLength(GeomFromText(@mls)) |
+-----+
| 4.2426406871193 |
+-----+
```

`GLength()` is a non-standard name. It corresponds to the OpenGIS `Length()` function.

- `IsClosed(mls)`

Returns 1 if the `MultiLineString` value `mls` is closed (that is, the `StartPoint()` and `EndPoint()` values are the same for each `LineString` in `mls`). Returns 0 if `mls` is not closed, and -1 if it is `NULL`.

```
mysql> SET @mls = 'MultiLineString((1 1,2 2,3 3),(4 4,5 5))';
mysql> SELECT IsClosed(GeomFromText(@mls));
+-----+
| IsClosed(GeomFromText(@mls)) |
+-----+
| 0 |
+-----+
```

### 16.5.2.5. Polygon Functions

- `Area(poly)`

Returns as a double-precision number the area of the `Polygon` value `poly`, as measured in its spatial reference system.

```
mysql> SET @poly = 'Polygon((0 0,0 3,3 0,0 0),(1 1,1 2,2 1,1 1))';
mysql> SELECT Area(GeomFromText(@poly));
+-----+
| Area(GeomFromText(@poly)) |
+-----+
| 4 |
+-----+
```

- `ExteriorRing(poly)`

Returns the exterior ring of the `Polygon` value `poly` as a `LineString`.

```
mysql> SET @poly =
-> 'Polygon((0 0,0 3,3 3,3 0,0 0),(1 1,1 2,2 2,2 1,1 1))';
mysql> SELECT AsText(ExteriorRing(GeomFromText(@poly)));
+-----+
| AsText(ExteriorRing(GeomFromText(@poly))) |
+-----+
```

```
+-----+
| LINESTRING(0 0,0 3,3 3,3 0,0 0) |
+-----+
```

- `InteriorRingN(poly,N)`

Returns the *N*-th interior ring for the `Polygon` value *poly* as a `LineString`. Rings are numbered beginning with 1.

```
mysql> SET @poly =
-> 'Polygon((0 0,0 3,3 3,3 0,0 0),(1 1,1 2,2 2,2 1,1 1))';
mysql> SELECT AsText(InteriorRingN(GeomFromText(@poly),1));
+-----+
| AsText(InteriorRingN(GeomFromText(@poly),1)) |
+-----+
| LINESTRING(1 1,1 2,2 2,2 1,1 1) |
+-----+
```

- `NumInteriorRings(poly)`

Returns the number of interior rings in the `Polygon` value *poly*.

```
mysql> SET @poly =
-> 'Polygon((0 0,0 3,3 3,3 0,0 0),(1 1,1 2,2 2,2 1,1 1))';
mysql> SELECT NumInteriorRings(GeomFromText(@poly));
+-----+
| NumInteriorRings(GeomFromText(@poly)) |
+-----+
| 1 |
+-----+
```

### 16.5.2.6. MultiPolygon Functions

- `Area(mpoly)`

Returns as a double-precision number the area of the `MultiPolygon` value *mpoly*, as measured in its spatial reference system.

```
mysql> SET @mpoly =
-> 'MultiPolygon(((0 0,0 3,3 3,3 0,0 0),(1 1,1 2,2 2,2 1,1 1)))';
mysql> SELECT Area(GeomFromText(@mpoly));
+-----+
| Area(GeomFromText(@mpoly)) |
+-----+
| 8 |
+-----+
```

The OpenGIS specification also defines the following functions, which MySQL does not implement:

- `Centroid(mpoly)`

Returns the mathematical centroid for the `MultiPolygon` value *mpoly* as a `Point`. The result is not guaranteed to be on the `MultiPolygon`.

- `PointOnSurface(mpoly)`

Returns a `Point` value that is guaranteed to be on the `MultiPolygon` value *mpoly*.

### 16.5.2.7. GeometryCollection Functions

- `GeometryN(gc,N)`

Returns the *N*-th geometry in the `GeometryCollection` value *gc*. Geometries are numbered beginning with 1.

```
mysql> SET @gc = 'GeometryCollection(Point(1 1),LineString(2 2, 3 3))';
mysql> SELECT AsText(GeometryN(GeomFromText(@gc),1));
+-----+
| AsText(GeometryN(GeomFromText(@gc),1)) |
+-----+
| POINT(1 1)                             |
+-----+
```

- `NumGeometries(gc)`

Returns the number of geometries in the `GeometryCollection` value *gc*.

```
mysql> SET @gc = 'GeometryCollection(Point(1 1),LineString(2 2, 3 3))';
mysql> SELECT NumGeometries(GeomFromText(@gc));
+-----+
| NumGeometries(GeomFromText(@gc)) |
+-----+
| 2                                 |
+-----+
```

## 16.5.3. Functions That Create New Geometries from Existing Ones

### 16.5.3.1. Geometry Functions That Produce New Geometries

[Section 16.5.2, “Geometry Functions”](#), discusses several functions that construct new geometries from existing ones. See that section for descriptions of these functions:

- `Envelope(g)`
- `StartPoint(ls)`
- `EndPoint(ls)`
- `PointN(ls,N)`
- `ExteriorRing(poly)`
- `InteriorRingN(poly,N)`
- `GeometryN(gc,N)`

### 16.5.3.2. Spatial Operators

OpenGIS proposes a number of other functions that can produce geometries. They are designed to implement spatial operators.

These functions are not implemented in MySQL. They may appear in future releases.

- `Buffer(g,d)`  
Returns a geometry that represents all points whose distance from the geometry value *g* is less than or equal to a distance of *d*.
- `ConvexHull(g)`  
Returns a geometry that represents the convex hull of the geometry value *g*.
- `Difference(g1,g2)`  
Returns a geometry that represents the point set difference of the geometry value *g1* with *g2*.
- `Intersection(g1,g2)`  
Returns a geometry that represents the point set intersection of the geometry values *g1* with *g2*.
- `SymDifference(g1,g2)`  
Returns a geometry that represents the point set symmetric difference of the geometry value *g1* with *g2*.
- `Union(g1,g2)`  
Returns a geometry that represents the point set union of the geometry values *g1* and *g2*.

## 16.5.4. Functions for Testing Spatial Relations Between Geometric Objects

The functions described in these sections take two geometries as input parameters and return a qualitative or quantitative relation between them.

## 16.5.5. Relations on Geometry Minimal Bounding Rectangles (MBRs)

MySQL provides several functions that test relations between minimal bounding rectangles of two geometries *g1* and *g2*. The return values 1 and 0 indicate true and false, respectively.

- `MBRContains(g1,g2)`  
Returns 1 or 0 to indicate whether the Minimum Bounding Rectangle of *g1* contains the Minimum Bounding Rectangle of *g2*.  

```
mysql> SET @g1 = GeomFromText('Polygon((0 0,0 3,3 3,3 0,0 0))');
mysql> SET @g2 = GeomFromText('Point(1 1)');
mysql> SELECT MBRContains(@g1,@g2), MBRContains(@g2,@g1);
```

MBRContains(@g1,@g2)	MBRContains(@g2,@g1)
1	0
- `MBRDisjoint(g1,g2)`  
Returns 1 or 0 to indicate whether the Minimum Bounding Rectangles of the two geometries *g1* and *g2* are disjoint (do not intersect).

- `MBREqual(g1,g2)`  
Returns 1 or 0 to indicate whether the Minimum Bounding Rectangles of the two geometries *g1* and *g2* are the same.
- `MBRIntersects(g1,g2)`  
Returns 1 or 0 to indicate whether the Minimum Bounding Rectangles of the two geometries *g1* and *g2* intersect.
- `MBROverlaps(g1,g2)`  
Returns 1 or 0 to indicate whether the Minimum Bounding Rectangles of the two geometries *g1* and *g2* overlap.
- `MBRTouches(g1,g2)`  
Returns 1 or 0 to indicate whether the Minimum Bounding Rectangles of the two geometries *g1* and *g2* touch.
- `MBRWithin(g1,g2)`  
Returns 1 or 0 to indicate whether the Minimum Bounding Rectangle of *g1* is within the Minimum Bounding Rectangle of *g2*.

```
mysql> SET @g1 = GeomFromText('Polygon((0 0,0 3,3 3,3 0,0 0))');
mysql> SET @g2 = GeomFromText('Polygon((0 0,0 5,5 5,5 0,0 0))');
mysql> SELECT MBRWithin(@g1,@g2), MBRWithin(@g2,@g1);
```

MBRWithin(@g1,@g2)	MBRWithin(@g2,@g1)
1	0

## 16.5.6. Functions That Test Spatial Relationships Between Geometries

The OpenGIS specification defines the following functions. They test the relationship between two geometry values *g1* and *g2*.

Currently, MySQL does not implement these functions according to the specification. Those that are implemented return the same result as the corresponding MBR-based functions. This includes functions in the following list other than `Distance()` and `Related()`.

These functions may be implemented in future releases with full support for spatial analysis, not just MBR-based support.

The return values 1 and 0 indicate true and false, respectively.

- `Contains(g1,g2)`  
Returns 1 or 0 to indicate whether *g1* completely contains *g2*.
- `Crosses(g1,g2)`  
Returns 1 if *g1* spatially crosses *g2*. Returns `NULL` if *g1* is a `Polygon` or a `MultiPolygon`, or if *g2* is a `Point` or a `MultiPoint`. Otherwise, returns 0.

The term *spatially crosses* denotes a spatial relation between two given geometries that has the following properties:

- The two geometries intersect
- Their intersection results in a geometry that has a dimension that is one less than the maximum dimension of the two given geometries
- Their intersection is not equal to either of the two given geometries
- `Disjoint(g1,g2)`  
Returns 1 or 0 to indicate whether *g1* is spatially disjoint from (does not intersect) *g2*.
- `Distance(g1,g2)`  
Returns as a double-precision number the shortest distance between any two points in the two geometries.
- `Equals(g1,g2)`  
Returns 1 or 0 to indicate whether *g1* is spatially equal to *g2*.
- `Intersects(g1,g2)`  
Returns 1 or 0 to indicate whether *g1* spatially intersects *g2*.
- `Overlaps(g1,g2)`  
Returns 1 or 0 to indicate whether *g1* spatially overlaps *g2*. The term *spatially overlaps* is used if two geometries intersect and their intersection results in a geometry of the same dimension but not equal to either of the given geometries.
- `Related(g1,g2,pattern_matrix)`  
Returns 1 or 0 to indicate whether the spatial relationship specified by *pattern\_matrix* exists between *g1* and *g2*. Returns -1 if the arguments are `NULL`. The pattern matrix is a string. Its specification will be noted here if this function is implemented.
- `Touches(g1,g2)`  
Returns 1 or 0 to indicate whether *g1* spatially touches *g2*. Two geometries *spatially touch* if the interiors of the geometries do not intersect, but the boundary of one of the geometries intersects either the boundary or the interior of the other.
- `Within(g1,g2)`  
Returns 1 or 0 to indicate whether *g1* is spatially within *g2*.

## 16.6. Optimizing Spatial Analysis

Search operations in non-spatial databases can be optimized using indexes. This is true for spatial databases as well. With the help of a great variety of multi-dimensional indexing methods that have previously been designed, it is possible to optimize spatial searches. The most typical of these are:

- Point queries that search for all objects that contain a given point



- Region queries that search for all objects that overlap a given region

MySQL uses **R-Trees with quadratic splitting** to index spatial columns. A spatial index is built using the MBR of a geometry. For most geometries, the MBR is a minimum rectangle that surrounds the geometries. For a horizontal or a vertical linestring, the MBR is a rectangle degenerated into the linestring. For a point, the MBR is a rectangle degenerated into the point.

## 16.6.1. Creating Spatial Indexes

MySQL can create spatial indexes using syntax similar to that for creating regular indexes, but extended with the `SPATIAL` keyword. Currently, spatial columns that are indexed must be declared `NOT NULL`. The following examples demonstrate how to create spatial indexes:

- With `CREATE TABLE`:

```
CREATE TABLE geom (g GEOMETRY NOT NULL, SPATIAL INDEX(g));
```

- With `ALTER TABLE`:

```
ALTER TABLE geom ADD SPATIAL INDEX(g);
```

- With `CREATE INDEX`:

```
CREATE SPATIAL INDEX sp_index ON geom (g);
```

To drop spatial indexes, use `ALTER TABLE` or `DROP INDEX`:

- With `ALTER TABLE`:

```
ALTER TABLE geom DROP INDEX g;
```

- With `DROP INDEX`:

```
DROP INDEX sp_index ON geom;
```

Example: Suppose that a table `geom` contains more than 32,000 geometries, which are stored in the column `g` of type `GEOMETRY`. The table also has an `AUTO_INCREMENT` column `fid` for storing object ID values.

```
mysql> DESCRIBE geom;
+-----+-----+-----+-----+-----+-----+
| Field | Type      | Null | Key | Default | Extra          |
+-----+-----+-----+-----+-----+-----+
| fid   | int(11)   |      | PRI | NULL    | auto_increment |
| g     | geometry  |      |     |         |                |
+-----+-----+-----+-----+-----+-----+
2 rows in set (0.00 sec)

mysql> SELECT COUNT(*) FROM geom;
+-----+
| count(*) |
+-----+
| 32376    |
+-----+
1 row in set (0.00 sec)
```

To add a spatial index on the column `g`, use this statement:

```
mysql> ALTER TABLE geom ADD SPATIAL INDEX(g);
Query OK, 32376 rows affected (4.05 sec)
Records: 32376 Duplicates: 0 Warnings: 0
```

## 16.6.2. Using a Spatial Index

The optimizer investigates whether available spatial indexes can be involved in the search for queries that use a function such as `MBRContains()` or `MBRWithin()` in the `WHERE` clause. The following query finds all objects that are in the given rectangle:

```
mysql> SET @poly =
-> 'Polygon((30000 15000,31000 15000,31000 16000,30000 16000,30000 15000))';
mysql> SELECT fid,AsText(g) FROM geom WHERE
-> MBRContains(GeomFromText(@poly),g);
```

fid	AsText(g)
21	LINESTRING(30350.4 15828.8,30350.6 15845,30333.8 15845,30 ...
22	LINESTRING(30350.6 15871.4,30350.6 15887.8,30334 15887.8, ...
23	LINESTRING(30350.6 15914.2,30350.6 15930.4,30334 15930.4, ...
24	LINESTRING(30290.2 15823,30290.2 15839.4,30273.4 15839.4, ...
25	LINESTRING(30291.4 15866.2,30291.6 15882.4,30274.8 15882. ...
26	LINESTRING(30291.6 15918.2,30291.6 15934.4,30275 15934.4, ...
249	LINESTRING(30337.8 15938.6,30337.8 15946.8,30320.4 15946. ...
1	LINESTRING(30250.4 15129.2,30248.8 15138.4,30238.2 15136. ...
2	LINESTRING(30220.2 15122.8,30217.2 15137.8,30207.6 15136. ...
3	LINESTRING(30179 15114.4,30176.6 15129.4,30167 15128,3016 ...
4	LINESTRING(30155.2 15121.4,30140.4 15118.6,30142 15109,30 ...
5	LINESTRING(30192.4 15085,30177.6 15082.2,30179.2 15072.4, ...
6	LINESTRING(30244 15087,30229 15086.2,30229.4 15076.4,3024 ...
7	LINESTRING(30200.6 15059.4,30185.6 15058.6,30186 15048.8, ...
10	LINESTRING(30179.6 15017.8,30181 15002.8,30190.8 15003.6, ...
11	LINESTRING(30154.2 15000.4,30168.6 15004.8,30166 15014.2, ...
13	LINESTRING(30105 15065.8,30108.4 15050.8,30118 15053,3011 ...
154	LINESTRING(30276.2 15143.8,30261.4 15141,30263 15131.4,30 ...
155	LINESTRING(30269.8 15084,30269.4 15093.4,30258.6 15093,30 ...
157	LINESTRING(30128.2 15011,30113.2 15010.2,30113.6 15000.4, ...

```
20 rows in set (0.00 sec)
```

Use `EXPLAIN` to check the way this query is executed:

```
mysql> SET @poly =
-> 'Polygon((30000 15000,31000 15000,31000 16000,30000 16000,30000 15000))';
mysql> EXPLAIN SELECT fid,AsText(g) FROM geom WHERE
-> MBRContains(GeomFromText(@poly),g)\G
***** 1. row *****
      id: 1
  select_type: SIMPLE
        table: geom
         type: range
possible_keys: g
          key: g
         key_len: 32
          ref: NULL
         rows: 50
      Extra: Using where
1 row in set (0.00 sec)
```

Check what would happen without a spatial index:

```
mysql> SET @poly =
-> 'Polygon((30000 15000,31000 15000,31000 16000,30000 16000,30000 15000))';
mysql> EXPLAIN SELECT fid,AsText(g) FROM g IGNORE INDEX (g) WHERE
-> MBRContains(GeomFromText(@poly),g)\G
***** 1. row *****
      id: 1
  select_type: SIMPLE
        table: geom
         type: ALL
```

```
possible_keys: NULL
  key: NULL
  key_len: NULL
  ref: NULL
  rows: 32376
  Extra: Using where
1 row in set (0.00 sec)
```

Executing the `SELECT` statement without the spatial index yields the same result but causes the execution time to rise from 0.00 seconds to 0.46 seconds:

```
mysql> SET @poly =
-> 'Polygon((30000 15000,31000 15000,31000 16000,30000 16000,30000 15000))';
mysql> SELECT fid,AsText(g) FROM geom IGNORE INDEX (g) WHERE
-> MBRContains(GeomFromText(@poly),g);
```

fid	AsText(g)
1	LINESTRING(30250.4 15129.2,30248.8 15138.4,30238.2 15136. ...
2	LINESTRING(30220.2 15122.8,30217.2 15137.8,30207.6 15136. ...
3	LINESTRING(30179 15114.4,30176.6 15129.4,30167 15128,3016 ...
4	LINESTRING(30155.2 15121.4,30140.4 15118.6,30142 15109,30 ...
5	LINESTRING(30192.4 15085,30177.6 15082.2,30179.2 15072.4, ...
6	LINESTRING(30244 15087,30229 15086.2,30229.4 15076.4,3024 ...
7	LINESTRING(30200.6 15059.4,30185.6 15058.6,30186 15048.8, ...
10	LINESTRING(30179.6 15017.8,30181 15002.8,30190.8 15003.6, ...
11	LINESTRING(30154.2 15000.4,30168.6 15004.8,30166 15014.2, ...
13	LINESTRING(30105 15065.8,30108.4 15050.8,30118 15053,3011 ...
21	LINESTRING(30350.4 15828.8,30350.6 15845,30333.8 15845,30 ...
22	LINESTRING(30350.6 15871.4,30350.6 15887.8,30334 15887.8, ...
23	LINESTRING(30350.6 15914.2,30350.6 15930.4,30334 15930.4, ...
24	LINESTRING(30290.2 15823,30290.2 15839.4,30273.4 15839.4, ...
25	LINESTRING(30291.4 15866.2,30291.6 15882.4,30274.8 15882. ...
26	LINESTRING(30291.6 15918.2,30291.6 15934.4,30275 15934.4, ...
154	LINESTRING(30276.2 15143.8,30261.4 15141,30263 15131.4,30 ...
155	LINESTRING(30269.8 15084,30269.4 15093.4,30258.6 15093,30 ...
157	LINESTRING(30128.2 15011,30113.2 15010.2,30113.6 15000.4, ...
249	LINESTRING(30337.8 15938.6,30337.8 15946.8,30320.4 15946. ...

```
20 rows in set (0.46 sec)
```

In future releases, spatial indexes may also be used for optimizing other functions. See [Section 16.5.4](#), “Functions for Testing Spatial Relations Between Geometric Objects”.

## 16.7. MySQL Conformance and Compatibility

MySQL does not yet implement the following GIS features:

- Additional Metadata Views

OpenGIS specifications propose several additional metadata views. For example, a system view named `GEOMETRY_COLUMNS` contains a description of geometry columns, one row for each geometry column in the database.

- The OpenGIS function `Length()` on `LineString` and `MultiLineString` currently should be called in MySQL as `GLength()`

The problem is that there is an existing SQL function `Length()` that calculates the length of string values, and sometimes it is not possible to distinguish whether the function is called in a textual or spatial context. We need either to solve this somehow, or decide on another function name.

---

# Chapter 17. APIs and Libraries

This chapter describes the APIs available for MySQL, where to get them, and how to use them. The C API is the most extensively covered, because it was developed by the MySQL team, and is the basis for most of the other APIs. This chapter also covers the `libmysqld` library (the embedded server), as well as some programs that are useful for application developers.

## 17.1. libmysqld, the Embedded MySQL Server Library

### 17.1.1. Overview of the Embedded MySQL Server Library

The embedded MySQL server library makes it possible to run a full-featured MySQL server inside a client application. The main benefits are increased speed and more simple management for embedded applications.

The embedded server library is based on the client/server version of MySQL, which is written in C/C++. Consequently, the embedded server also is written in C/C++. There is no embedded server available in other languages.

The API is identical for the embedded MySQL version and the client/server version. To change an old threaded application to use the embedded library, you normally only have to add calls to the following functions:

Function	When to Call
<code>mysql_server_init()</code>	Should be called before any other MySQL function is called, preferably early in the <code>main()</code> function.
<code>mysql_server_end()</code>	Should be called before your program exits.
<code>mysql_thread_init()</code>	Should be called in each thread you create that accesses MySQL.
<code>mysql_thread_end()</code>	Should be called before calling <code>pthread_exit()</code>

Then you must link your code with `libmysqld.a` instead of `libmysqlclient.a`.

The `mysql_server_XXX()` functions are also included in `libmysqlclient.a` to allow you to change between the embedded and the client/server version by just linking your application with the right library. See [Section 17.2.12.1, “mysql\\_server\\_init\(\)”](#).

One difference between the embedded server and the standalone server is that for the embedded server, authentication for connections is disabled by default. To use authentication for the embedded server, specify the `--with-embedded-privilege-control` option when you invoke `configure` to configure your MySQL distribution. This option is available as of MySQL 4.1.3.

### 17.1.2. Compiling Programs with `libmysqld`

To get a `libmysqld` library you should configure MySQL with the `--with-embedded-server` option. See [Section 2.9.2, “Typical configure Options”](#).

When you link your program with `libmysqld`, you must also include the system-specific `pthread` libraries and some libraries that the MySQL server uses. You can get the full list of libraries by executing `mysql_config --libmysqld-libs`.

The correct flags for compiling and linking a threaded program must be used, even if you do not directly call any thread functions in your code.

To compile a C program to include the necessary files to embed the MySQL server library into a compiled version of a program, use the GNU C compiler (`gcc`). The compiler will need to know where to find various files and need instructions on how to compile the program. The following example shows how a program could be compiled from the command line:

```
gcc mysql_test.c -o mysql_test -lz \  
`/usr/local/mysql/bin/mysql_config --include --libmysqld-libs`
```

Immediately following the `gcc` command is the name of the uncompiled C program file. After it, the `-o` option is given to indicate that the file name that follows is the name that the compiler is to give to the output file, the compiled program. The next line of code tells the compiler to obtain the location of the include files and libraries and other settings for the system on which it's compiled. Because of a problem with `mysql_config`, the option `-lz` (for compression) is added here. The `mysql_config` piece is contained in backticks, not single quotes.

### 17.1.3. Restrictions When Using the Embedded MySQL Server

The embedded server has the following limitations:

- No support for `ISAM` tables. (This is done mainly to make the library smaller.)
- No user-defined functions (UDFs).
- No stack trace on core dump.
- No internal RAID support. (This is not normally needed as most current operating systems support big files.)
- You cannot set this up as a master or a slave (no replication).
- Very large result sets may be unusable on low memory systems.
- You cannot connect to an embedded server from an outside process with sockets or TCP/IP. However, you can connect to an intermediate application, which in turn can connect to an embedded server on the behalf of a remote client or outside process.

Some of these limitations can be changed by editing the `mysql_embed.h` include file and recompiling MySQL.

### 17.1.4. Options with the Embedded Server

Any options that may be given with the `mysqld` server daemon, may be used with an embedded server library. Server options may be given in an array as an argument to the `mysql_server_init()`, which initializes the server. They also may be given in an option file like `my.cnf`. To specify an option file for a C program, use the `--defaults-file` option as one of the elements of the second argument of the `mysql_server_init()` function. See [Section 17.2.12.1](#), “`mysql_server_init()`”, for more information on the `mysql_server_init()` function.

Using option files can make it easier to switch between a client/server application and one where MySQL is embedded. Put common options under the `[server]` group. These are read by both MySQL versions. Client/server-specific options should go under the `[mysqld]` section. Put options

specific to the embedded MySQL server library in the [ [embedded](#) ] section. Options specific to applications go under section labeled [ [ApplicationName\\_SERVER](#) ]. See [Section 4.3.2, “Using Option Files”](#).

## 17.1.5. Embedded Server Examples

These two example programs should work without any changes on a Linux or FreeBSD system. For other operating systems, minor changes are needed, mostly with file paths. These examples are designed to give enough details for you to understand the problem, without the clutter that is a necessary part of a real application. The first example is very straightforward. The second example is a little more advanced with some error checking. The first is followed by a command-line entry for compiling the program. The second is followed by a GNUmake file that may be used for compiling instead.

### Example 1

`test1_libmysqld.c`

```
#include <stdio.h>
#include <stdlib.h>
#include <stdarg.h>
#include "mysql.h"

MYSQL *mysql;
MYSQL_RES *results;
MYSQL_ROW record;

static char *server_options[] = { "mysql_test", "--defaults-file=my.cnf", NULL };
int num_elements = sizeof(server_options)/ sizeof(char *);

static char *server_groups[] = { "libmysqld_server", "libmysqld_client", NULL };

int main(void)
{
    mysql_server_init(num_elements, server_options, server_groups);
    mysql = mysql_init(NULL);
    mysql_options(mysql, MYSQL_READ_DEFAULT_GROUP, "libmysqld_client");
    mysql_options(mysql, MYSQL_OPT_USE_EMBEDDED_CONNECTION, NULL);

    mysql_real_connect(mysql, NULL, NULL, NULL, "database1", 0, NULL, 0);

    mysql_query(mysql, "SELECT column1, column2 FROM table1");

    results = mysql_store_result(mysql);

    while((record = mysql_fetch_row(results))) {
        printf("%s - %s \n", record[0], record[1]);
    }

    mysql_free_result(results);
    mysql_close(mysql);
    mysql_server_end();

    return 0;
}
```

Here is the command line for compiling the above program:

```
gcc test1_libmysqld.c -o test1_libmysqld -lz \
`/usr/local/mysql/bin/mysql_config --include --libmysqld-libs`
```

### Example 2

To try out the example, create a `test2_libmysqld` directory at the same level as the MySQL source directory. Save the `test2_libmysqld.c` source and the [GNUmakefile](#) in the directory, and run GNU `make` from inside the `test2_libmysqld` directory.

`test2_libmysqld.c`

```
/*
```

```

* A simple example client, using the embedded MySQL server library
*/

#include <mysql.h>
#include <stdarg.h>
#include <stdio.h>
#include <stdlib.h>

MYSQL *db_connect(const char *dbname);
void db_disconnect(MYSQL *db);
void db_do_query(MYSQL *db, const char *query);

const char *server_groups[] = {
    "test2_libmysqld_SERVER", "embedded", "server", NULL
};

int
main(int argc, char **argv)
{
    MYSQL *one, *two;

    /* mysql_server_init() must be called before any other mysql
     * functions.
     *
     * You can use mysql_server_init(0, NULL, NULL), and it
     * initializes the server using groups = {
     *     "server", "embedded", NULL
     * }.
     *
     * In your $HOME/.my.cnf file, you probably want to put:
     [test2_libmysqld_SERVER]
     language = /path/to/source/of/mysql/sql/share/english
     *
     * You could, of course, modify argc and argv before passing
     * them to this function. Or you could create new ones in any
     * way you like. But all of the arguments in argv (except for
     * argv[0], which is the program name) should be valid options
     * for the MySQL server.
     *
     * If you link this client against the normal mysqlclient
     * library, this function is just a stub that does nothing.
     */
    mysql_server_init(argc, argv, (char **)server_groups);

    one = db_connect("test");
    two = db_connect(NULL);

    db_do_query(one, "SHOW TABLE STATUS");
    db_do_query(two, "SHOW DATABASES");

    mysql_close(two);
    mysql_close(one);

    /* This must be called after all other mysql functions */
    mysql_server_end();

    exit(EXIT_SUCCESS);
}

static void
die(MYSQL *db, char *fmt, ...)
{
    va_list ap;
    va_start(ap, fmt);
    vfprintf(stderr, fmt, ap);
    va_end(ap);
    (void)putc('\n', stderr);
    if (db)
        db_disconnect(db);
    exit(EXIT_FAILURE);
}

MYSQL *
db_connect(const char *dbname)
{
    MYSQL *db = mysql_init(NULL);
    if (!db)
        die(db, "mysql_init failed: no memory");
    /*

```

```

* Notice that the client and server use separate group names.
* This is critical, because the server does not accept the
* client's options, and vice versa.
*/
mysql_options(db, MYSQL_READ_DEFAULT_GROUP, "test2_libmysqld_CLIENT");
if (!mysql_real_connect(db, NULL, NULL, NULL, dbname, 0, NULL, 0))
    die(db, "mysql_real_connect failed: %s", mysql_error(db));

return db;
}

void
db_disconnect(MYSQL *db)
{
    mysql_close(db);
}

void
db_do_query(MYSQL *db, const char *query)
{
    if (mysql_query(db, query) != 0)
        goto err;

    if (mysql_field_count(db) > 0)
    {
        MYSQL_RES *res;
        MYSQL_ROW row, end_row;
        int num_fields;

        if (!(res = mysql_store_result(db)))
            goto err;
        num_fields = mysql_num_fields(res);
        while ((row = mysql_fetch_row(res)))
        {
            (void)fputs(">> ", stdout);
            for (end_row = row + num_fields; row < end_row; ++row)
                (void)printf("%s\t", row ? (char*)*row : "NULL");
            (void)fputc('\n', stdout);
        }
        (void)fputc('\n', stdout);
        mysql_free_result(res);
    }
    else
        (void)printf("Affected rows: %lld\n", mysql_affected_rows(db));

    return;
err:
    die(db, "db_do_query failed: %s [%s]", mysql_error(db), query);
}

```

## GNUMakefile

```

# This assumes the MySQL software is installed in /usr/local/mysql
inc      := /usr/local/mysql/include/mysql
lib      := /usr/local/mysql/lib

# If you have not installed the MySQL software yet, try this instead
#inc     := $(HOME)/mysql-4.0/include
#lib     := $(HOME)/mysql-4.0/libmysqld

CC       := gcc
CPPFLAGS := -I$(inc) -D_THREAD_SAFE -D_REENTRANT
CFLAGS  := -g -W -Wall
LDFLAGS := -static
# You can change -lmysqld to -lmysqlclient to use the
# client/server library
LDLIBS  = -L$(lib) -lmysqld -lz -lm -lcrypt

ifneq (,$(shell grep FreeBSD /COPYRIGHT 2>/dev/null))
# FreeBSD
LDFLAGS += -pthread
else
# Assume Linux
LDLIBS += -lpthread
endif

# This works for simple one-file test programs

```



```
sources := $(wildcard *.c)
objects := $(patsubst %c,%o,$(sources))
targets := $(basename $(sources))

all: $(targets)

clean:
    rm -f $(targets) $(objects) *.core
```

## 17.1.6. Licensing the Embedded Server

We encourage everyone to promote free software by releasing code under the GPL or a compatible license. For those who are not able to do this, another option is to purchase a commercial license for the MySQL code from MySQL AB. For details, please see <http://www.mysql.com/company/legal/licensing/>.

## 17.2. MySQL C API

The C API code is distributed with MySQL. It is included in the `mysqlclient` library and allows C programs to access a database.

Many of the clients in the MySQL source distribution are written in C. If you are looking for examples that demonstrate how to use the C API, take a look at these clients. You can find these in the `clients` directory in the MySQL source distribution.

Most of the other client APIs (all except Connector/J and Connector/NET) use the `mysqlclient` library to communicate with the MySQL server. This means that, for example, you can take advantage of many of the same environment variables that are used by other client programs, because they are referenced from the library. See [Chapter 8, Client and Utility Programs](#), for a list of these variables.

The client has a maximum communication buffer size. The size of the buffer that is allocated initially (16KB) is automatically increased up to the maximum size (the maximum is 16MB). Because buffer sizes are increased only as demand warrants, simply increasing the default maximum limit does not in itself cause more resources to be used. This size check is mostly a check for erroneous statements and communication packets.

The communication buffer must be large enough to contain a single SQL statement (for client-to-server traffic) and one row of returned data (for server-to-client traffic). Each thread's communication buffer is dynamically enlarged to handle any query or row up to the maximum limit. For example, if you have `BLOB` values that contain up to 16MB of data, you must have a communication buffer limit of at least 16MB (in both server and client). The client's default maximum is 16MB, but the default maximum in the server is 1MB. You can increase this by changing the value of the `max_allowed_packet` parameter when the server is started. See [Section 7.5.2, "Tuning Server Parameters"](#).

The MySQL server shrinks each communication buffer to `net_buffer_length` bytes after each query. For clients, the size of the buffer associated with a connection is not decreased until the connection is closed, at which time client memory is reclaimed.

For programming with threads, see [Section 17.2.15, "How to Make a Threaded Client"](#). For creating a standalone application which includes the "server" and "client" in the same program (and does not communicate with an external MySQL server), see [Section 17.1, "libmysqld, the Embedded MySQL Server Library"](#).

### 17.2.1. C API Data types

- `MYSQL`

This structure represents a handle to one database connection. It is used for almost all MySQL functions. You should not try to make a copy of a `MYSQL` structure. There is no guarantee that such a copy will be usable.

- `MYSQL_RES`

This structure represents the result of a query that returns rows (`SELECT`, `SHOW`, `DESCRIBE`, `EXPLAIN`). The information returned from a query is called the *result set* in the remainder of this section.

- `MYSQL_ROW`

This is a type-safe representation of one row of data. It is currently implemented as an array of counted byte strings. (You cannot treat these as null-terminated strings if field values may contain binary data, because such values may contain null bytes internally.) Rows are obtained by calling `mysql_fetch_row()`.

- `MYSQL_FIELD`

This structure contains information about a field, such as the field's name, type, and size. Its members are described in more detail here. You may obtain the `MYSQL_FIELD` structures for each field by calling `mysql_fetch_field()` repeatedly. Field values are not part of this structure; they are contained in a `MYSQL_ROW` structure.

- `MYSQL_FIELD_OFFSET`

This is a type-safe representation of an offset into a MySQL field list. (Used by `mysql_field_seek()`.) Offsets are field numbers within a row, beginning at zero.

- `my_ulonglong`

The type used for the number of rows and for `mysql_affected_rows()`, `mysql_num_rows()`, and `mysql_insert_id()`. This type provides a range of 0 to `1.84e19`.

On some systems, attempting to print a value of type `my_ulonglong` does not work. To print such a value, convert it to `unsigned long` and use a `%lu` print format. Example:

```
printf ("Number of rows: %lu\n", (unsigned long) mysql_num_rows(result));
```

- `my_bool`

A boolean type, for values that are true (non-zero) or false (zero).

The `MYSQL_FIELD` structure contains the members listed here:

- `char * name`

The name of the field, as a null-terminated string. If the field was given an alias with an `AS` clause, the value of `name` is the alias.

- `char * org_name`

The name of the field, as a null-terminated string. Aliases are ignored. This member was added in MySQL 4.1.0.

- `char * table`

The name of the table containing this field, if it isn't a calculated field. For calculated fields, the `table` value is an empty string. If the table was given an alias with an `AS` clause, the value of `table` is the alias.

- `char * org_table`

The name of the table, as a null-terminated string. Aliases are ignored. This member was added in MySQL 4.1.0.

- `char * db`

The name of the database that the field comes from, as a null-terminated string. If the field is a calculated field, `db` is an empty string. This member was added in MySQL 4.1.0.

- `char * catalog`

The catalog name. This value is always `"def"`. This member was added in MySQL 4.1.1.

- `char * def`

The default value of this field, as a null-terminated string. This is set only if you use `mysql_list_fields()`.

- `unsigned long length`

The width of the field, as specified in the table definition.

- `unsigned long max_length`

The maximum width of the field for the result set (the length of the longest field value for the rows actually in the result set). If you use `mysql_store_result()` or `mysql_list_fields()`, this contains the maximum length for the field. If you use `mysql_use_result()`, the value of this variable is zero.

- `unsigned int name_length`

The length of `name`. This member was added in MySQL 4.1.0.

- `unsigned int org_name_length`

The length of `org_name`. This member was added in MySQL 4.1.0.

- `unsigned int table_length`

The length of `table`. This member was added in MySQL 4.1.0.

- `unsigned int org_table_length`

The length of `org_table`. This member was added in MySQL 4.1.0.

- `unsigned int db_length`

The length of `db`. This member was added in MySQL 4.1.0.

- `unsigned int catalog_length`

The length of `catalog`. This member was added in MySQL 4.1.1.

- `unsigned int def_length`

The length of `def`. This member was added in MySQL 4.1.0.

- `unsigned int flags`

Different bit-flags for the field. The `flags` value may have zero or more of the following bits set:

Flag Value	Flag Description
<code>NOT_NULL_FLAG</code>	Field can't be <code>NULL</code>
<code>PRI_KEY_FLAG</code>	Field is part of a primary key
<code>UNIQUE_KEY_FLAG</code>	Field is part of a unique key
<code>MULTIPLE_KEY_FLAG</code>	Field is part of a non-unique key
<code>UNSIGNED_FLAG</code>	Field has the <code>UNSIGNED</code> attribute
<code>ZEROFILL_FLAG</code>	Field has the <code>ZEROFILL</code> attribute
<code>BINARY_FLAG</code>	Field has the <code>BINARY</code> attribute
<code>AUTO_INCREMENT_FLAG</code>	Field has the <code>AUTO_INCREMENT</code> attribute
<code>ENUM_FLAG</code>	Field is an <code>ENUM</code> (deprecated)
<code>SET_FLAG</code>	Field is a <code>SET</code> (deprecated)
<code>BLOB_FLAG</code>	Field is a <code>BLOB</code> or <code>TEXT</code> (deprecated)
<code>TIMESTAMP_FLAG</code>	Field is a <code>TIMESTAMP</code> (deprecated)

Use of the `BLOB_FLAG`, `ENUM_FLAG`, `SET_FLAG`, and `TIMESTAMP_FLAG` flags is deprecated because they indicate the type of a field rather than an attribute of its type. It is preferable to test `field->type` against `MYSQL_TYPE_BLOB`, `MYSQL_TYPE_ENUM`, `MYSQL_TYPE_SET`, or `MYSQL_TYPE_TIMESTAMP` instead.

The following example illustrates a typical use of the `flags` value:

```
if (field->flags & NOT_NULL_FLAG)
    printf("Field can't be null\n");
```

You may use the following convenience macros to determine the boolean status of the `flags` value:

Flag Status	Description
<code>IS_NOT_NULL(flags)</code>	True if this field is defined as <code>NOT NULL</code>
<code>IS_PRI_KEY(flags)</code>	True if this field is a primary key
<code>IS_BLOB(flags)</code>	True if this field is a <code>BLOB</code> or <code>TEXT</code> (deprecated; test <code>field-&gt;type</code> instead)

- `unsigned int decimals`

The number of decimals for numeric fields.

- `unsigned int charsetnr`

The character set number for the field. This member was added in MySQL 4.1.0.

- `enum enum_field_types type`

The type of the field. The `type` value may be one of the `MYSQL_TYPE_` symbols shown in the following table. Before MySQL 4.1, the symbol names begin with `FIELD_TYPE_` rather than `MYSQL_TYPE_`. The older types still are recognized for backward compatibility.

Type Value	Type Description
<code>MYSQL_TYPE_TINY</code>	TINYINT field
<code>MYSQL_TYPE_SHORT</code>	SMALLINT field
<code>MYSQL_TYPE_LONG</code>	INTEGER field
<code>MYSQL_TYPE_INT24</code>	MEDIUMINT field
<code>MYSQL_TYPE_LONGLONG</code>	BIGINT field
<code>MYSQL_TYPE_DECIMAL</code>	DECIMAL or NUMERIC field
<code>MYSQL_TYPE_FLOAT</code>	FLOAT field
<code>MYSQL_TYPE_DOUBLE</code>	DOUBLE or REAL field
<code>MYSQL_TYPE_TIMESTAMP</code>	TIMESTAMP field
<code>MYSQL_TYPE_DATE</code>	DATE field
<code>MYSQL_TYPE_TIME</code>	TIME field
<code>MYSQL_TYPE_DATETIME</code>	DATETIME field
<code>MYSQL_TYPE_YEAR</code>	YEAR field
<code>MYSQL_TYPE_STRING</code>	CHAR or BINARY field
<code>MYSQL_TYPE_VAR_STRING</code>	VARCHAR or VARBINARY field
<code>MYSQL_TYPE_BLOB</code>	BLOB or TEXT field (use <code>max_length</code> to determine the maximum length)
<code>MYSQL_TYPE_SET</code>	SET field
<code>MYSQL_TYPE_ENUM</code>	ENUM field
<code>MYSQL_TYPE_GEOMETRY</code>	Spatial field (MySQL 4.1.0 and up)
<code>MYSQL_TYPE_NULL</code>	NULL-type field
<code>MYSQL_TYPE_CHAR</code>	Deprecated; use <code>MYSQL_TYPE_TINY</code> instead

You can use the `IS_NUM()` macro to test whether a field has a numeric type. Pass the `type` value to `IS_NUM()` and it evaluates to TRUE if the field is numeric:

```
if (IS_NUM(field->type))
    printf("Field is numeric\n");
```

To distinguish between binary and non-binary data for string data types, check whether the `charsetnr` value is 63. If so, the character set is `binary`, which indicates binary rather than non-binary data. This is how to distinguish between `BINARY` and `CHAR`, `VARBINARY` and `VARCHAR`, and `BLOB` and `TEXT`.

## 17.2.2. C API Function Overview

The functions available in the C API are summarized here and described in greater detail in a later section. See [Section 17.2.3, “C API Function Descriptions”](#).

Function	Description
<code>mysql_affected_rows()</code>	Returns the number of rows changed/deleted/inserted by the last <code>UP-</code>

	DATE, DELETE, or INSERT query.
<code>mysql_autocommit()</code>	Toggles autocommit mode on/off.
<code>mysql_change_user()</code>	Changes user and database on an open connection.
<code>mysql_close()</code>	Closes a server connection.
<code>mysql_commit()</code>	Commits the transaction.
<code>mysql_connect()</code>	Connects to a MySQL server. This function is deprecated; use <code>mysql_real_connect()</code> instead.
<code>mysql_create_db()</code>	Creates a database. This function is deprecated; use the SQL statement <code>CREATE DATABASE</code> instead.
<code>mysql_data_seek()</code>	Seeks to an arbitrary row number in a query result set.
<code>mysql_debug()</code>	Does a <code>DEBUG_PUSH</code> with the given string.
<code>mysql_drop_db()</code>	Drops a database. This function is deprecated; use the SQL statement <code>DROP DATABASE</code> instead.
<code>mysql_dump_debug_info()</code>	Makes the server write debug information to the log.
<code>mysql_eof()</code>	Determines whether the last row of a result set has been read. This function is deprecated; <code>mysql_errno()</code> or <code>mysql_error()</code> may be used instead.
<code>mysql_errno()</code>	Returns the error number for the most recently invoked MySQL function.
<code>mysql_error()</code>	Returns the error message for the most recently invoked MySQL function.
<code>mysql_escape_string()</code>	Escapes special characters in a string for use in an SQL statement.
<code>mysql_fetch_field()</code>	Returns the type of the next table field.
<code>mysql_fetch_field_direct()</code>	Returns the type of a table field, given a field number.
<code>mysql_fetch_fields()</code>	Returns an array of all field structures.
<code>mysql_fetch_lengths()</code>	Returns the lengths of all columns in the current row.
<code>mysql_fetch_row()</code>	Fetches the next row from the result set.
<code>mysql_field_seek()</code>	Puts the column cursor on a specified column.
<code>mysql_field_count()</code>	Returns the number of result columns for the most recent statement.
<code>mysql_field_tell()</code>	Returns the position of the field cursor used for the last <code>mysql_fetch_field()</code> .
<code>mysql_free_result()</code>	Frees memory used by a result set.
<code>mysql_get_client_info()</code>	Returns client version information as a string.
<code>mysql_get_client_version()</code>	Returns client version information as an integer.
<code>mysql_get_host_info()</code>	Returns a string describing the connection.
<code>mysql_get_server_version()</code>	Returns version number of server as an integer (new in 4.1).
<code>mysql_get_proto_info()</code>	Returns the protocol version used by the connection.
<code>mysql_get_server_info()</code>	Returns the server version number.
<code>mysql_info()</code>	Returns information about the most recently executed query.
<code>mysql_init()</code>	Gets or initializes a <code>MYSQL</code> structure.
<code>mysql_insert_id()</code>	Returns the ID generated for an <code>AUTO_INCREMENT</code> column by the previous query.
<code>mysql_kill()</code>	Kills a given thread.
<code>mysql_library_end()</code>	Finalize MySQL C API library.

<code>mysql_library_init()</code>	Initialize MySQL C API library.
<code>mysql_list_dbs()</code>	Returns database names matching a simple regular expression.
<code>mysql_list_fields()</code>	Returns field names matching a simple regular expression.
<code>mysql_list_processes()</code>	Returns a list of the current server threads.
<code>mysql_list_tables()</code>	Returns table names matching a simple regular expression.
<code>mysql_more_results()</code>	Checks whether any more results exist.
<code>mysql_next_result()</code>	Returns/initiates the next result in multiple-statement executions.
<code>mysql_num_fields()</code>	Returns the number of columns in a result set.
<code>mysql_num_rows()</code>	Returns the number of rows in a result set.
<code>mysql_options()</code>	Sets connect options for <code>mysql_connect()</code> .
<code>mysql_ping()</code>	Checks whether the connection to the server is working, reconnecting as necessary.
<code>mysql_query()</code>	Executes an SQL query specified as a null-terminated string.
<code>mysql_real_connect()</code>	Connects to a MySQL server.
<code>mysql_real_escape_string()</code>	Escapes special characters in a string for use in an SQL statement, taking into account the current character set of the connection.
<code>mysql_real_query()</code>	Executes an SQL query specified as a counted string.
<code>mysql_refresh()</code>	Flush or reset tables and caches.
<code>mysql_reload()</code>	Tells the server to reload the grant tables.
<code>mysql_rollback()</code>	Rolls back the transaction.
<code>mysql_row_seek()</code>	Seeks to a row offset in a result set, using value returned from <code>mysql_row_tell()</code> .
<code>mysql_row_tell()</code>	Returns the row cursor position.
<code>mysql_select_db()</code>	Selects a database.
<code>mysql_server_end()</code>	Finalize embedded server library.
<code>mysql_server_init()</code>	Initialize embedded server library.
<code>mysql_set_local_infile_default()</code>	Set the <code>LOAD DATA LOCAL INFILE</code> handler callbacks to their default values.
<code>mysql_set_local_infile_handler()</code>	Install application-specific <code>LOAD DATA LOCAL INFILE</code> handler callbacks.
<code>mysql_set_server_option()</code>	Sets an option for the connection (like <code>multi-statements</code> ).
<code>mysql_sqlstate()</code>	Returns the SQLSTATE error code for the last error.
<code>mysql_shutdown()</code>	Shuts down the database server.
<code>mysql_stat()</code>	Returns the server status as a string.
<code>mysql_store_result()</code>	Retrieves a complete result set to the client.
<code>mysql_thread_id()</code>	Returns the current thread ID.
<code>mysql_thread_safe()</code>	Returns 1 if the clients are compiled as thread-safe.
<code>mysql_use_result()</code>	Initiates a row-by-row result set retrieval.
<code>mysql_warning_count()</code>	Returns the warning count for the previous SQL statement.

Application programs should use this general outline for interacting with MySQL:

1. Initialize the MySQL library by calling `mysql_library_init()`. The library can be either the

`mysqlclient` C client library or the `mysqld` embedded server library, depending on whether the application was linked with the `-libmysqlclient` or `-libmysqld` flag.

2. Initialize a connection handler by calling `mysql_init()` and connect to the server by calling `mysql_real_connect()`.
3. Issue SQL statements and process their results. (The following discussion provides more information about how to do this.)
4. Close the connection to the MySQL server by calling `mysql_close()`.
5. End use of the MySQL library by calling `mysql_library_end()`.

The purpose of calling `mysql_library_init()` and `mysql_library_end()` is to provide proper initialization and finalization of the MySQL library. For applications that are linked with the client library, they provide improved memory management. If you don't call `mysql_library_end()`, a block of memory remains allocated. (This does not increase the amount of memory used by the application, but some memory leak detectors will complain about it.) For applications that are linked with the embedded server, these calls start and stop the server.

`mysql_library_init()` and `mysql_library_end()` are available as of MySQL 4.1.10. These are actually `#define` symbols that make them equivalent to `mysql_server_init()` and `mysql_server_end()`, but the names more clearly indicate that they should be called when beginning and ending use of a MySQL library no matter whether the application uses the `mysqlclient` or `mysqld` library. For older versions of MySQL, you can call `mysql_server_init()` and `mysql_server_end()` instead.

In a non-multi-threaded environment, the call to `mysql_library_init()` may be omitted, because `mysql_init()` will invoke it automatically as necessary. However, a race condition is possible if `mysql_library_init()` is invoked by `mysql_init()` in a multi-threaded environment: `mysql_library_init()` is not thread-safe, so it should be called prior to any other client library call.

To connect to the server, call `mysql_init()` to initialize a connection handler, then call `mysql_real_connect()` with that handler (along with other information such as the hostname, username, and password). Upon connection, `mysql_real_connect()` sets the `reconnect` flag (part of the `MYSQL` structure) to a value of `1`. A value of `1` for this flag indicates that if a statement cannot be performed because of a lost connection, to try reconnecting to the server before giving up. When you are done with the connection, call `mysql_close()` to terminate it.

While a connection is active, the client may send SQL statements to the server using `mysql_query()` or `mysql_real_query()`. The difference between the two is that `mysql_query()` expects the query to be specified as a null-terminated string whereas `mysql_real_query()` expects a counted string. If the string contains binary data (which may include null bytes), you must use `mysql_real_query()`.

For each non-`SELECT` query (for example, `INSERT`, `UPDATE`, `DELETE`), you can find out how many rows were changed (affected) by calling `mysql_affected_rows()`.

For `SELECT` queries, you retrieve the selected rows as a result set. (Note that some statements are `SELECT`-like in that they return rows. These include `SHOW`, `DESCRIBE`, and `EXPLAIN`. They should be treated the same way as `SELECT` statements.)

There are two ways for a client to process result sets. One way is to retrieve the entire result set all at once by calling `mysql_store_result()`. This function acquires from the server all the rows returned by the query and stores them in the client. The second way is for the client to initiate a row-by-row result set retrieval by calling `mysql_use_result()`. This function initializes the retrieval, but does not actually get any rows from the server.



In both cases, you access rows by calling `mysql_fetch_row()`. With `mysql_store_result()`, `mysql_fetch_row()` accesses rows that have previously been fetched from the server. With `mysql_use_result()`, `mysql_fetch_row()` actually retrieves the row from the server. Information about the size of the data in each row is available by calling `mysql_fetch_lengths()`.

After you are done with a result set, call `mysql_free_result()` to free the memory used for it.

The two retrieval mechanisms are complementary. Client programs should choose the approach that is most appropriate for their requirements. In practice, clients tend to use `mysql_store_result()` more commonly.

An advantage of `mysql_store_result()` is that because the rows have all been fetched to the client, you not only can access rows sequentially, you can move back and forth in the result set using `mysql_data_seek()` or `mysql_row_seek()` to change the current row position within the result set. You can also find out how many rows there are by calling `mysql_num_rows()`. On the other hand, the memory requirements for `mysql_store_result()` may be very high for large result sets and you are more likely to encounter out-of-memory conditions.

An advantage of `mysql_use_result()` is that the client requires less memory for the result set because it maintains only one row at a time (and because there is less allocation overhead, `mysql_use_result()` can be faster). Disadvantages are that you must process each row quickly to avoid tying up the server, you don't have random access to rows within the result set (you can only access rows sequentially), and you don't know how many rows are in the result set until you have retrieved them all. Furthermore, you **must** retrieve all the rows even if you determine in mid-retrieval that you've found the information you were looking for.

The API makes it possible for clients to respond appropriately to statements (retrieving rows only as necessary) without knowing whether the statement is a `SELECT`. You can do this by calling `mysql_store_result()` after each `mysql_query()` (or `mysql_real_query()`). If the result set call succeeds, the statement was a `SELECT` and you can read the rows. If the result set call fails, call `mysql_field_count()` to determine whether a result was actually to be expected. If `mysql_field_count()` returns zero, the statement returned no data (indicating that it was an `INSERT`, `UPDATE`, `DELETE`, and so forth), and was not expected to return rows. If `mysql_field_count()` is non-zero, the statement should have returned rows, but didn't. This indicates that the statement was a `SELECT` that failed. See the description for `mysql_field_count()` for an example of how this can be done.

Both `mysql_store_result()` and `mysql_use_result()` allow you to obtain information about the fields that make up the result set (the number of fields, their names and types, and so forth). You can access field information sequentially within the row by calling `mysql_fetch_field()` repeatedly, or by field number within the row by calling `mysql_fetch_field_direct()`. The current field cursor position may be changed by calling `mysql_field_seek()`. Setting the field cursor affects subsequent calls to `mysql_fetch_field()`. You can also get information for fields all at once by calling `mysql_fetch_fields()`.

For detecting and reporting errors, MySQL provides access to error information by means of the `mysql_errno()` and `mysql_error()` functions. These return the error code or error message for the most recently invoked function that can succeed or fail, allowing you to determine when an error occurred and what it was.

### 17.2.3. C API Function Descriptions

In the descriptions here, a parameter or return value of `NULL` means `NULL` in the sense of the C programming language, not a MySQL `NULL` value.

Functions that return a value generally return a pointer or an integer. Unless specified otherwise, functions returning a pointer return a non-`NULL` value to indicate success or a `NULL` value to indicate an error, and functions returning an integer return zero to indicate success or non-zero to indicate an error.

Note that “non-zero” means just that. Unless the function description says otherwise, do not test against a value other than zero:

```
if (result)                /* correct */
    ... error ...

if (result < 0)            /* incorrect */
    ... error ...

if (result == -1)         /* incorrect */
    ... error ...
```

When a function returns an error, the **Errors** subsection of the function description lists the possible types of errors. You can find out which of these occurred by calling `mysql_errno()`. A string representation of the error may be obtained by calling `mysql_error()`.

### 17.2.3.1. `mysql_affected_rows()`

```
my_ulonglong mysql_affected_rows(MYSQL *mysql)
```

#### Description

Returns the number of rows changed by the last `UPDATE`, deleted by the last `DELETE` or inserted by the last `INSERT` statement. May be called immediately after `mysql_query()` for `UPDATE`, `DELETE`, or `INSERT` statements. For `SELECT` statements, `mysql_affected_rows()` works like `mysql_num_rows()`.

#### Return Values

An integer greater than zero indicates the number of rows affected or retrieved. Zero indicates that no records were updated for an `UPDATE` statement, no rows matched the `WHERE` clause in the query or that no query has yet been executed. -1 indicates that the query returned an error or that, for a `SELECT` query, `mysql_affected_rows()` was called prior to calling `mysql_store_result()`. Because `mysql_affected_rows()` returns an unsigned value, you can check for -1 by comparing the return value to `(my_ulonglong)-1` (or to `(my_ulonglong)~0`, which is equivalent).

#### Errors

None.

#### Example

```
mysql_query(&mysql, "UPDATE products SET cost=cost*1.25 WHERE group=10");
printf("%ld products updated", (long) mysql_affected_rows(&mysql));
```

If you specify the flag `CLIENT_FOUND_ROWS` when connecting to `mysqld`, `mysql_affected_rows()` returns the number of rows matched by the `WHERE` statement for `UPDATE` statements. Otherwise, it returns the number of rows actually changed.

Note that when you use a `REPLACE` command, `mysql_affected_rows()` returns 2 if the new row replaced an old row, because in this case, one row was inserted after the duplicate was deleted.

If you use `INSERT ... ON DUPLICATE KEY UPDATE` to insert a row, `mysql_affected_rows()` returns 1 if the row is inserted as a new row and 2 if an existing row is updated.

### 17.2.3.2. `mysql_autocommit()`

```
my_bool mysql_autocommit(MYSQL *mysql, my_bool mode)
```

**Description**

Sets autocommit mode on if `mode` is 1, off if `mode` is 0.

This function was added in MySQL 4.1.0.

**Return Values**

Zero if successful. Non-zero if an error occurred.

**Errors**

None.

### 17.2.3.3. `mysql_change_user()`

```
my_bool mysql_change_user(MYSQL *mysql, const char *user, const char
*password, const char *db)
```

**Description**

Changes the user and causes the database specified by `db` to become the default (current) database on the connection specified by `mysql`. In subsequent queries, this database is the default for table references that do not include an explicit database specifier.

This function was introduced in MySQL 3.23.3.

`mysql_change_user()` fails if the connected user cannot be authenticated or doesn't have permission to use the database. In this case, the user and database are not changed.

The `db` parameter may be set to `NULL` if you don't want to have a default database.

Starting from MySQL 4.0.6 this command always performs a `ROLLBACK` of any active transactions, closes all temporary tables, unlocks all locked tables and resets the state as if one had done a new connect. This happens even if the user didn't change.

**Return Values**

Zero for success. Non-zero if an error occurred.

**Errors**

The same that you can get from `mysql_real_connect()`.

- `CR_COMMANDS_OUT_OF_SYNC`  
Commands were executed in an improper order.
- `CR_SERVER_GONE_ERROR`  
The MySQL server has gone away.
- `CR_SERVER_LOST`  
The connection to the server was lost during the query.
- `CR_UNKNOWN_ERROR`

An unknown error occurred.

- `ER_UNKNOWN_COM_ERROR`

The MySQL server doesn't implement this command (probably an old server).

- `ER_ACCESS_DENIED_ERROR`

The user or password was wrong.

- `ER_BAD_DB_ERROR`

The database didn't exist.

- `ER_DBACCESS_DENIED_ERROR`

The user did not have access rights to the database.

- `ER_WRONG_DB_NAME`

The database name was too long.

#### Example

```
if (mysql_change_user(&mysql, "user", "password", "new_database"))
{
    fprintf(stderr, "Failed to change user. Error: %s\n",
            mysql_error(&mysql));
}
```

### 17.2.3.4. `mysql_character_set_name()`

```
const char *mysql_character_set_name(MYSQL *mysql)
```

#### Description

Returns the default character set for the current connection.

#### Return Values

The default character set

#### Errors

None.

### 17.2.3.5. `mysql_close()`

```
void mysql_close(MYSQL *mysql)
```

#### Description

Closes a previously opened connection. `mysql_close()` also deallocates the connection handle pointed to by `mysql` if the handle was allocated automatically by `mysql_init()` or `mysql_connect()`.

#### Return Values

None.

**Errors**

None.

### 17.2.3.6. `mysql_commit()`

```
my_bool mysql_commit(MYSQL *mysql)
```

**Description**

Commits the current transaction.

This function was added in MySQL 4.1.0.

**Return Values**

Zero if successful. Non-zero if an error occurred.

**Errors**

None.

### 17.2.3.7. `mysql_connect()`

```
MYSQL *mysql_connect(MYSQL *mysql, const char *host, const char *user,  
const char *passwd)
```

**Description**

This function is deprecated. It is preferable to use `mysql_real_connect()` instead.

`mysql_connect()` attempts to establish a connection to a MySQL database engine running on `host`. `mysql_connect()` must complete successfully before you can execute any of the other API functions, with the exception of `mysql_get_client_info()`.

The meanings of the parameters are the same as for the corresponding parameters for `mysql_real_connect()` with the difference that the connection parameter may be `NULL`. In this case, the C API allocates memory for the connection structure automatically and frees it when you call `mysql_close()`. The disadvantage of this approach is that you can't retrieve an error message if the connection fails. (To get error information from `mysql_errno()` or `mysql_error()`, you must provide a valid `MYSQL` pointer.)

**Return Values**

Same as for `mysql_real_connect()`.

**Errors**

Same as for `mysql_real_connect()`.

### 17.2.3.8. `mysql_create_db()`

```
int mysql_create_db(MYSQL *mysql, const char *db)
```

**Description**

Creates the database named by the `db` parameter.

This function is deprecated. It is preferable to use `mysql_query()` to issue an SQL `CREATE DATABASE` statement instead.

### Return Values

Zero if the database was created successfully. Non-zero if an error occurred.

### Errors

- `CR_COMMANDS_OUT_OF_SYNC`  
Commands were executed in an improper order.
- `CR_SERVER_GONE_ERROR`  
The MySQL server has gone away.
- `CR_SERVER_LOST`  
The connection to the server was lost during the query.
- `CR_UNKNOWN_ERROR`  
An unknown error occurred.

### Example

```
if(mysql_create_db(&mysql, "my_database"))
{
    fprintf(stderr, "Failed to create new database. Error: %s\n",
            mysql_error(&mysql));
}
```

## 17.2.3.9. `mysql_data_seek()`

```
void mysql_data_seek(MYSQL_RES *result, my_ulonglong offset)
```

### Description

Seeks to an arbitrary row in a query result set. The `offset` value is a row number and should be in the range from 0 to `mysql_num_rows(result)-1`.

This function requires that the result set structure contains the entire result of the query, so `mysql_data_seek()` may be used only in conjunction with `mysql_store_result()`, not with `mysql_use_result()`.

### Return Values

None.

### Errors

None.

## 17.2.3.10. `mysql_debug()`

```
void mysql_debug(const char *debug)
```

**Description**

Does a `DEBUG_PUSH` with the given string. `mysql_debug()` uses the Fred Fish debug library. To use this function, you must compile the client library to support debugging. See [Section E.1, “Debugging a MySQL Server”](#), and [Section E.2, “Debugging a MySQL Client”](#).

**Return Values**

None.

**Errors**

None.

**Example**

The call shown here causes the client library to generate a trace file in `/tmp/client.trace` on the client machine:

```
mysql_debug("d:t:O,/tmp/client.trace");
```

### 17.2.3.11. `mysql_drop_db()`

```
int mysql_drop_db(MYSQL *mysql, const char *db)
```

**Description**

Drops the database named by the `db` parameter.

This function is deprecated. It is preferable to use `mysql_query()` to issue an SQL `DROP DATABASE` statement instead.

**Return Values**

Zero if the database was dropped successfully. Non-zero if an error occurred.

**Errors**

- `CR_COMMANDS_OUT_OF_SYNC`  
Commands were executed in an improper order.
- `CR_SERVER_GONE_ERROR`  
The MySQL server has gone away.
- `CR_SERVER_LOST`  
The connection to the server was lost during the query.
- `CR_UNKNOWN_ERROR`  
An unknown error occurred.

**Example**

```
if(mysql_drop_db(&mysql, "my_database"))
    fprintf(stderr, "Failed to drop the database: Error: %s\n",
            mysql_error(&mysql));
```

### 17.2.3.12. `mysql_dump_debug_info()`

```
int mysql_dump_debug_info(MYSQL *mysql)
```

#### Description

Instructs the server to write some debug information to the log. For this to work, the connected user must have the `SUPER` privilege.

#### Return Values

Zero if the command was successful. Non-zero if an error occurred.

#### Errors

- `CR_COMMANDS_OUT_OF_SYNC`  
Commands were executed in an improper order.
- `CR_SERVER_GONE_ERROR`  
The MySQL server has gone away.
- `CR_SERVER_LOST`  
The connection to the server was lost during the query.
- `CR_UNKNOWN_ERROR`  
An unknown error occurred.

### 17.2.3.13. `mysql_eof()`

```
my_bool mysql_eof(MYSQL_RES *result)
```

#### Description

This function is deprecated. `mysql_errno()` or `mysql_error()` may be used instead.

`mysql_eof()` determines whether the last row of a result set has been read.

If you acquire a result set from a successful call to `mysql_store_result()`, the client receives the entire set in one operation. In this case, a `NULL` return from `mysql_fetch_row()` always means the end of the result set has been reached and it is unnecessary to call `mysql_eof()`. When used with `mysql_store_result()`, `mysql_eof()` always returns true.

On the other hand, if you use `mysql_use_result()` to initiate a result set retrieval, the rows of the set are obtained from the server one by one as you call `mysql_fetch_row()` repeatedly. Because an error may occur on the connection during this process, a `NULL` return value from `mysql_fetch_row()` does not necessarily mean the end of the result set was reached normally. In this case, you can use `mysql_eof()` to determine what happened. `mysql_eof()` returns a non-zero value if the end of the result set was reached and zero if an error occurred.



Historically, `mysql_eof()` predates the standard MySQL error functions `mysql_errno()` and `mysql_error()`. Because those error functions provide the same information, their use is preferred over `mysql_eof()`, which is deprecated. (In fact, they provide more information, because `mysql_eof()` returns only a boolean value whereas the error functions indicate a reason for the error when one occurs.)

### Return Values

Zero if no error occurred. Non-zero if the end of the result set has been reached.

### Errors

None.

### Example

The following example shows how you might use `mysql_eof()`:

```
mysql_query(&mysql, "SELECT * FROM some_table");
result = mysql_use_result(&mysql);
while((row = mysql_fetch_row(result)))
{
    // do something with data
}
if(!mysql_eof(result)) // mysql_fetch_row() failed due to an error
{
    fprintf(stderr, "Error: %s\n", mysql_error(&mysql));
}
```

However, you can achieve the same effect with the standard MySQL error functions:

```
mysql_query(&mysql, "SELECT * FROM some_table");
result = mysql_use_result(&mysql);
while((row = mysql_fetch_row(result)))
{
    // do something with data
}
if(mysql_errno(&mysql)) // mysql_fetch_row() failed due to an error
{
    fprintf(stderr, "Error: %s\n", mysql_error(&mysql));
}
```

## 17.2.3.14. `mysql_errno()`

```
unsigned int mysql_errno(MYSQL *mysql)
```

### Description

For the connection specified by `mysql`, `mysql_errno()` returns the error code for the most recently invoked API function that can succeed or fail. A return value of zero means that no error occurred. Client error message numbers are listed in the MySQL `errmsg.h` header file. Server error message numbers are listed in `mysqld_error.h`. Errors also are listed at [Appendix B, Error Codes and Messages](#).

Note that some functions like `mysql_fetch_row()` don't set `mysql_errno()` if they succeed.

A rule of thumb is that all functions that have to ask the server for information reset `mysql_errno()` if they succeed.

MySQL-specific error numbers returned by `mysql_errno()` differ from SQLSTATE values returned by `mysql_sqlstate()`. For example, the `mysql` client program displays errors using the following format, where 1146 is the `mysql_errno()` value and '42S02' is the corresponding `mysql_sqlstate()` value:

```
shell> SELECT * FROM no_such_table;
ERROR 1146 (42S02): Table 'test.no_such_table' doesn't exist
```

### Return Values

An error code value for the last `mysql_xxx()` call, if it failed. zero means no error occurred.

### Errors

None.

## 17.2.3.15. `mysql_error()`

```
const char *mysql_error(MYSQL *mysql)
```

### Description

For the connection specified by `mysql`, `mysql_error()` returns a null-terminated string containing the error message for the most recently invoked API function that failed. If a function didn't fail, the return value of `mysql_error()` may be the previous error or an empty string to indicate no error.

A rule of thumb is that all functions that have to ask the server for information reset `mysql_error()` if they succeed.

For functions that reset `mysql_errno()`, the following two tests are equivalent:

```
if(mysql_errno(&mysql))
{
    // an error occurred
}

if(mysql_error(&mysql)[0] != '\0')
{
    // an error occurred
}
```

The language of the client error messages may be changed by recompiling the MySQL client library. Currently, you can choose error messages in several different languages. See [Section 5.10.2, “Setting the Error Message Language”](#).

### Return Values

A null-terminated character string that describes the error. An empty string if no error occurred.

### Errors

None.

## 17.2.3.16. `mysql_escape_string()`

You should use `mysql_real_escape_string()` instead!

This function is identical to `mysql_real_escape_string()` except that `mysql_real_escape_string()` takes a connection handler as its first argument and escapes the string according to the current character set. `mysql_escape_string()` does not take a connection argument and does not respect the current character set.

## 17.2.3.17. `mysql_fetch_field()`

```
MYSQL_FIELD *mysql_fetch_field(MYSQL_RES *result)
```

### Description

Returns the definition of one column of a result set as a `MYSQL_FIELD` structure. Call this function repeatedly to retrieve information about all columns in the result set. `mysql_fetch_field()` returns `NULL` when no more fields are left.

`mysql_fetch_field()` is reset to return information about the first field each time you execute a new `SELECT` query. The field returned by `mysql_fetch_field()` is also affected by calls to `mysql_field_seek()`.

If you've called `mysql_query()` to perform a `SELECT` on a table but have not called `mysql_store_result()`, MySQL returns the default blob length (8KB) if you call `mysql_fetch_field()` to ask for the length of a `BLOB` field. (The 8KB size is chosen because MySQL doesn't know the maximum length for the `BLOB`. This should be made configurable sometime.) Once you've retrieved the result set, `field->max_length` contains the length of the largest value for this column in the specific query.

### Return Values

The `MYSQL_FIELD` structure for the current column. `NULL` if no columns are left.

### Errors

None.

### Example

```
MYSQL_FIELD *field;

while((field = mysql_fetch_field(result)))
{
    printf("field name %s\n", field->name);
}
```

## 17.2.3.18. `mysql_fetch_field_direct()`

```
MYSQL_FIELD *mysql_fetch_field_direct(MYSQL_RES *result, unsigned int
fieldnr)
```

### Description

Given a field number `fieldnr` for a column within a result set, returns that column's field definition as a `MYSQL_FIELD` structure. You may use this function to retrieve the definition for an arbitrary column. The value of `fieldnr` should be in the range from 0 to `mysql_num_fields(result)-1`.

### Return Values

The `MYSQL_FIELD` structure for the specified column.

### Errors

None.

### Example

```
unsigned int num_fields;
unsigned int i;
MYSQL_FIELD *field;

num_fields = mysql_num_fields(result);
for(i = 0; i < num_fields; i++)
{
```

```

    field = mysql_fetch_field_direct(result, i);
    printf("Field %u is %s\n", i, field->name);
}

```

### 17.2.3.19. `mysql_fetch_fields()`

```
MYSQL_FIELD *mysql_fetch_fields(MYSQL_RES *result)
```

#### Description

Returns an array of all `MYSQL_FIELD` structures for a result set. Each structure provides the field definition for one column of the result set.

#### Return Values

An array of `MYSQL_FIELD` structures for all columns of a result set.

#### Errors

None.

#### Example

```

unsigned int num_fields;
unsigned int i;
MYSQL_FIELD *fields;

num_fields = mysql_num_fields(result);
fields = mysql_fetch_fields(result);
for(i = 0; i < num_fields; i++)
{
    printf("Field %u is %s\n", i, fields[i].name);
}

```

### 17.2.3.20. `mysql_fetch_lengths()`

```
unsigned long *mysql_fetch_lengths(MYSQL_RES *result)
```

#### Description

Returns the lengths of the columns of the current row within a result set. If you plan to copy field values, this length information is also useful for optimization, because you can avoid calling `strlen()`. In addition, if the result set contains binary data, you **must** use this function to determine the size of the data, because `strlen()` returns incorrect results for any field containing null characters.

The length for empty columns and for columns containing `NULL` values is zero. To see how to distinguish these two cases, see the description for `mysql_fetch_row()`.

#### Return Values

An array of unsigned long integers representing the size of each column (not including any terminating null characters). `NULL` if an error occurred.

#### Errors

`mysql_fetch_lengths()` is valid only for the current row of the result set. It returns `NULL` if you call it before calling `mysql_fetch_row()` or after retrieving all rows in the result.

#### Example

```
MYSQL_ROW row;
```

```

unsigned long *lengths;
unsigned int num_fields;
unsigned int i;

row = mysql_fetch_row(result);
if (row)
{
    num_fields = mysql_num_fields(result);
    lengths = mysql_fetch_lengths(result);
    for(i = 0; i < num_fields; i++)
    {
        printf("Column %u is %lu bytes in length.\n", i, lengths[i]);
    }
}

```

### 17.2.3.21. `mysql_fetch_row()`

`MYSQL_ROW mysql_fetch_row(MYSQL_RES *result)`

#### Description

Retrieves the next row of a result set. When used after `mysql_store_result()`, `mysql_fetch_row()` returns `NULL` when there are no more rows to retrieve. When used after `mysql_use_result()`, `mysql_fetch_row()` returns `NULL` when there are no more rows to retrieve or if an error occurred.

The number of values in the row is given by `mysql_num_fields(result)`. If `row` holds the return value from a call to `mysql_fetch_row()`, pointers to the values are accessed as `row[0]` to `row[mysql_num_fields(result)-1]`. `NULL` values in the row are indicated by `NULL` pointers.

The lengths of the field values in the row may be obtained by calling `mysql_fetch_lengths()`. Empty fields and fields containing `NULL` both have length 0; you can distinguish these by checking the pointer for the field value. If the pointer is `NULL`, the field is `NULL`; otherwise, the field is empty.

#### Return Values

A `MYSQL_ROW` structure for the next row. `NULL` if there are no more rows to retrieve or if an error occurred.

#### Errors

Note that error is not reset between calls to `mysql_fetch_row()`

- `CR_SERVER_LOST`

The connection to the server was lost during the query.

- `CR_UNKNOWN_ERROR`

An unknown error occurred.

#### Example

```

MYSQL_ROW row;
unsigned int num_fields;
unsigned int i;

num_fields = mysql_num_fields(result);
while ((row = mysql_fetch_row(result)))
{
    unsigned long *lengths;
    lengths = mysql_fetch_lengths(result);
    for(i = 0; i < num_fields; i++)

```

```

{
    printf("[%.*s] ", (int) lengths[i], row[i] ? row[i] : "NULL");
}
printf("\n");
}

```

### 17.2.3.22. `mysql_field_count()`

```
unsigned int mysql_field_count(MYSQL *mysql)
```

If you are using a version of MySQL earlier than 3.22.24, you should use `unsigned int mysql_num_fields(MYSQL *mysql)` instead.

#### Description

Returns the number of columns for the most recent query on the connection.

The normal use of this function is when `mysql_store_result()` returned `NULL` (and thus you have no result set pointer). In this case, you can call `mysql_field_count()` to determine whether `mysql_store_result()` should have produced a non-empty result. This allows the client program to take proper action without knowing whether the query was a `SELECT` (or `SELECT`-like) statement. The example shown here illustrates how this may be done.

See [Section 17.2.13.1, “Why `mysql\_store\_result\(\)` Sometimes Returns `NULL` After `mysql\_query\(\)` Returns Success”](#).

#### Return Values

An unsigned integer representing the number of columns in a result set.

#### Errors

None.

#### Example

```

MYSQL_RES *result;
unsigned int num_fields;
unsigned int num_rows;

if (mysql_query(&mysql, query_string))
{
    // error
}
else // query succeeded, process any data returned by it
{
    result = mysql_store_result(&mysql);
    if (result) // there are rows
    {
        num_fields = mysql_num_fields(result);
        // retrieve rows, then call mysql_free_result(result)
    }
    else // mysql_store_result() returned nothing; should it have?
    {
        if(mysql_field_count(&mysql) == 0)
        {
            // query does not return data
            // (it was not a SELECT)
            num_rows = mysql_affected_rows(&mysql);
        }
        else // mysql_store_result() should have returned data
        {
            fprintf(stderr, "Error: %s\n", mysql_error(&mysql));
        }
    }
}

```

An alternative is to replace the `mysql_field_count(&mysql)` call with `mysql_errno(&mysql)`. In this case, you are checking directly for an error from `mysql_store_result()` rather than inferring from the value of `mysql_field_count()` whether the statement was a `SELECT`.

### 17.2.3.23. `mysql_field_seek()`

```
MYSQL_FIELD_OFFSET mysql_field_seek(MYSQL_RES *result,  
MYSQL_FIELD_OFFSET offset)
```

#### Description

Sets the field cursor to the given offset. The next call to `mysql_fetch_field()` retrieves the field definition of the column associated with that offset.

To seek to the beginning of a row, pass an `offset` value of zero.

#### Return Values

The previous value of the field cursor.

#### Errors

None.

### 17.2.3.24. `mysql_field_tell()`

```
MYSQL_FIELD_OFFSET mysql_field_tell(MYSQL_RES *result)
```

#### Description

Returns the position of the field cursor used for the last `mysql_fetch_field()`. This value can be used as an argument to `mysql_field_seek()`.

#### Return Values

The current offset of the field cursor.

#### Errors

None.

### 17.2.3.25. `mysql_free_result()`

```
void mysql_free_result(MYSQL_RES *result)
```

#### Description

Frees the memory allocated for a result set by `mysql_store_result()`, `mysql_use_result()`, `mysql_list_dbs()`, and so forth. When you are done with a result set, you must free the memory it uses by calling `mysql_free_result()`.

Do not attempt to access a result set after freeing it.

#### Return Values

None.

**Errors**

None.

**17.2.3.26. `mysql_get_client_info()`**

```
const char *mysql_get_client_info(void)
```

**Description**

Returns a string that represents the client library version.

**Return Values**

A character string that represents the MySQL client library version.

**Errors**

None.

**17.2.3.27. `mysql_get_client_version()`**

```
unsigned long mysql_get_client_version(void)
```

**Description**

Returns an integer that represents the client library version. The value has the format `XYZZ` where `X` is the major version, `YY` is the release level, and `ZZ` is the version number within the release level. For example, a value of `40102` represents a client library version of `4.1.2`.

This function was added in MySQL 4.0.16.

**Return Values**

An integer that represents the MySQL client library version.

**Errors**

None.

**17.2.3.28. `mysql_get_host_info()`**

```
const char *mysql_get_host_info(MYSQL *mysql)
```

**Description**

Returns a string describing the type of connection in use, including the server hostname.

**Return Values**

A character string representing the server hostname and the connection type.

**Errors**

None.

**17.2.3.29. `mysql_get_proto_info()`**



```
unsigned int mysql_get_proto_info(MYSQL *mysql)
```

**Description**

Returns the protocol version used by current connection.

**Return Values**

An unsigned integer representing the protocol version used by the current connection.

**Errors**

None.

### 17.2.3.30. `mysql_get_server_info()`

```
const char *mysql_get_server_info(MYSQL *mysql)
```

**Description**

Returns a string that represents the server version number.

**Return Values**

A character string that represents the server version number.

**Errors**

None.

### 17.2.3.31. `mysql_get_server_version()`

```
unsigned long mysql_get_server_version(MYSQL *mysql)
```

**Description**

Returns the version number of the server as an integer.

This function was added in MySQL 4.1.0.

**Return Values**

A number that represents the MySQL server version in this format:

```
major_version*10000 + minor_version *100 + sub_version
```

For example, 4.1.2 is returned as 40102.

This function is useful in client programs for quickly determining whether some version-specific server capability exists.

**Errors**

None.

### 17.2.3.32. `mysql_hex_string()`

```
unsigned long mysql_hex_string(char *to, const char *from, unsigned
```

---

long length)

### Description

This function is used to create a legal SQL string that you can use in an SQL statement. See [Section 9.1.1, “Strings”](#).

The string in `from` is encoded to hexadecimal format, with each character encoded as two hexadecimal digits. The result is placed in `to` and a terminating null byte is appended.

The string pointed to by `from` must be `length` bytes long. You must allocate the `to` buffer to be at least `length*2+1` bytes long. When `mysql_hex_string()` returns, the contents of `to` is a null-terminated string. The return value is the length of the encoded string, not including the terminating null character.

The return value can be placed into an SQL statement using either `0xvalue` or `X'value'` format. However, the return value does not include the `0x` or `X'...`. The caller must supply whichever of those is desired.

`mysql_hex_string()` was added in MySQL 4.0.23 and 4.1.8.

### Example

```
char query[1000],*end;

end = strmov(query,"INSERT INTO test_table values(");
end = strmov(end,"0x");
end += mysql_hex_string(end,"What's this",11);
end = strmov(end,"0x");
end += mysql_hex_string(end,"binary data: \0\r\n",16);
*end++ = '\0';

if (mysql_real_query(&mysql,query,(unsigned int) (end - query)))
{
    fprintf(stderr, "Failed to insert row, Error: %s\n",
            mysql_error(&mysql));
}
```

The `strmov()` function used in the example is included in the `mysqlclient` library and works like `strcpy()` but returns a pointer to the terminating null of the first parameter.

### Return Values

The length of the value placed into `to`, not including the terminating null character.

### Errors

None.

## 17.2.3.33. `mysql_info()`

```
const char *mysql_info(MYSQL *mysql)
```

### Description

Retrieves a string providing information about the most recently executed query, but only for the statements listed here. For other statements, `mysql_info()` returns `NULL`. The format of the string varies depending on the type of query, as described here. The numbers are illustrative only; the string contains values appropriate for the query.

- `INSERT INTO ... SELECT ...`

String format: Records: 100 Duplicates: 0 Warnings: 0

- `INSERT INTO ... VALUES (...),(...),(...)`...

String format: Records: 3 Duplicates: 0 Warnings: 0

- `LOAD DATA INFILE ...`

String format: Records: 1 Deleted: 0 Skipped: 0 Warnings: 0

- `ALTER TABLE`

String format: Records: 3 Duplicates: 0 Warnings: 0

- `UPDATE`

String format: Rows matched: 40 Changed: 40 Warnings: 0

Note that `mysql_info()` returns a non-NULL value for `INSERT ... VALUES` only for the multiple-row form of the statement (that is, only if multiple value lists are specified).

#### Return Values

A character string representing additional information about the most recently executed query. `NULL` if no information is available for the query.

#### Errors

None.

### 17.2.3.34. `mysql_init()`

```
MYSQL *mysql_init(MYSQL *mysql)
```

#### Description

Allocates or initializes a `MYSQL` object suitable for `mysql_real_connect()`. If `mysql` is a `NULL` pointer, the function allocates, initializes, and returns a new object. Otherwise, the object is initialized and the address of the object is returned. If `mysql_init()` allocates a new object, it is freed when `mysql_close()` is called to close the connection.

#### Return Values

An initialized `MYSQL*` handle. `NULL` if there was insufficient memory to allocate a new object.

#### Errors

In case of insufficient memory, `NULL` is returned.

### 17.2.3.35. `mysql_insert_id()`

```
my_ulonglong mysql_insert_id(MYSQL *mysql)
```

#### Description

Returns the value generated for an `AUTO_INCREMENT` column by the previous `INSERT` or `UPDATE` statement. Use this function after you have performed an `INSERT` statement into a table that contains an

`AUTO_INCREMENT` field.

More precisely, `mysql_insert_id()` is updated under these conditions:

- `INSERT` statements that store a value into an `AUTO_INCREMENT` column. This is true whether the value is automatically generated by storing the special values `NULL` or `0` into the column, or is an explicit non-special value.
- In the case of a multiple-row `INSERT` statement, `mysql_insert_id()` returns the **first** automatically generated `AUTO_INCREMENT` value; if no such value is generated, it returns the last **last** explicit value inserted into the `AUTO_INCREMENT` column.
- `INSERT` statements that generate an `AUTO_INCREMENT` value by inserting `LAST_INSERT_ID(expr)` into any column.
- `INSERT` statements that generate an `AUTO_INCREMENT` value by updating any column to `LAST_INSERT_ID(expr)`.
- The value of `mysql_insert_id()` is not affected by statements such as `SELECT` that return a result set.
- If the previous statement returned an error, the value of `mysql_insert_id()` is undefined.

Note that `mysql_insert_id()` returns `0` if the previous statement does not use an `AUTO_INCREMENT` value. If you need to save the value for later, be sure to call `mysql_insert_id()` immediately after the statement that generates the value.

The value of `mysql_insert_id()` is affected only by statements issued within the current client connection. It is not affected by statements issued by other clients.

See [Section 12.9.3, “Information Functions”](#).

Also note that the value of the SQL `LAST_INSERT_ID()` function always contains the most recently generated `AUTO_INCREMENT` value, and is not reset between statements because the value of that function is maintained in the server. Another difference is that `LAST_INSERT_ID()` is not updated if you set an `AUTO_INCREMENT` column to a specific non-special value.

The reason for the difference between `LAST_INSERT_ID()` and `mysql_insert_id()` is that `LAST_INSERT_ID()` is made easy to use in scripts while `mysql_insert_id()` tries to provide a little more exact information of what happens to the `AUTO_INCREMENT` column.

### Return Values

Described in the preceding discussion.

### Errors

None.

## 17.2.3.36. `mysql_kill()`

```
int mysql_kill(MYSQL *mysql, unsigned long pid)
```

### Description

Asks the server to kill the thread specified by `pid`.

### Return Values

Zero for success. Non-zero if an error occurred.

### Errors

- `CR_COMMANDS_OUT_OF_SYNC`  
Commands were executed in an improper order.
- `CR_SERVER_GONE_ERROR`  
The MySQL server has gone away.
- `CR_SERVER_LOST`  
The connection to the server was lost during the query.
- `CR_UNKNOWN_ERROR`  
An unknown error occurred.

### 17.2.3.37. `mysql_library_end()`

```
void mysql_library_end(void)
```

#### Description

This is a synonym for the `mysql_server_end()` function. It was added in MySQL 4.1.10.

See [Section 17.2.2, “C API Function Overview”](#), for usage information.

### 17.2.3.38. `mysql_library_init()`

```
int mysql_library_init(int argc, char **argv, char **groups)
```

#### Description

This is a synonym for the `mysql_server_init()` function. It was added in MySQL 4.1.10.

See [Section 17.2.2, “C API Function Overview”](#), for usage information.

### 17.2.3.39. `mysql_list_dbs()`

```
MYSQL_RES *mysql_list_dbs(MYSQL *mysql, const char *wild)
```

#### Description

Returns a result set consisting of database names on the server that match the simple regular expression specified by the `wild` parameter. `wild` may contain the wildcard characters ‘%’ or ‘\_’, or may be a `NULL` pointer to match all databases. Calling `mysql_list_dbs()` is similar to executing the query `SHOW databases [LIKE wild]`.

You must free the result set with `mysql_free_result()`.

#### Return Values

A `MYSQL_RES` result set for success. `NULL` if an error occurred.

**Errors**

- `CR_COMMANDS_OUT_OF_SYNC`  
Commands were executed in an improper order.
- `CR_OUT_OF_MEMORY`  
Out of memory.
- `CR_SERVER_GONE_ERROR`  
The MySQL server has gone away.
- `CR_SERVER_LOST`  
The connection to the server was lost during the query.
- `CR_UNKNOWN_ERROR`  
An unknown error occurred.

**17.2.3.40. `mysql_list_fields()`**

```
MYSQL_RES *mysql_list_fields(MYSQL *mysql, const char *table, const char *wild)
```

**Description**

Returns a result set consisting of field names in the given table that match the simple regular expression specified by the *wild* parameter. *wild* may contain the wildcard characters '%' or '\_', or may be a `NULL` pointer to match all fields. Calling `mysql_list_fields()` is similar to executing the query `SHOW COLUMNS FROM tbl_name [LIKE wild]`.

Note that it's recommended that you use `SHOW COLUMNS FROM tbl_name` instead of `mysql_list_fields()`.

You must free the result set with `mysql_free_result()`.

**Return Values**

A `MYSQL_RES` result set for success. `NULL` if an error occurred.

**Errors**

- `CR_COMMANDS_OUT_OF_SYNC`  
Commands were executed in an improper order.
- `CR_SERVER_GONE_ERROR`  
The MySQL server has gone away.
- `CR_SERVER_LOST`  
The connection to the server was lost during the query.

- `CR_UNKNOWN_ERROR`

An unknown error occurred.

### 17.2.3.41. `mysql_list_processes()`

```
MYSQL_RES *mysql_list_processes(MYSQL *mysql)
```

#### Description

Returns a result set describing the current server threads. This is the same kind of information as that reported by `mysqladmin processlist` or a `SHOW PROCESSLIST` query.

You must free the result set with `mysql_free_result()`.

#### Return Values

A `MYSQL_RES` result set for success. `NULL` if an error occurred.

#### Errors

- `CR_COMMANDS_OUT_OF_SYNC`

Commands were executed in an improper order.

- `CR_SERVER_GONE_ERROR`

The MySQL server has gone away.

- `CR_SERVER_LOST`

The connection to the server was lost during the query.

- `CR_UNKNOWN_ERROR`

An unknown error occurred.

### 17.2.3.42. `mysql_list_tables()`

```
MYSQL_RES *mysql_list_tables(MYSQL *mysql, const char *wild)
```

#### Description

Returns a result set consisting of table names in the current database that match the simple regular expression specified by the `wild` parameter. `wild` may contain the wildcard characters `'%'` or `'_'`, or may be a `NULL` pointer to match all tables. Calling `mysql_list_tables()` is similar to executing the query `SHOW tables [LIKE wild]`.

You must free the result set with `mysql_free_result()`.

#### Return Values

A `MYSQL_RES` result set for success. `NULL` if an error occurred.

#### Errors

- `CR_COMMANDS_OUT_OF_SYNC`  
Commands were executed in an improper order.
- `CR_SERVER_GONE_ERROR`  
The MySQL server has gone away.
- `CR_SERVER_LOST`  
The connection to the server was lost during the query.
- `CR_UNKNOWN_ERROR`  
An unknown error occurred.

### 17.2.3.43. `mysql_more_results()`

```
my_bool mysql_more_results(MYSQL *mysql)
```

#### Description

Returns true if more results exist from the currently executed query, and the application must call `mysql_next_result()` to fetch the results.

This function was added in MySQL 4.1.0.

#### Return Values

`TRUE` (1) if more results exist. `FALSE` (0) if no more results exist.

In most cases, you can call `mysql_next_result()` instead to test whether more results exist and initiate retrieval if so.

See [Section 17.2.9, “C API Handling of Multiple Statement Execution”](#), and [Section 17.2.3.44, “`mysql\_next\_result\(\)`”](#).

#### Errors

None.

### 17.2.3.44. `mysql_next_result()`

```
int mysql_next_result(MYSQL *mysql)
```

#### Description

If more query results exist, `mysql_next_result()` reads the next query results and returns the status back to application.

You must call `mysql_free_result()` for the preceding query if it returned a result set.

After calling `mysql_next_result()` the state of the connection is as if you had called `mysql_real_query()` or `mysql_query()` for the next query. This means that you can call `mysql_store_result()`, `mysql_warning_count()`, `mysql_affected_rows()`, and so forth.

If `mysql_next_result()` returns an error, no other statements are executed and there are no more



results to fetch.

For an example that shows how to use `mysql_next_result()`, see [Section 17.2.9, “C API Handling of Multiple Statement Execution”](#). See [Section 17.2.9, “C API Handling of Multiple Statement Execution”](#).

This function was added in MySQL 4.1.0.

### Return Values

Return Value	Description
0	Successful and there are more results
-1	Successful and there are no more results
>0	An error occurred

### Errors

- [CR\\_COMMANDS\\_OUT\\_OF\\_SYNC](#)  
Commands were executed in an improper order. For example if you didn't call `mysql_use_result()` for a previous result set.
- [CR\\_SERVER\\_GONE\\_ERROR](#)  
The MySQL server has gone away.
- [CR\\_SERVER\\_LOST](#)  
The connection to the server was lost during the query.
- [CR\\_UNKNOWN\\_ERROR](#)  
An unknown error occurred.

### 17.2.3.45. `mysql_num_fields()`

```
unsigned int mysql_num_fields(MYSQL_RES *result)
```

Or:

```
unsigned int mysql_num_fields(MYSQL *mysql)
```

The second form doesn't work on MySQL 3.22.24 or newer. To pass a `MYSQL*` argument, you must use `unsigned int mysql_field_count(MYSQL *mysql)` instead.

#### Description

Returns the number of columns in a result set.

Note that you can get the number of columns either from a pointer to a result set or to a connection handle. You would use the connection handle if `mysql_store_result()` or `mysql_use_result()` returned `NULL` (and thus you have no result set pointer). In this case, you can call `mysql_field_count()` to determine whether `mysql_store_result()` should have produced a non-empty result. This allows the client program to take proper action without knowing whether the query was a `SELECT` (or `SELECT`-like) statement. The example shown here illustrates how this may

be done.

See [Section 17.2.13.1, “Why `mysql\_store\_result\(\)` Sometimes Returns NULL After `mysql\_query\(\)` Returns Success”](#).

### Return Values

An unsigned integer representing the number of columns in a result set.

### Errors

None.

### Example

```
MYSQL_RES *result;
unsigned int num_fields;
unsigned int num_rows;

if (mysql_query(&mysql, query_string))
{
    // error
}
else // query succeeded, process any data returned by it
{
    result = mysql_store_result(&mysql);
    if (result) // there are rows
    {
        num_fields = mysql_num_fields(result);
        // retrieve rows, then call mysql_free_result(result)
    }
    else // mysql_store_result() returned nothing; should it have?
    {
        if (mysql_errno(&mysql))
        {
            fprintf(stderr, "Error: %s\n", mysql_error(&mysql));
        }
        else if (mysql_field_count(&mysql) == 0)
        {
            // query does not return data
            // (it was not a SELECT)
            num_rows = mysql_affected_rows(&mysql);
        }
    }
}
```

An alternative (if you know that your query should have returned a result set) is to replace the `mysql_errno(&mysql)` call with a check whether `mysql_field_count(&mysql)` is = 0. This happens only if something went wrong.

## 17.2.3.46. `mysql_num_rows()`

```
my_ulonglong mysql_num_rows(MYSQL_RES *result)
```

### Description

Returns the number of rows in the result set.

The use of `mysql_num_rows()` depends on whether you use `mysql_store_result()` or `mysql_use_result()` to return the result set. If you use `mysql_store_result()`, `mysql_num_rows()` may be called immediately. If you use `mysql_use_result()`, `mysql_num_rows()` does not return the correct value until all the rows in the result set have been retrieved.

### Return Values

The number of rows in the result set.

**Errors**

None.

**17.2.3.47. `mysql_options()`**

```
int mysql_options(MYSQL *mysql, enum mysql_option option, const char *arg)
```

**Description**

Can be used to set extra connect options and affect behavior for a connection. This function may be called multiple times to set several options.

`mysql_options()` should be called after `mysql_init()` and before `mysql_connect()` or `mysql_real_connect()`.

The `option` argument is the option that you want to set; the `arg` argument is the value for the option. If the option is an integer, then `arg` should point to the value of the integer.

Possible option values:

Option	Argument Type	Function
<code>MYSQL_INIT_COMMAND</code>	<code>char *</code>	Command to execute when connecting to the MySQL server. Will automatically be re-executed when reconnecting.
<code>MYSQL_OPT_COMPRESS</code>	Not used	Use the compressed client/server protocol.
<code>MYSQL_OPT_CONNECT_TIMEOUT</code>	<code>unsigned int *</code>	Connect timeout in seconds.
<code>MYSQL_OPT_GUESS_CONNECTION</code>	Not used	For an application linked against <code>libmysqld</code> , this allows the library to “guess” whether to use the embedded server or a remote server. “Guess” means that if the hostname is set and is not <code>localhost</code> , it uses a remote server. This behavior is the default. <code>MYSQL_OPT_USE_EMBEDDED_CONNECTION</code> and <code>MYSQL_OPT_USE_REMOTE_CONNECTION</code> can be used to override it. This option is ignored for applications linked against <code>libmysqlclient</code> . Available starting in 4.1.1.
<code>MYSQL_OPT_LOCAL_INFILE</code>	optional pointer to <code>uint</code>	If no pointer is given or if pointer points to an <code>unsigned int != 0</code> the command <code>LOAD LOCAL INFILE</code> is enabled.
<code>MYSQL_OPT_NAMED_PIPE</code>	Not used	Use named pipes to connect to a MySQL server on NT.
<code>MYSQL_OPT_PROTOCOL</code>	<code>unsigned int *</code>	Type of protocol to use. Should be one of the enum values of <code>mysql_protocol_type</code> defined in <code>mysql.h</code> . Added in 4.1.0.

<code>MYSQL_OPT_READ_TIMEOUT</code>	<code>unsigned int *</code>	Timeout for reads from server (works currently only on Windows on TCP/IP connections). Added in 4.1.1.
<code>MYSQL_OPT_SET_CLIENT_IP</code>	<code>char *</code>	For an application linked against <code>libmysqld</code> (with <code>libmysqld</code> compiled with authentication support), this means that the user is considered to have connected from the specified IP address (specified as a string) for authentication purposes. This option is ignored for applications linked against <code>libmysqlclient</code> . Added in 4.1.1.
<code>MYSQL_OPT_USE_EMBEDDED_CONNECTION</code>	Not used	For an application linked against <code>libmysqld</code> , this forces the use of the embedded server for the connection. This option is ignored for applications linked against <code>libmysqlclient</code> . Added in 4.1.1.
<code>MYSQL_OPT_USE_REMOTE_CONNECTION</code>	Not used	For an application linked against <code>libmysqld</code> , this forces the use of a remote server for the connection. This option is ignored for applications linked against <code>libmysqlclient</code> . Added in 4.1.1.
<code>MYSQL_OPT_USE_RESULT</code>	Not used	This option is available beginning with MySQL 4.1.1, but is unused.
<code>MYSQL_OPT_WRITE_TIMEOUT</code>	<code>unsigned int *</code>	Timeout for writes to server (works currently only on Windows on TCP/IP connections). Added in 4.1.1.
<code>MYSQL_READ_DEFAULT_FILE</code>	<code>char *</code>	Read options from the named option file instead of from <code>my.cnf</code> .
<code>MYSQL_READ_DEFAULT_GROUP</code>	<code>char *</code>	Read options from the named group from <code>my.cnf</code> or the file specified with <code>MYSQL_READ_DEFAULT_FILE</code> .
<code>MYSQL_REPORT_DATA_TRUNCATION</code>	<code>my_bool *</code>	Enable or disable reporting of data truncation errors for prepared statements via <code>MYSQL_BIND.error</code> . (Default: disabled.)
<code>MYSQL_SECURE_AUTH</code>	<code>my_bool*</code>	Whether to connect to a server that does not support the improved password hashing available beginning in MySQL 4.1.1. This option was added in MySQL 4.1.1.
<code>MYSQL_SET_CHARSET_DIR</code>	<code>char*</code>	The pathname to the directory that contains character set definition files.
<code>MYSQL_SET_CHARSET_NAME</code>	<code>char*</code>	The name of the character set to use as the default character set.
<code>MYSQL_SHARED_MEMORY_BASE_NAME</code>	<code>char*</code>	Named of shared-memory object for communication to server. Should be same as the option <code>-shared-memory-base-name</code>

		used for the <code>mysqld</code> server you want to connect to. Added in 4.1.0.
--	--	---

Note that the `client` group is always read if you use `MYSQL_READ_DEFAULT_FILE` or `MYSQL_READ_DEFAULT_GROUP`.

The specified group in the option file may contain the following options:

Option	Description
<code>connect-timeout</code>	Connect timeout in seconds. On Linux this timeout is also used for waiting for the first answer from the server.
<code>compress</code>	Use the compressed client/server protocol.
<code>database</code>	Connect to this database if no database was specified in the connect command.
<code>debug</code>	Debug options.
<code>disable-local-infile</code>	Disable use of <code>LOAD DATA LOCAL</code> .
<code>host</code>	Default hostname.
<code>init-command</code>	Command to execute when connecting to MySQL server. Will automatically be re-executed when reconnecting.
<code>interactive-timeout</code>	Same as specifying <code>CLIENT_INTERACTIVE</code> to <code>mysql_real_connect()</code> . See <a href="#">Section 17.2.3.50</a> , “ <code>mysql_real_connect()</code> ”.
<code>local-infile[=(0 1)]</code>	If no argument or argument <code>!= 0</code> then enable use of <code>LOAD DATA LOCAL</code> .
<code>max_allowed_packet</code>	Max size of packet client can read from server.
<code>multi-results</code>	Allow multiple result sets from multiple-statement executions or stored procedures. Added in 4.1.1.
<code>multi-statements</code>	Allow the client to send multiple statements in a single string (separated by ‘;’). Added in 4.1.9.
<code>password</code>	Default password.
<code>pipe</code>	Use named pipes to connect to a MySQL server on NT.
<code>protocol={TCP SOCKET PIPE MEMORY}</code>	The protocol to use when connecting to the server. (Added in MySQL 4.1)
<code>port</code>	Default port number.
<code>return-found-rows</code>	Tell <code>mysql_info()</code> to return found rows instead of updated rows when using <code>UPDATE</code> .
<code>shared-memory-base-name=me</code>	Shared-memory name to use to connect to server (default is “MYSQL”). Added in MySQL 4.1.
<code>socket</code>	Default socket file.
<code>user</code>	Default user.

Note that `timeout` has been replaced by `connect-timeout`, but `timeout` still works for a while.

For more information about option files, see [Section 4.3.2](#), “Using Option Files”.

**Return Values**

Zero for success. Non-zero if you used an unknown option.

**Example**

```
MYSQL mysql;

mysql_init(&mysql);
mysql_options(&mysql,MYSQL_OPT_COMPRESS,0);
mysql_options(&mysql,MYSQL_READ_DEFAULT_GROUP,"odbc");
if (!mysql_real_connect(&mysql,"host","user","passwd","database",0,NULL,0))
{
    fprintf(stderr,"Failed to connect to database: Error: %s\n",
            mysql_error(&mysql));
}
```

This code requests the client to use the compressed client/server protocol and read the additional options from the `odbc` section in the `my.cnf` file.

**17.2.3.48. `mysql_ping()`**

```
int mysql_ping(MYSQL *mysql)
```

**Description**

Checks whether the connection to the server is working. If the connection has gone down, an automatic reconnection is attempted.

This function can be used by clients that remain idle for a long while, to check whether the server has closed the connection and reconnect if necessary.

**Return Values**

Zero if the connection to the server is alive. Non-zero if an error occurred. A non-zero return does not indicate whether the MySQL server itself is down; the connection might be broken for other reasons such as network problems.

**Errors**

- `CR_COMMANDS_OUT_OF_SYNC`  
Commands were executed in an improper order.
- `CR_SERVER_GONE_ERROR`  
The MySQL server has gone away.
- `CR_UNKNOWN_ERROR`  
An unknown error occurred.

**17.2.3.49. `mysql_query()`**

```
int mysql_query(MYSQL *mysql, const char *query)
```

**Description**

Executes the SQL query pointed to by the null-terminated string `query`. Normally, the string must con-

---

sist of a single SQL statement and you should not add a terminating semicolon (;) or \g to the statement. If multiple-statement execution has been enabled, the string can contain several statements separated by semicolons. See [Section 17.2.9, “C API Handling of Multiple Statement Execution”](#).

`mysql_query()` cannot be used for queries that contain binary data; you should use `mysql_real_query()` instead. (Binary data may contain the '\0' character, which `mysql_query()` interprets as the end of the query string.)

If you want to know whether the query should return a result set, you can use `mysql_field_count()` to check for this. See [Section 17.2.3.22, “mysql\\_field\\_count\(\)”](#).

### Return Values

Zero if the query was successful. Non-zero if an error occurred.

### Errors

- `CR_COMMANDS_OUT_OF_SYNC`  
Commands were executed in an improper order.
- `CR_SERVER_GONE_ERROR`  
The MySQL server has gone away.
- `CR_SERVER_LOST`  
The connection to the server was lost during the query.
- `CR_UNKNOWN_ERROR`  
An unknown error occurred.

## 17.2.3.50. `mysql_real_connect()`

```
MYSQL *mysql_real_connect(MYSQL *mysql, const char *host, const char *user, const char *passwd, const char *db, unsigned int port, const char *unix_socket, unsigned long client_flag)
```

### Description

`mysql_real_connect()` attempts to establish a connection to a MySQL database engine running on `host`. `mysql_real_connect()` must complete successfully before you can execute any other API functions that require a valid `MYSQL` connection handle structure.

The parameters are specified as follows:

- The first parameter should be the address of an existing `MYSQL` structure. Before calling `mysql_real_connect()` you must call `mysql_init()` to initialize the `MYSQL` structure. You can change a lot of connect options with the `mysql_options()` call. See [Section 17.2.3.47, “mysql\\_options\(\)”](#).
- The value of `host` may be either a hostname or an IP address. If `host` is `NULL` or the string "localhost", a connection to the local host is assumed. If the OS supports sockets (Unix) or named pipes (Windows), they are used instead of TCP/IP to connect to the server.
- The `user` parameter contains the user's MySQL login ID. If `user` is `NULL` or the empty string "",

the current user is assumed. Under Unix, this is the current login name. Under Windows ODBC, the current username must be specified explicitly. See the MyODBC section of [Chapter 18, Connectors](#).

- The `passwd` parameter contains the password for `user`. If `passwd` is `NULL`, only entries in the `user` table for the user that have a blank (empty) password field are checked for a match. This allows the database administrator to set up the MySQL privilege system in such a way that users get different privileges depending on whether they have specified a password.

**Note:** Do not attempt to encrypt the password before calling `mysql_real_connect()`; password encryption is handled automatically by the client API.

- `db` is the database name. If `db` is not `NULL`, the connection sets the default database to this value.
- If `port` is not 0, the value is used as the port number for the TCP/IP connection. Note that the `host` parameter determines the type of the connection.
- If `unix_socket` is not `NULL`, the string specifies the socket or named pipe that should be used. Note that the `host` parameter determines the type of the connection.
- The value of `client_flag` is usually 0, but can be set to a combination of the following flags to enable certain features:

Flag Name	Flag Description
<code>CLIENT_COMPRESS</code>	Use compression protocol.
<code>CLIENT_FOUND_ROWS</code>	Return the number of found (matched) rows, not the number of changed rows.
<code>CLIENT_IGNORE_SPACE</code>	Allow spaces after function names. Makes all functions names reserved words.
<code>CLIENT_INTERACTIVE</code>	Allow <code>interactive_timeout</code> seconds (instead of <code>wait_timeout</code> seconds) of inactivity before closing the connection. The client's session <code>wait_timeout</code> variable is set to the value of the session <code>interactive_timeout</code> variable.
<code>CLIENT_LOCAL_FILES</code>	Enable <code>LOAD DATA LOCAL</code> handling.
<code>CLIENT_MULTI_STATEMENTS</code>	Tell the server that the client may send multiple statements in a single string (separated by <code>;</code> ). If this flag is not set, multiple-statement execution is disabled. See the note following this table for more information about this flag. Added in MySQL 4.1.
<code>CLIENT_MULTI_RESULTS</code>	Tell the server that the client can handle multiple result sets from multiple-statement executions or stored procedures. This is automatically set if <code>CLIENT_MULTI_STATEMENTS</code> is set. See the note following this table for more information about this flag. Added in MySQL 4.1.
<code>CLIENT_NO_SCHEMA</code>	Don't allow the <code>db_name.tbl_name.col_name</code> syntax. This is for ODBC. It causes the parser to generate an error if you use that syntax, which is useful for trapping bugs in some ODBC programs.
<code>CLIENT_ODBC</code>	The client is an ODBC client. This changes <code>mysqld</code> to be more ODBC-friendly.
<code>CLIENT_SSL</code>	Use SSL (encrypted protocol). This option should not be set by application programs; it is set internally in the client library. Instead, use <code>mysql_ssl_set()</code> before calling <code>mysql_real_connect()</code> .

If you enable `CLIENT_MULTI_STATEMENTS` or `CLIENT_MULTI_RESULTS`, you should process the result for every call to `mysql_query()` or `mysql_real_query()` by using a loop that calls



`mysql_next_result()` to determine whether there are more results. For an example, see [Section 17.2.9, “C API Handling of Multiple Statement Execution”](#).

For some parameters, it is possible to have the value taken from an option file rather than from an explicit value in the `mysql_real_connect()` call. To do this, call `mysql_options()` with the `MYSQL_READ_DEFAULT_FILE` or `MYSQL_READ_DEFAULT_GROUP` option before calling `mysql_real_connect()`. Then, in the `mysql_real_connect()` call, specify the “no-value” value for each parameter to be read from an option file:

- For `host`, specify a value of `NULL` or the empty string (`" "`).
- For `user`, specify a value of `NULL` or the empty string.
- For `passwd`, specify a value of `NULL`. (For the password, a value of the empty string in the `mysql_real_connect()` call cannot be overridden in an option file, because the empty string indicates explicitly that the MySQL account must have an empty password.)
- For `db`, specify a value of `NULL` or the empty string.
- For `port`, specify a value of 0.
- For `unix_socket`, specify a value of `NULL`.

If no value is found in an option file for a parameter, its default value is used as indicated in the descriptions given earlier in this section.

### Return Values

A `MYSQL*` connection handle if the connection was successful, `NULL` if the connection was unsuccessful. For a successful connection, the return value is the same as the value of the first parameter.

### Errors

- `CR_CONN_HOST_ERROR`  
Failed to connect to the MySQL server.
- `CR_CONNECTION_ERROR`  
Failed to connect to the local MySQL server.
- `CR_IPSOCK_ERROR`  
Failed to create an IP socket.
- `CR_OUT_OF_MEMORY`  
Out of memory.
- `CR_SOCKET_CREATE_ERROR`  
Failed to create a Unix socket.
- `CR_UNKNOWN_HOST`  
Failed to find the IP address for the hostname.
- `CR_VERSION_ERROR`

A protocol mismatch resulted from attempting to connect to a server with a client library that uses a different protocol version. This can happen if you use a very old client library to connect to a new server that wasn't started with the `--old-protocol` option.

- `CR_NAMEDPIPEOPEN_ERROR`

Failed to create a named pipe on Windows.

- `CR_NAMEDPIPEWAIT_ERROR`

Failed to wait for a named pipe on Windows.

- `CR_NAMEDPIPESETSTATE_ERROR`

Failed to get a pipe handler on Windows.

- `CR_SERVER_LOST`

If `connect_timeout > 0` and it took longer than `connect_timeout` seconds to connect to the server or if the server died while executing the `init-command`.

### Example

```
MYSQL mysql;

mysql_init(&mysql);
mysql_options(&mysql,MYSQL_READ_DEFAULT_GROUP,"your_prog_name");
if (!mysql_real_connect(&mysql,"host","user","passwd","database",0,NULL,0))
{
    fprintf(stderr,"Failed to connect to database: Error: %s\n",
            mysql_error(&mysql));
}
```

By using `mysql_options()` the MySQL library reads the `[client]` and `[your_prog_name]` sections in the `my.cnf` file which ensures that your program works, even if someone has set up MySQL in some non-standard way.

Note that upon connection, `mysql_real_connect()` sets the `reconnect` flag (part of the `MYSQL` structure) to a value of `1`. A value of `1` for this flag indicates that if a statement cannot be performed because of a lost connection, to try reconnecting to the server before giving up.

## 17.2.3.51. `mysql_real_escape_string()`

```
unsigned long mysql_real_escape_string(MYSQL *mysql, char *to, const
char *from, unsigned long length)
```

Note that `mysql` must be a valid, open connection. This is needed because the escaping depends on the character set in use by the server.

### Description

This function is used to create a legal SQL string that you can use in an SQL statement. See [Section 9.1.1, “Strings”](#).

The string in `from` is encoded to an escaped SQL string, taking into account the current character set of the connection. The result is placed in `to` and a terminating null byte is appended. Characters encoded are `NUL` (ASCII 0), `'\n'`, `'\r'`, `'\'`, `'\''`, `'\"'`, and Control-Z (see [Section 9.1, “Literal Values”](#)). (Strictly speaking, MySQL requires only that backslash and the quote character used to quote the string in the query be escaped. This function quotes the other characters to make them easier to read in log files.)

The string pointed to by `from` must be `length` bytes long. You must allocate the `to` buffer to be at least `length*2+1` bytes long. (In the worst case, each character may need to be encoded as using two bytes, and you need room for the terminating null byte.) When `mysql_real_escape_string()` returns, the contents of `to` is a null-terminated string. The return value is the length of the encoded string, not including the terminating null character.

If you need to change the character set of the connection, you should use the `mysql_set_character_set()` function rather than executing a `SET NAMES` (or `SET CHARACTER SET`) statement. `mysql_set_character_set()` works like `SET NAMES` but also affects the character set used by `mysql_real_escape_string()`, which `SET NAMES` does not.

### Example

```
char query[1000],*end;

end = strmov(query,"INSERT INTO test_table values(");
*end++ = '\\';
end += mysql_real_escape_string(&mysql, end,"What's this",11);
*end++ = '\\';
*end++ = ',';
*end++ = '\\';
end += mysql_real_escape_string(&mysql, end,"binary data: \\0\\r\\n",16);
*end++ = '\\';
*end++ = ')';

if (mysql_real_query(&mysql,query,(unsigned int) (end - query)))
{
    fprintf(stderr, "Failed to insert row, Error: %s\\n",
            mysql_error(&mysql));
}
```

The `strmov()` function used in the example is included in the `mysqlclient` library and works like `strcpy()` but returns a pointer to the terminating null of the first parameter.

### Return Values

The length of the value placed into `to`, not including the terminating null character.

### Errors

None.

## 17.2.3.52. `mysql_real_query()`

```
int mysql_real_query(MYSQL *mysql, const char *query, unsigned long
length)
```

### Description

Executes the SQL query pointed to by `query`, which should be a string `length` bytes long. Normally, the string must consist of a single SQL statement and you should not add a terminating semicolon (`;`) or `\g` to the statement. If multiple-statement execution has been enabled, the string can contain several statements separated by semicolons. See [Section 17.2.9, “C API Handling of Multiple Statement Execution”](#).

You **must** use `mysql_real_query()` rather than `mysql_query()` for queries that contain binary data, because binary data may contain the `\0` character. In addition, `mysql_real_query()` is faster than `mysql_query()` because it does not call `strlen()` on the query string.

If you want to know whether the query should return a result set, you can use `mysql_field_count()` to check for this. See [Section 17.2.3.22, “mysql\\_field\\_count\(\)”](#).

**Return Values**

Zero if the query was successful. Non-zero if an error occurred.

**Errors**

- `CR_COMMANDS_OUT_OF_SYNC`  
Commands were executed in an improper order.
- `CR_SERVER_GONE_ERROR`  
The MySQL server has gone away.
- `CR_SERVER_LOST`  
The connection to the server was lost during the query.
- `CR_UNKNOWN_ERROR`  
An unknown error occurred.

**17.2.3.53. `mysql_refresh()`**

```
int mysql_refresh(MYSQL *mysql, unsigned int options)
```

**Description**

This function flushes tables or caches, or resets replication server information. The connected user must have the `RELOAD` privilege.

The `options` argument is a bit mask composed from any combination of the following values. Multiple values can be OR'ed together to perform multiple operations with a single call.

- `REFRESH_GRANT`  
Refresh the grant tables, like `FLUSH PRIVILEGES`.
- `REFRESH_LOG`  
Flush the logs, like `FLUSH LOGS`.
- `REFRESH_TABLES`  
Flush the table cache, like `FLUSH TABLES`.
- `REFRESH_HOSTS`  
Flush the host cache, like `FLUSH HOSTS`.
- `REFRESH_STATUS`  
Reset status variables, like `FLUSH STATUS`.
- `REFRESH_THREADS`  
Flush the thread cache.

- `REFRESH_SLAVE`

On a slave replication server, reset the master server information and restart the slave, like `RESET SLAVE`.

- `REFRESH_MASTER`

On a master replication server, remove the binary log files listed in the binary log index and truncate the index file, like `RESET MASTER`.

### Return Values

Zero for success. Non-zero if an error occurred.

### Errors

- `CR_COMMANDS_OUT_OF_SYNC`

Commands were executed in an improper order.

- `CR_SERVER_GONE_ERROR`

The MySQL server has gone away.

- `CR_SERVER_LOST`

The connection to the server was lost during the query.

- `CR_UNKNOWN_ERROR`

An unknown error occurred.

## 17.2.3.54. `mysql_reload()`

```
int mysql_reload(MYSQL *mysql)
```

### Description

Asks the MySQL server to reload the grant tables. The connected user must have the `RELOAD` privilege.

This function is deprecated. It is preferable to use `mysql_query()` to issue an SQL `FLUSH PRIVILEGES` statement instead.

### Return Values

Zero for success. Non-zero if an error occurred.

### Errors

- `CR_COMMANDS_OUT_OF_SYNC`

Commands were executed in an improper order.

- `CR_SERVER_GONE_ERROR`

The MySQL server has gone away.

- `CR_SERVER_LOST`

The connection to the server was lost during the query.

- `CR_UNKNOWN_ERROR`

An unknown error occurred.

### 17.2.3.55. `mysql_rollback()`

```
my_bool mysql_rollback(MYSQL *mysql)
```

#### **Description**

Rolls back the current transaction.

This function was added in MySQL 4.1.0.

#### **Return Values**

Zero if successful. Non-zero if an error occurred.

#### **Errors**

None.

### 17.2.3.56. `mysql_row_seek()`

```
MYSQL_ROW_OFFSET mysql_row_seek(MYSQL_RES *result, MYSQL_ROW_OFFSET  
offset)
```

#### **Description**

Sets the row cursor to an arbitrary row in a query result set. The `offset` value is a row offset that should be a value returned from `mysql_row_tell()` or from `mysql_row_seek()`. This value is not a row number; if you want to seek to a row within a result set by number, use `mysql_data_seek()` instead.

This function requires that the result set structure contains the entire result of the query, so `mysql_row_seek()` may be used only in conjunction with `mysql_store_result()`, not with `mysql_use_result()`.

#### **Return Values**

The previous value of the row cursor. This value may be passed to a subsequent call to `mysql_row_seek()`.

#### **Errors**

None.

### 17.2.3.57. `mysql_row_tell()`

```
MYSQL_ROW_OFFSET mysql_row_tell(MYSQL_RES *result)
```

#### **Description**

Returns the current position of the row cursor for the last `mysql_fetch_row()`. This value can be used as an argument to `mysql_row_seek()`.

You should use `mysql_row_tell()` only after `mysql_store_result()`, not after `mysql_use_result()`.

**Return Values**

The current offset of the row cursor.

**Errors**

None.

### 17.2.3.58. `mysql_select_db()`

```
int mysql_select_db(MYSQL *mysql, const char *db)
```

**Description**

Causes the database specified by `db` to become the default (current) database on the connection specified by `mysql`. In subsequent queries, this database is the default for table references that do not include an explicit database specifier.

`mysql_select_db()` fails unless the connected user can be authenticated as having permission to use the database.

**Return Values**

Zero for success. Non-zero if an error occurred.

**Errors**

- `CR_COMMANDS_OUT_OF_SYNC`  
Commands were executed in an improper order.
- `CR_SERVER_GONE_ERROR`  
The MySQL server has gone away.
- `CR_SERVER_LOST`  
The connection to the server was lost during the query.
- `CR_UNKNOWN_ERROR`  
An unknown error occurred.

### 17.2.3.59. `mysql_set_character_set()`

```
int mysql_set_character_set(MYSQL *mysql, char *csname)
```

**Description**

This function is used to set the default character set for the current connection. The string `csname` specifies a valid character set name. The connection collation becomes the default collation of the character

set. This function works like the `SET NAMES` statement, but also sets the value of `mysql->charset`, and thus affects the character set used by `mysql_real_escape_string()`

This function was added in MySQL 4.1.13.

### Return Values

Zero for success. Non-zero if an error occurred.

### Example

```
MYSQL mysql;

mysql_init(&mysql);
if (!mysql_real_connect(&mysql, "host", "user", "passwd", "database", 0, NULL, 0))
{
    fprintf(stderr, "Failed to connect to database: Error: %s\n",
            mysql_error(&mysql));
}

if (!mysql_set_character_set(&mysql, "utf8"))
{
    printf("New client character set: %s\n", mysql_character_set_name(&mysql));
}
```

## 17.2.3.60. `mysql_set_local_infile_default()`

```
void
mysql_set_local_infile_default(MYSQL *mysql);
```

### Description

Sets the `LOAD LOCAL DATA INFILE` handler callback functions to the defaults used internally by the C client library. The library calls this function automatically if `mysql_set_local_infile_handler()` has not been called or does not supply valid functions for each of its callbacks.

The `mysql_set_local_infile_default()` function was added in MySQL 4.1.2.

### Return Values

None.

### Errors

None.

## 17.2.3.61. `mysql_set_local_infile_handler()`

```
void
mysql_set_local_infile_handler(MYSQL *mysql,
                               int (*local_infile_init)(void **, const char *, void *),
                               int (*local_infile_read)(void *, char *, unsigned int),
                               void (*local_infile_end)(void *),
                               int (*local_infile_error)(void *, char*, unsigned int),
                               void *userdata);
```

### Description

This function installs callbacks to be used during the execution of `LOAD DATA LOCAL INFILE` statements. It enables application programs to exert control over local (client-side) datafile reading. The arguments are the connection handler, a set of pointers to callback functions, and a pointer to a data area that the callbacks can use to share information.



To use `mysql_set_local_infile_handler()`, you must write the following callback functions:

```
int
local_infile_init(void **ptr, const char *filename, void *userdata);
```

The initialization function. This is called once to do any setup necessary, open the datafile, allocate data structures, and so forth. The first `void**` argument is a pointer to a pointer. You can set the pointer (that is, `*ptr`) to a value that will be passed to each of the other callbacks (as a `void*`). The callbacks can use this pointed-to value to maintain state information. The `userdata` argument is the same value that is passed to `mysql_set_local_infile_handler()`.

The initialization function should return zero for success, non-zero for an error.

```
int
local_infile_read(void *ptr, char *buf, unsigned int buf_len);
```

The data-reading function. This is called repeatedly to read the data file. `buf` points to the buffer where the read data should be stored, and `buf_len` is the maximum number of bytes that the callback can read and store in the buffer. (It can read fewer bytes, but should not read more.)

The return value is the number of bytes read, or zero when no more data could be read (this indicates EOF). Return a value less than zero if an error occurs.

```
void
local_infile_end(void *ptr)
```

The termination function. This is called once after `local_infile_read()` has returned zero (EOF) or an error. This function should deallocate any memory allocated by `local_infile_init()` and perform any other cleanup necessary. It is invoked even if the initialization function returns an error.

```
int
local_infile_error(void *ptr, char *error_msg, unsigned int error_msg_len);
```

The error-handling function. This is called to get a textual error message to return to the user in case any of your other functions returns an error. `error_msg` points to the buffer into which the message should be written, and `error_msg_len` is the length of the buffer. The message should be written as a null-terminated string, so the message can be at most `error_msg_len-1` bytes long.

The return value is the error number.

Typically, the other callbacks store the error message in the data structure pointed to by `ptr`, so that `local_infile_error()` can copy the message from there into `error_msg`.

After calling `mysql_set_local_infile_handler()` in your C code and passing pointers to your callback functions, you can then issue a `LOAD DATA LOCAL INFILE` statement (for example, by using `mysql_query()`). The client library automatically invokes your callbacks. The filename specified in `LOAD DATA LOCAL INFILE` will be passed as the second parameter to the `local_infile_init()` callback.

The `mysql_set_local_infile_handler()` function was added in MySQL 4.1.2.

## Return Values

None.

## Errors

None.

### 17.2.3.62. `mysql_set_server_option()`

```
int mysql_set_server_option(MYSQL *mysql, enum enum_mysql_set_option option)
```

#### Description

Enables or disables an option for the connection. `option` can have one of the following values:

<code>MYSQL_OPTION_MULTI_STATEMENTS_ON</code>	Enable multiple-statement support.
<code>MYSQL_OPTION_MULTI_STATEMENTS_OFF</code>	Disable multiple-statement support.

If you enable multiple-statement support, you should retrieve results from calls to `mysql_query()` or `mysql_real_query()` by using a loop that calls `mysql_next_result()` to determine whether there are more results. For an example, see [Section 17.2.9, “C API Handling of Multiple Statement Execution”](#).

Enabling multiple-statement support with `MYSQL_OPTION_MULTI_STATEMENTS_ON` does not have quite the same effect as enabling it by passing the `CLIENT_MULTI_STATEMENTS` flag to `mysql_real_connect()`: `CLIENT_MULTI_STATEMENTS` also enables `CLIENT_MULTI_RESULTS`.

This function was added in MySQL 4.1.1.

#### Return Values

Zero for success. Non-zero if an error occurred.

#### Errors

- `CR_COMMANDS_OUT_OF_SYNC`  
Commands were executed in an improper order.
- `CR_SERVER_GONE_ERROR`  
The MySQL server has gone away.
- `CR_SERVER_LOST`  
The connection to the server was lost during the query.
- `ER_UNKNOWN_COM_ERROR`  
The server didn't support `mysql_set_server_option()` (which is the case that the server is older than 4.1.1) or the server didn't support the option one tried to set.

### 17.2.3.63. `mysql_shutdown()`

```
int mysql_shutdown(MYSQL *mysql, enum mysql_enum_shutdown_level shutdown_level)
```

## Description

Asks the database server to shut down. The connected user must have `SHUTDOWN` privileges. The `shutdown_level` argument was added in MySQL 4.1.3. MySQL 4.1 supports only one type of shutdown; `shutdown_level` must be equal to `SHUTDOWN_DEFAULT`. Additional shutdown levels are planned to make it possible to choose the desired level. Dynamically linked executables which have been compiled with older versions of the `libmysqlclient` headers and call `mysql_shutdown()` need to be used with the old `libmysqlclient` dynamic library.

The shutdown process is described in [Section 5.2.6, “The Shutdown Process”](#).

## Return Values

Zero for success. Non-zero if an error occurred.

## Errors

- `CR_COMMANDS_OUT_OF_SYNC`  
Commands were executed in an improper order.
- `CR_SERVER_GONE_ERROR`  
The MySQL server has gone away.
- `CR_SERVER_LOST`  
The connection to the server was lost during the query.
- `CR_UNKNOWN_ERROR`  
An unknown error occurred.

### 17.2.3.64. `mysql_sqlstate()`

```
const char *mysql_sqlstate(MYSQL *mysql)
```

#### Description

Returns a null-terminated string containing the SQLSTATE error code for the last error. The error code consists of five characters. `'00000'` means “no error.” The values are specified by ANSI SQL and ODBC. For a list of possible values, see [Appendix B, \*Error Codes and Messages\*](#).

SQLSTATE values returned by `mysql_sqlstate()` differ from MySQL-specific error numbers returned by `mysql_errno()`. For example, the `mysql` client program displays errors using the following format, where `1146` is the `mysql_errno()` value and `'42S02'` is the corresponding `mysql_sqlstate()` value:

```
shell> SELECT * FROM no_such_table;  
ERROR 1146 (42S02): Table 'test.no_such_table' doesn't exist
```

Not all MySQL error numbers are mapped to SQLSTATE error codes. The value `'HY000'` (general error) is used for unmapped error numbers.

This function was added in MySQL 4.1.1.

#### Return Values

A null-terminated character string containing the SQLSTATE error code.

**See Also**

See [Section 17.2.3.14](#), “`mysql_errno()`”, [Section 17.2.3.15](#), “`mysql_error()`”, and [Section 17.2.7.26](#), “`mysql_stmt_sqlstate()`”.

### 17.2.3.65. `mysql_ssl_set()`

```
int mysql_ssl_set(MYSQL *mysql, const char *key, const char *cert,
const char *ca, const char *capath, const char *cipher)
```

**Description**

`mysql_ssl_set()` is used for establishing secure connections using SSL. It must be called before `mysql_real_connect()`.

`mysql_ssl_set()` does nothing unless OpenSSL support is enabled in the client library.

`mysql` is the connection handler returned from `mysql_init()`. The other parameters are specified as follows:

- `key` is the pathname to the key file.
- `cert` is the pathname to the certificate file.
- `ca` is the pathname to the certificate authority file.
- `capath` is the pathname to a directory that contains trusted SSL CA certificates in pem format.
- `cipher` is a list of allowable ciphers to use for SSL encryption.

Any unused SSL parameters may be given as `NULL`.

**Return Values**

This function always returns 0. If SSL setup is incorrect, `mysql_real_connect()` returns an error when you attempt to connect.

### 17.2.3.66. `mysql_stat()`

```
const char *mysql_stat(MYSQL *mysql)
```

**Description**

Returns a character string containing information similar to that provided by the `mysqladmin status` command. This includes uptime in seconds and the number of running threads, questions, reloads, and open tables.

**Return Values**

A character string describing the server status. `NULL` if an error occurred.

**Errors**

- `CR_COMMANDS_OUT_OF_SYNC`

Commands were executed in an improper order.

- `CR_SERVER_GONE_ERROR`

The MySQL server has gone away.

- `CR_SERVER_LOST`

The connection to the server was lost during the query.

- `CR_UNKNOWN_ERROR`

An unknown error occurred.

### 17.2.3.67. `mysql_store_result()`

```
MYSQL_RES *mysql_store_result(MYSQL *mysql)
```

#### Description

You must call `mysql_store_result()` or `mysql_use_result()` for every query that successfully retrieves data (`SELECT`, `SHOW`, `DESCRIBE`, `EXPLAIN`, `CHECK TABLE`, and so forth).

You don't have to call `mysql_store_result()` or `mysql_use_result()` for other queries, but it does not do any harm or cause any notable performance degradation if you call `mysql_store_result()` in all cases. You can detect if the query didn't have a result set by checking if `mysql_store_result()` returns 0 (more about this later on).

If you want to know whether the query should return a result set, you can use `mysql_field_count()` to check for this. See [Section 17.2.3.22, “mysql\\_field\\_count\(\)”](#).

`mysql_store_result()` reads the entire result of a query to the client, allocates a `MYSQL_RES` structure, and places the result into this structure.

`mysql_store_result()` returns a null pointer if the query didn't return a result set (if the query was, for example, an `INSERT` statement).

`mysql_store_result()` also returns a null pointer if reading of the result set failed. You can check whether an error occurred by checking if `mysql_error()` returns a non-empty string, if `mysql_errno()` returns non-zero, or if `mysql_field_count()` returns zero.

An empty result set is returned if there are no rows returned. (An empty result set differs from a null pointer as a return value.)

Once you have called `mysql_store_result()` and got a result back that isn't a null pointer, you may call `mysql_num_rows()` to find out how many rows are in the result set.

You can call `mysql_fetch_row()` to fetch rows from the result set, or `mysql_row_seek()` and `mysql_row_tell()` to obtain or set the current row position within the result set.

You must call `mysql_free_result()` once you are done with the result set.

See [Section 17.2.13.1, “Why mysql\\_store\\_result\(\) Sometimes Returns NULL After mysql\\_query\(\) Returns Success”](#).

#### Return Values

A `MYSQL_RES` result structure with the results. `NULL` if an error occurred.

### Errors

`mysql_store_result()` resets `mysql_error()` and `mysql_errno()` if it succeeds.

- `CR_COMMANDS_OUT_OF_SYNC`  
Commands were executed in an improper order.
- `CR_OUT_OF_MEMORY`  
Out of memory.
- `CR_SERVER_GONE_ERROR`  
The MySQL server has gone away.
- `CR_SERVER_LOST`  
The connection to the server was lost during the query.
- `CR_UNKNOWN_ERROR`  
An unknown error occurred.

### 17.2.3.68. `mysql_thread_id()`

```
unsigned long mysql_thread_id(MYSQL *mysql)
```

#### Description

Returns the thread ID of the current connection. This value can be used as an argument to `mysql_kill()` to kill the thread.

If the connection is lost and you reconnect with `mysql_ping()`, the thread ID changes. This means you should not get the thread ID and store it for later. You should get it when you need it.

#### Return Values

The thread ID of the current connection.

#### Errors

None.

### 17.2.3.69. `mysql_use_result()`

```
MYSQL_RES *mysql_use_result(MYSQL *mysql)
```

#### Description

You must call `mysql_store_result()` or `mysql_use_result()` for every query that successfully retrieves data (`SELECT`, `SHOW`, `DESCRIBE`, `EXPLAIN`).

`mysql_use_result()` initiates a result set retrieval but does not actually read the result set into the client like `mysql_store_result()` does. Instead, each row must be retrieved individually by mak-

ing calls to `mysql_fetch_row()`. This reads the result of a query directly from the server without storing it in a temporary table or local buffer, which is somewhat faster and uses much less memory than `mysql_store_result()`. The client allocates memory only for the current row and a communication buffer that may grow up to `max_allowed_packet` bytes.

On the other hand, you shouldn't use `mysql_use_result()` if you are doing a lot of processing for each row on the client side, or if the output is sent to a screen on which the user may type a ^S (stop scroll). This ties up the server and prevent other threads from updating any tables from which the data is being fetched.

When using `mysql_use_result()`, you must execute `mysql_fetch_row()` until a `NULL` value is returned, otherwise, the unfetched rows are returned as part of the result set for your next query. The C API gives the error `Commands out of sync; you can't run this command now` if you forget to do this!

You may not use `mysql_data_seek()`, `mysql_row_seek()`, `mysql_row_tell()`, `mysql_num_rows()`, or `mysql_affected_rows()` with a result returned from `mysql_use_result()`, nor may you issue other queries until `mysql_use_result()` has finished. (However, after you have fetched all the rows, `mysql_num_rows()` accurately returns the number of rows fetched.)

You must call `mysql_free_result()` once you are done with the result set.

When using the `libmysqld` embedded server, the memory benefits are essentially lost because memory usage incrementally increases with each row retrieved until `mysql_free_result()` is called.

### Return Values

A `MYSQL_RES` result structure. `NULL` if an error occurred.

### Errors

`mysql_use_result()` resets `mysql_error()` and `mysql_errno()` if it succeeds.

- `CR_COMMANDS_OUT_OF_SYNC`  
Commands were executed in an improper order.
- `CR_OUT_OF_MEMORY`  
Out of memory.
- `CR_SERVER_GONE_ERROR`  
The MySQL server has gone away.
- `CR_SERVER_LOST`  
The connection to the server was lost during the query.
- `CR_UNKNOWN_ERROR`  
An unknown error occurred.

## 17.2.3.70. `mysql_warning_count()`

```
unsigned int mysql_warning_count(MYSQL *mysql)
```

**Description**

Returns the number of warnings generated during execution of the previous SQL statement.

This function was added in MySQL 4.1.0.

**Return Values**

The warning count.

**Errors**

None.

## 17.2.4. C API Prepared Statements

As of MySQL 4.1, the client/server protocol provides for the use of prepared statements. This capability uses the `MYSQL_STMT` statement handler data structure returned by the `mysql_stmt_init()` initialization function. Prepared execution is an efficient way to execute a statement more than once. The statement is first parsed to prepare it for execution. Then it is executed one or more times at a later time, using the statement handle returned by the initialization function.

Prepared execution is faster than direct execution for statements executed more than once, primarily because the query is parsed only once. In the case of direct execution, the query is parsed every time it is executed. Prepared execution also can provide a reduction of network traffic because for each execution of the prepared statement, it is necessary only to send the data for the parameters.

Prepared statements might not provide a performance increase in some situations. For best results, test your application both with prepared and non-prepared statements and choose whichever yields best performance.

Another advantage of prepared statements is that it uses a binary protocol that makes data transfer between client and server more efficient.

The following statements can be used as prepared statements: `CREATE TABLE`, `DELETE`, `DO`, `INSERT`, `REPLACE`, `SELECT`, `SET`, `UPDATE`, and most `SHOW` statements. Other statements are not yet supported.

## 17.2.5. C API Prepared Statement Data types

**Note:** Some incompatible changes were made in MySQL 4.1.2. See [Section 17.2.7, “C API Prepared Statement Function Descriptions”](#), for details.

Prepared statements mainly use the `MYSQL_STMT` and `MYSQL_BIND` data structures. A third structure, `MYSQL_TIME`, is used to transfer temporal data.

- `MYSQL_STMT`

This structure represents a prepared statement. A statement is created by calling `mysql_stmt_init()`, which returns a statement handle (that is, a pointer to a `MYSQL_STMT`). The handle is used for all subsequent statement-related functions until you close it with `mysql_stmt_close()`.

The `MYSQL_STMT` structure has no members that are for application use. Also, you should not try to make a copy of a `MYSQL_STMT` structure. There is no guarantee that such a copy will be usable.

Multiple statement handles can be associated with a single connection. The limit on the number of



handles depends on the available system resources.

- `MYSQL_BIND`

This structure is used both for statement input (data values sent to the server) and output (result values returned from the server). For input, it is used with `mysql_stmt_bind_param()` to bind parameter data values to buffers for use by `mysql_stmt_execute()`. For output, it is used with `mysql_stmt_bind_result()` to bind result set buffers for use in fetching rows with `mysql_stmt_fetch()`.

To use a `MYSQL_BIND` structure, you should zero its contents to initialize it, and then set its members appropriately. For example, to declare and initialize an array of three `MYSQL_BIND` structures, use this code:

```
MYSQL_BIND  bind[3];
memset(bind, 0, sizeof(bind));
```

The `MYSQL_BIND` structure contains the following members for use by application programs. Each is used both for input and for output, although sometimes for different purposes depending on the direction of data transfer.

- `enum enum_field_types buffer_type`

The type of the buffer. The allowable `buffer_type` values are listed later in this section. For input, `buffer_type` indicates what type of value you are binding to a statement parameter. For output, it indicates what type of value you expect to receive in a result buffer.

- `void *buffer`

For input, this is a pointer to the buffer in which a statement parameter's data value is stored. For output, it is a pointer to the buffer in which to return a result set column value. For numeric data types, `buffer` should point to a variable of the proper C type. (If you are associating the variable with a column that has the `UNSIGNED` attribute, the variable should be an `unsigned` C type. Indicate whether the variable is signed or unsigned by using the `is_unsigned` member, described later in this list.) For date and time data types, `buffer` should point to a `MYSQL_TIME` structure. For character and binary string data types, `buffer` should point to a character buffer.

- `unsigned long buffer_length`

The actual size of `*buffer` in bytes. This indicates the maximum amount of data that can be stored in the buffer. For character and binary C data, the `buffer_length` value specifies the length of `*buffer` when used with `mysql_stmt_bind_param()`, or the maximum number of data bytes that can be fetched into the buffer when used with `mysql_stmt_bind_result()`.

- `unsigned long *length`

A pointer to an `unsigned long` variable that indicates the actual number of bytes of data stored in `*buffer`. `length` is used for character or binary C data. For input parameter data binding, `length` points to an `unsigned long` variable that indicates the length of the parameter value stored in `*buffer`; this is used by `mysql_stmt_execute()`. For output value binding, `mysql_stmt_fetch()` places the length of the column value that is returned into the variable that `length` points to.

`length` is ignored for numeric and temporal data types because the length of the data value is determined by the `buffer_type` value.

- `my_bool *is_null`

This member points to a `my_bool` variable that is true if a value is `NULL`, false if it is not `NULL`. For input, set `*is_null` to true to indicate that you are passing a `NULL` value as a statement parameter. For output, this value is set to true after you fetch a row if the result set column value returned from the statement is `NULL`.

`is_null` is a pointer to a boolean rather than a boolean scalar so that it can be used in the following way:

- If your data values are always `NULL`, use `MYSQL_TYPE_NULL` to bind the column.
- If your data values are always `NOT NULL`, set `is_null = (my_bool*) 0`.
- In all other cases, you should set `is_null` to the address of a `my_bool` variable and change that variable's value appropriately between executions to indicate whether data values are `NULL` or `NOT NULL`.

- `my_bool is_unsigned`

This member is used for integer types. (These correspond to the `MYSQL_TYPE_TINY`, `MYSQL_TYPE_SHORT`, `MYSQL_TYPE_LONG`, and `MYSQL_TYPE_LONGLONG` type codes.) `is_unsigned` should be set to true for unsigned types and false for signed types.

- `MYSQL_TIME`

This structure is used to send and receive `DATE`, `TIME`, `DATETIME`, and `TIMESTAMP` data directly to and from the server. This is done by setting the `buffer_type` member of a `MYSQL_BIND` structure to one of the temporal types, and setting the `buffer` member to point to a `MYSQL_TIME` structure.

The `MYSQL_TIME` structure contains the following members:

- `unsigned int year`

The year.

- `unsigned int month`

The month of the year.

- `unsigned int day`

The day of the month.

- `unsigned int hour`

The hour of the day.

- `unsigned int minute`

The minute of the hour.

- `unsigned int second`

The second of the minute.

- `my_bool neg`

A boolean flag to indicate whether the time is negative.

- `unsigned long second_part`

The fractional part of the second. This member currently is unused.

Only those parts of a `MYSQL_TIME` structure that apply to a given type of temporal value are used: The `year`, `month`, and `day` elements are used for `DATE`, `DATETIME`, and `TIMESTAMP` values. The `hour`, `minute`, and `second` elements are used for `TIME`, `DATETIME`, and `TIMESTAMP` values. See [Section 17.2.10, “C API Handling of Date and Time Values”](#).

The following table shows the allowable values that may be specified in the `buffer_type` member of `MYSQL_BIND` structures. The table also shows those SQL types that correspond most closely to each `buffer_type` value, and, for numeric and temporal types, the corresponding recommended C type.

The types are “recommended” because implicit type conversion may be performed in both directions. The `buffer_type` value controls the conversion that will be performed. For example, to fetch a SQL `MEDIUMINT` column value, you can specify a `buffer_type` value of `MYSQL_TYPE_LONG` and use a C variable of type `int` as the destination buffer. If you fetch a numeric column with a value of 255 into a `char[4]` character array, specify a `buffer_type` value of `MYSQL_TYPE_STRING` and the resulting value in the array will be a 4-byte string containing `'255\0'`.

To distinguish between binary and non-binary data for string data types, check whether the `charsetnr` value of the result set metadata is 63. If so, the character set is `binary`, which indicates binary rather than non-binary data. This is how to distinguish between `BINARY` and `CHAR`, `VARBINARY` and `VARCHAR`, and `BLOB` and `TEXT`.

<code>buffer_type</code> Value	SQL Type	Recommended C Type
<code>MYSQL_TYPE_TINY</code>	<code>TINYINT</code>	<code>unsigned char</code>
<code>MYSQL_TYPE_SHORT</code>	<code>SMALLINT</code>	<code>short int</code>
<code>MYSQL_TYPE_LONG</code>	<code>INT</code>	<code>int</code>
<code>MYSQL_TYPE_LONGLONG</code>	<code>BIGINT</code>	<code>long long int</code>
<code>MYSQL_TYPE_FLOAT</code>	<code>FLOAT</code>	<code>float</code>
<code>MYSQL_TYPE_DOUBLE</code>	<code>DOUBLE</code>	<code>double</code>
<code>MYSQL_TYPE_TIME</code>	<code>TIME</code>	<code>MYSQL_TIME</code>
<code>MYSQL_TYPE_DATE</code>	<code>DATE</code>	<code>MYSQL_TIME</code>
<code>MYSQL_TYPE_DATETIME</code>	<code>DATETIME</code>	<code>MYSQL_TIME</code>
<code>MYSQL_TYPE_TIMESTAMP</code>	<code>TIMESTAMP</code>	<code>MYSQL_TIME</code>
<code>MYSQL_TYPE_STRING</code>	<code>CHAR/BINARY</code>	<code>char[]</code>
<code>MYSQL_TYPE_VAR_STRING</code>	<code>VARCHAR/VARBINARY</code>	<code>char[]</code>
<code>MYSQL_TYPE_TINY_BLOB</code>	<code>TINYBLOB/TINYTEXT</code>	<code>char[]</code>
<code>MYSQL_TYPE_BLOB</code>	<code>BLOB/TEXT</code>	<code>char[]</code>
<code>MYSQL_TYPE_MEDIUM_BLOB</code>	<code>MEDIUMBLOB/MEDIUMTEXT</code>	<code>char[]</code>
<code>MYSQL_TYPE_LONG_BLOB</code>	<code>LOB/LOBTEXT</code>	<code>char[]</code>

## 17.2.6. C API Prepared Statement Function Overview

**Note:** Some incompatible changes were made in MySQL 4.1.2. See [Section 17.2.7, “C API Prepared Statement Function Descriptions”](#), for details.

The functions available for prepared statement processing are summarized here and described in greater detail in a later section. See [Section 17.2.7, “C API Prepared Statement Function Descriptions”](#).

Function	Description
<a href="#">mysql_stmt_affected_rows()</a>	Returns the number of rows changes, deleted, or inserted by prepared <a href="#">UPDATE</a> , <a href="#">DELETE</a> , or <a href="#">INSERT</a> statement.
<a href="#">mysql_stmt_attr_get()</a>	Get value of an attribute for a prepared statement.
<a href="#">mysql_stmt_attr_set()</a>	Sets an attribute for a prepared statement.
<a href="#">mysql_stmt_bind_param()</a>	Associates application data buffers with the parameter markers in a prepared SQL statement.
<a href="#">mysql_stmt_bind_result()</a>	Associates application data buffers with columns in the result set.
<a href="#">mysql_stmt_close()</a>	Frees memory used by prepared statement.
<a href="#">mysql_stmt_data_seek()</a>	Seeks to an arbitrary row number in a statement result set.
<a href="#">mysql_stmt_errno()</a>	Returns the error number for the last statement execution.
<a href="#">mysql_stmt_error()</a>	Returns the error message for the last statement execution.
<a href="#">mysql_stmt_execute()</a>	Executes the prepared statement.
<a href="#">mysql_stmt_fetch()</a>	Fetches the next row of data from the result set and returns data for all bound columns.
<a href="#">mysql_stmt_fetch_column()</a>	Fetch data for one column of the current row of the result set.
<a href="#">mysql_stmt_field_count()</a>	Returns the number of result columns for the most recent statement.
<a href="#">mysql_stmt_free_result()</a>	Free the resources allocated to the statement handle.
<a href="#">mysql_stmt_init()</a>	Allocates memory for <a href="#">MYSQL_STMT</a> structure and initializes it.
<a href="#">mysql_stmt_insert_id()</a>	Returns the ID generated for an <a href="#">AUTO_INCREMENT</a> column by prepared statement.
<a href="#">mysql_stmt_num_rows()</a>	Returns total rows from the statement buffered result set.
<a href="#">mysql_stmt_param_count()</a>	Returns the number of parameters in a prepared SQL statement.
<a href="#">mysql_stmt_param_metadata()</a>	(Return parameter metadata in the form of a result set.) Currently, this function does nothing.
<a href="#">mysql_stmt_prepare()</a>	Prepares an SQL string for execution.
<a href="#">mysql_stmt_reset()</a>	Reset the statement buffers in the server.
<a href="#">mysql_stmt_result_metadata()</a>	Returns prepared statement metadata in the form of a result set.
<a href="#">mysql_stmt_row_seek()</a>	Seeks to a row offset in a statement result set, using value returned from <a href="#">mysql_stmt_row_tell()</a> .
<a href="#">mysql_stmt_row_tell()</a>	Returns the statement row cursor position.
<a href="#">mysql_stmt_send_long_data()</a>	Sends long data in chunks to server.
<a href="#">mysql_stmt_sqlstate()</a>	Returns the SQLSTATE error code for the last statement execution.
<a href="#">mysql_stmt_store_result()</a>	Retrieves the complete result set to the client.

Call [mysql\\_stmt\\_init\(\)](#) to create a statement handle, then [mysql\\_stmt\\_prepare](#) to prepare it, [mysql\\_stmt\\_bind\\_param\(\)](#) to supply the parameter data, and [mysql\\_stmt\\_execute\(\)](#) to execute the statement. You can repeat the [mysql\\_stmt\\_execute\(\)](#) by changing parameter values in the respective buffers supplied through [mysql\\_stmt\\_bind\\_param\(\)](#).

If the statement is a `SELECT` or any other statement that produces a result set, `mysql_stmt_prepare()` also returns the result set metadata information in the form of a `MYSQL_RES` result set through `mysql_stmt_result_metadata()`.

You can supply the result buffers using `mysql_stmt_bind_result()`, so that the `mysql_stmt_fetch()` automatically returns data to these buffers. This is row-by-row fetching.

You can also send the text or binary data in chunks to server using `mysql_stmt_send_long_data()`. See [Section 17.2.7.25](#), “`mysql_stmt_send_long_data()`”.

When statement execution has been completed, the statement handle must be closed using `mysql_stmt_close()` so that all resources associated with it can be freed.

If you obtained a `SELECT` statement's result set metadata by calling `mysql_stmt_result_metadata()`, you should also free the metadata using `mysql_free_result()`.

### Execution Steps

To prepare and execute a statement, an application follows these steps:

1. Create a prepared statement handle with `mysql_stmt_init()`. To prepare the statement on the server, call `mysql_stmt_prepare()` and pass it a string containing the SQL statement.
2. If the statement produces a result set, call `mysql_stmt_result_metadata()` to obtain the result set metadata. This metadata is itself in the form of result set, albeit a separate one from the one that contains the rows returned by the query. The metadata result set indicates how many columns are in the result and contains information about each column.
3. Set the values of any parameters using `mysql_stmt_bind_param()`. All parameters must be set. Otherwise, statement execution returns an error or produces unexpected results.
4. Call `mysql_stmt_execute()` to execute the statement.
5. If the statement produces a result set, bind the data buffers to use for retrieving the row values by calling `mysql_stmt_bind_result()`.
6. Fetch the data into the buffers row by row by calling `mysql_stmt_fetch()` repeatedly until no more rows are found.
7. Repeat steps 3 through 6 as necessary, by changing the parameter values and re-executing the statement.

When `mysql_stmt_prepare()` is called, the MySQL client/server protocol performs these actions:

- The server parses the statement and sends the okay status back to the client by assigning a statement ID. It also sends total number of parameters, a column count, and its metadata if it is a result set oriented statement. All syntax and semantics of the statement are checked by the server during this call.
- The client uses this statement ID for the further operations, so that the server can identify the statement from among its pool of statements.

When `mysql_stmt_execute()` is called, the MySQL client/server protocol performs these actions:

- The client uses the statement handle and sends the parameter data to the server.
- The server identifies the statement using the ID provided by the client, replaces the parameter markers with the newly supplied data, and executes the statement. If the statement produces a result set, the server sends the data back to the client. Otherwise, it sends an okay status and total number of rows changed, deleted, or inserted.

When `mysql_stmt_fetch()` is called, the MySQL client/server protocol performs these actions:

- The client reads the data from the packet row by row and places it into the application data buffers by doing the necessary conversions. If the application buffer type is same as that of the field type returned from the server, the conversions are straightforward.

If an error occurs, you can get the statement error code, error message, and SQLSTATE value using `mysql_stmt_errno()`, `mysql_stmt_error()`, and `mysql_stmt_sqlstate()`, respectively.

### Prepared Statement Logging

For prepared statements that are executed with the `mysql_stmt_prepare()` and `mysql_stmt_execute()` C API functions, the server writes `Prepare` and `Execute` lines to the general query log so that you can tell when statements are prepared and executed.

Suppose that you prepare and execute a statement as follows:

1. Call `mysql_stmt_prepare()` to prepare the statement string "`SELECT ?`".
2. Call `mysql_stmt_bind_param()` to bind the value `3` to the parameter in the prepared statement.
3. Call `mysql_stmt_execute()` to execute the prepared statement.

As a result of the preceding calls, the server writes the following lines to the general query log:

```
Prepare [1] SELECT ?
Execute [1] SELECT 3
```

Each `Prepare` and `Execute` line in the log is tagged with a `[N]` statement identifier so that you can keep track of which prepared statement is being logged. `N` is a positive integer. If there are multiple prepared statements active simultaneously for the client, `N` may be greater than 1. Each `Execute` lines shows a prepared statement after substitution of data values for `?` parameters.

Version notes: `Prepare` lines are displayed without `[N]` before MySQL 4.1.10. `Execute` lines are not displayed at all before MySQL 4.1.10.

## 17.2.7. C API Prepared Statement Function Descriptions

To prepare and execute queries, use the functions in the following sections.

**Note:** In MySQL 4.1.2, the names of several prepared statement functions were changed, as shown here:

Old Name	New Name
<code>mysql_bind_param()</code>	<code>mysql_stmt_bind_param()</code>

<code>mysql_bind_result()</code>	<code>mysql_stmt_bind_result()</code>
<code>mysql_prepare()</code>	<code>mysql_stmt_prepare()</code>
<code>mysql_execute()</code>	<code>mysql_stmt_execute()</code>
<code>mysql_fetch()</code>	<code>mysql_stmt_fetch()</code>
<code>mysql_fetch_column()</code>	<code>mysql_stmt_fetch_column()</code>
<code>mysql_param_count()</code>	<code>mysql_stmt_param_count()</code>
<code>mysql_param_result()</code>	<code>mysql_stmt_param_metadata()</code>
<code>mysql_get_metadata()</code>	<code>mysql_stmt_result_metadata()</code>
<code>mysql_send_long_data()</code>	<code>mysql_stmt_send_long_data()</code>

All functions that operate with a `MYSQL_STMT` structure begin with the prefix `mysql_stmt_`.

Also in 4.1.2, the signature of the `mysql_stmt_prepare()` function was changed to `int mysql_stmt_prepare(MYSQL_STMT *stmt, const char *query, unsigned long length)`. To create a `MYSQL_STMT` handle, you should use the `mysql_stmt_init()` function.

### 17.2.7.1. `mysql_stmt_affected_rows()`

```
my_ulonglong mysql_stmt_affected_rows(MYSQL_STMT *stmt)
```

#### Description

Returns the total number of rows changed, deleted, or inserted by the last executed statement. May be called immediately after `mysql_stmt_execute()` for `UPDATE`, `DELETE`, or `INSERT` statements. For `SELECT` statements, `mysql_stmt_affected_rows()` works like `mysql_num_rows()`.

This function was added in MySQL 4.1.0.

#### Return Values

An integer greater than zero indicates the number of rows affected or retrieved. Zero indicates that no records were updated for an `UPDATE` statement, no rows matched the `WHERE` clause in the query, or that no query has yet been executed. -1 indicates that the query returned an error or that, for a `SELECT` query, `mysql_stmt_affected_rows()` was called prior to calling `mysql_stmt_store_result()`. Because `mysql_stmt_affected_rows()` returns an unsigned value, you can check for -1 by comparing the return value to `(my_ulonglong)-1` (or to `(my_ulonglong)~0`, which is equivalent).

See [Section 17.2.3.1, “mysql\\_affected\\_rows\(\)”](#), for additional information on the return value.

#### Errors

None.

#### Example

For the usage of `mysql_stmt_affected_rows()`, refer to the Example from [Section 17.2.7.10, “mysql\\_stmt\\_execute\(\)”](#).

### 17.2.7.2. `mysql_stmt_attr_get()`

```
int mysql_stmt_attr_get(MYSQL_STMT *stmt, enum enum_stmt_attr_type option, void *arg)
```

**Description**

Can be used to get the current value for a statement attribute.

The `option` argument is the option that you want to get; the `arg` should point to a variable that should contain the option value. If the option is an integer, then `arg` should point to the value of the integer.

See `mysql_stmt_attr_set()` for a list of options and option types. See [Section 17.2.7.3](#), “`mysql_stmt_attr_set()`”.

**Note:** In MySQL 4.1, `mysql_stmt_attr_get()` uses `unsigned int *`, not `my_bool *`, for `STMT_ATTR_UPDATE_MAX_LENGTH`. This is corrected in MySQL 5.1.7.

This function was added in MySQL 4.1.2.

**Return Values**

0 if okay. Non-zero if `option` is unknown.

**Errors**

None.

### 17.2.7.3. `mysql_stmt_attr_set()`

```
int mysql_stmt_attr_set(MYSQL_STMT *stmt, enum enum_stmt_attr_type option, const void *arg)
```

**Description**

Can be used to affect behavior for a prepared statement. In MySQL 4.1, the `option` argument can take the single value `STMT_ATTR_UPDATE_MAX_LENGTH`; the `arg` argument is a pointer of type `my_bool *`. If `arg` points to the value 1, then the metadata `MYSQL_FIELD->max_length` in `mysql_stmt_store_result()` is updated when the prepared statement is executed.

**Note:** In MySQL 4.1, `mysql_stmt_attr_get()` uses `unsigned int *`, not `my_bool *`, for `STMT_ATTR_UPDATE_MAX_LENGTH`. This is corrected in MySQL 5.1.7.

This function was added in MySQL 4.1.2. Additional options are planned for this function in later versions of MySQL.

**Return Values**

0 if okay. Non-zero if `option` is unknown.

**Errors**

None.

### 17.2.7.4. `mysql_stmt_bind_param()`

```
my_bool mysql_stmt_bind_param(MYSQL_STMT *stmt, MYSQL_BIND *bind)
```

**Description**

`mysql_stmt_bind_param()` is used to bind data for the parameter markers in the SQL statement that was passed to `mysql_stmt_prepare()`. It uses `MYSQL_BIND` structures to supply the data. `bind` is the address of an array of `MYSQL_BIND` structures. The client library expects the array to contain an element for each ‘?’ parameter marker that is present in the query.



Suppose that you prepare the following statement:

```
INSERT INTO mytbl VALUES(?, ?, ?)
```

When you bind the parameters, the array of `MYSQL_BIND` structures must contain three elements, and can be declared like this:

```
MYSQL_BIND bind[3];
```

The members of each `MYSQL_BIND` element that should be set are described in [Section 17.2.5, “C API Prepared Statement Data types”](#).

This function was added in MySQL 4.1.2.

### Return Values

Zero if the bind was successful. Non-zero if an error occurred.

### Errors

- `CR_INVALID_BUFFER_USE`  
Indicates if the bind is to supply the long data in chunks and if the buffer type is non string or binary.
- `CR_UNSUPPORTED_PARAM_TYPE`  
The conversion is not supported. Possibly the `buffer_type` value is illegal or is not one of the supported types.
- `CR_OUT_OF_MEMORY`  
Out of memory.
- `CR_UNKNOWN_ERROR`  
An unknown error occurred.

### Example

For the usage of `mysql_stmt_bind_param()`, refer to the Example from [Section 17.2.7.10, “mysql\\_stmt\\_execute\(\)”](#).

## 17.2.7.5. `mysql_stmt_bind_result()`

```
my_bool mysql_stmt_bind_result(MYSQL_STMT *stmt, MYSQL_BIND *bind)
```

### Description

`mysql_stmt_bind_result()` is used to associate (bind) columns in the result set to data buffers and length buffers. When `mysql_stmt_fetch()` is called to fetch data, the MySQL client/server protocol places the data for the bound columns into the specified buffers.

All columns must be bound to buffers prior to calling `mysql_stmt_fetch()`. `bind` is the address of an array of `MYSQL_BIND` structures. The client library expects the array to contain an element for each column of the result set. If you do not bind columns to `MYSQL_BIND` structures, `mysql_stmt_fetch()` simply ignores the data fetch. The buffers should be large enough to hold the

data values, because the protocol doesn't return data values in chunks.

A column can be bound or rebound at any time, even after a result set has been partially retrieved. The new binding takes effect the next time `mysql_stmt_fetch()` is called. Suppose that an application binds the columns in a result set and calls `mysql_stmt_fetch()`. The client/server protocol returns data in the bound buffers. Then suppose that the application binds the columns to a different set of buffers. The protocol does not place data into the newly bound buffers until the next call to `mysql_stmt_fetch()` occurs.

To bind a column, an application calls `mysql_stmt_bind_result()` and passes the type, address, and the address of the length buffer. The members of each `MYSQL_BIND` element that should be set are described in [Section 17.2.5, “C API Prepared Statement Data types”](#).

This function was added in MySQL 4.1.2.

### Return Values

Zero if the bind was successful. Non-zero if an error occurred.

### Errors

- `CR_UNSUPPORTED_PARAM_TYPE`

The conversion is not supported. Possibly the `buffer_type` value is illegal or is not one of the supported types.

- `CR_OUT_OF_MEMORY`

Out of memory.

- `CR_UNKNOWN_ERROR`

An unknown error occurred.

### Example

For the usage of `mysql_stmt_bind_result()`, refer to the Example from [Section 17.2.7.11, “mysql\\_stmt\\_fetch\(\)”](#).

## 17.2.7.6. `mysql_stmt_close()`

```
my_bool mysql_stmt_close(MYSQL_STMT *)
```

### Description

Closes the prepared statement. `mysql_stmt_close()` also deallocates the statement handle pointed to by `stmt`.

If the current statement has pending or unread results, this function cancels them so that the next query can be executed.

This function was added in MySQL 4.1.0.

### Return Values

Zero if the statement was freed successfully. Non-zero if an error occurred.

### Errors

- `CR_SERVER_GONE_ERROR`

The MySQL server has gone away.

- `CR_UNKNOWN_ERROR`

An unknown error occurred.

### Example

For the usage of `mysql_stmt_close()`, refer to the Example from [Section 17.2.7.10](#), “`mysql_stmt_execute()`”.

### 17.2.7.7. `mysql_stmt_data_seek()`

```
void mysql_stmt_data_seek(MYSQL_STMT *stmt, my_ulonglong offset)
```

#### Description

Seeks to an arbitrary row in a statement result set. The `offset` value is a row number and should be in the range from 0 to `mysql_stmt_num_rows(stmt) - 1`.

This function requires that the statement result set structure contains the entire result of the last executed query, so `mysql_stmt_data_seek()` may be used only in conjunction with `mysql_stmt_store_result()`.

This function was added in MySQL 4.1.1.

#### Return Values

None.

#### Errors

None.

### 17.2.7.8. `mysql_stmt_errno()`

```
unsigned int mysql_stmt_errno(MYSQL_STMT *stmt)
```

#### Description

For the statement specified by `stmt`, `mysql_stmt_errno()` returns the error code for the most recently invoked statement API function that can succeed or fail. A return value of zero means that no error occurred. Client error message numbers are listed in the MySQL `errmsg.h` header file. Server error message numbers are listed in `mysqld_error.h`. Errors also are listed at [Appendix B, \*Error Codes and Messages\*](#).

This function was added in MySQL 4.1.0.

#### Return Values

An error code value. Zero if no error occurred.

#### Errors

None.

### 17.2.7.9. `mysql_stmt_error()`

```
const char *mysql_stmt_error(MYSQL_STMT *stmt)
```

#### Description

For the statement specified by `stmt`, `mysql_stmt_error()` returns a null-terminated string containing the error message for the most recently invoked statement API function that can succeed or fail. An empty string ( " ") is returned if no error occurred. This means the following two tests are equivalent:

```
if (mysql_stmt_errno(stmt))
{
    // an error occurred
}

if (mysql_stmt_error(stmt)[0])
{
    // an error occurred
}
```

The language of the client error messages may be changed by recompiling the MySQL client library. Currently, you can choose error messages in several different languages.

This function was added in MySQL 4.1.0.

#### Return Values

A character string that describes the error. An empty string if no error occurred.

#### Errors

None.

### 17.2.7.10. `mysql_stmt_execute()`

```
int mysql_stmt_execute(MYSQL_STMT *stmt)
```

#### Description

`mysql_stmt_execute()` executes the prepared query associated with the statement handle. The currently bound parameter marker values are sent to server during this call, and the server replaces the markers with this newly supplied data.

If the statement is an `UPDATE`, `DELETE`, or `INSERT`, the total number of changed, deleted, or inserted rows can be found by calling `mysql_stmt_affected_rows()`. If this is a statement such as `SELECT` that generates a result set, you must call `mysql_stmt_fetch()` to fetch the data prior to calling any other functions that result in query processing. For more information on how to fetch the results, refer to [Section 17.2.7.11](#), “`mysql_stmt_fetch()`”.

For statements that generate a result set, you can request that `mysql_stmt_execute()` open a cursor for the statement by calling `mysql_stmt_attr_set()` before executing the statement. If you execute a statement multiple times, `mysql_stmt_execute()` closes any open cursor before opening a new one.

This function was added in MySQL 4.1.2.

#### Return Values

Zero if execution was successful. Non-zero if an error occurred.

## Errors

- `CR_COMMANDS_OUT_OF_SYNC`  
Commands were executed in an improper order.
- `CR_OUT_OF_MEMORY`  
Out of memory.
- `CR_SERVER_GONE_ERROR`  
The MySQL server has gone away.
- `CR_SERVER_LOST`  
The connection to the server was lost during the query.
- `CR_UNKNOWN_ERROR`  
An unknown error occurred.

## Example

The following example demonstrates how to create and populate a table using `mysql_stmt_init()`, `mysql_stmt_prepare()`, `mysql_stmt_param_count()`, `mysql_stmt_bind_param()`, `mysql_stmt_execute()`, and `mysql_stmt_affected_rows()`. The `mysql` variable is assumed to be a valid connection handle.

```
#define STRING_SIZE 50

#define DROP_SAMPLE_TABLE "DROP TABLE IF EXISTS test_table"
#define CREATE_SAMPLE_TABLE "CREATE TABLE test_table(col1 INT,\
                                                    col2 VARCHAR(40),\
                                                    col3 SMALLINT,\
                                                    col4 TIMESTAMP)"

#define INSERT_SAMPLE "INSERT INTO test_table(col1,col2,col3) VALUES(?,?,?)"

MYSQL_STMT      *stmt;
MYSQL_BIND      bind[3];
my_ulonglong    affected_rows;
int             param_count;
short          small_data;
int            int_data;
char           str_data[STRING_SIZE];
unsigned long  str_length;
my_bool       is_null;

if (mysql_query(mysql, DROP_SAMPLE_TABLE))
{
    fprintf(stderr, " DROP TABLE failed\n");
    fprintf(stderr, " %s\n", mysql_error(mysql));
    exit(0);
}

if (mysql_query(mysql, CREATE_SAMPLE_TABLE))
{
    fprintf(stderr, " CREATE TABLE failed\n");
    fprintf(stderr, " %s\n", mysql_error(mysql));
    exit(0);
}

/* Prepare an INSERT query with 3 parameters */
/* (the TIMESTAMP column is not named; the server */
/* sets it to the current date and time) */
stmt = mysql_stmt_init(mysql);
if (!stmt)
{
```

```

    fprintf(stderr, " mysql_stmt_init(), out of memory\n");
    exit(0);
}
if (mysql_stmt_prepare(stmt, INSERT_SAMPLE, strlen(INSERT_SAMPLE)))
{
    fprintf(stderr, " mysql_stmt_prepare(), INSERT failed\n");
    fprintf(stderr, " %s\n", mysql_stmt_error(stmt));
    exit(0);
}
fprintf(stdout, " prepare, INSERT successful\n");

/* Get the parameter count from the statement */
param_count= mysql_stmt_param_count(stmt);
fprintf(stdout, " total parameters in INSERT: %d\n", param_count);

if (param_count != 3) /* validate parameter count */
{
    fprintf(stderr, " invalid parameter count returned by MySQL\n");
    exit(0);
}

/* Bind the data for all 3 parameters */

memset(bind, 0, sizeof(bind));

/* INTEGER PARAM */
/* This is a number type, so there is no need to specify buffer_length */
bind[0].buffer_type= MYSQL_TYPE_LONG;
bind[0].buffer= (char *)&int_data;
bind[0].is_null= 0;
bind[0].length= 0;

/* STRING PARAM */
bind[1].buffer_type= MYSQL_TYPE_STRING;
bind[1].buffer= (char *)str_data;
bind[1].buffer_length= STRING_SIZE;
bind[1].is_null= 0;
bind[1].length= &str_length;

/* SMALLINT PARAM */
bind[2].buffer_type= MYSQL_TYPE_SHORT;
bind[2].buffer= (char *)&small_data;
bind[2].is_null= &is_null;
bind[2].length= 0;

/* Bind the buffers */
if (mysql_stmt_bind_param(stmt, bind))
{
    fprintf(stderr, " mysql_stmt_bind_param() failed\n");
    fprintf(stderr, " %s\n", mysql_stmt_error(stmt));
    exit(0);
}

/* Specify the data values for the first row */
int_data= 10; /* integer */
strncpy(str_data, "MySQL", STRING_SIZE); /* string */
str_length= strlen(str_data);

/* INSERT SMALLINT data as NULL */
is_null= 1;

/* Execute the INSERT statement - 1*/
if (mysql_stmt_execute(stmt))
{
    fprintf(stderr, " mysql_stmt_execute(), 1 failed\n");
    fprintf(stderr, " %s\n", mysql_stmt_error(stmt));
    exit(0);
}

/* Get the total number of affected rows */
affected_rows= mysql_stmt_affected_rows(stmt);
fprintf(stdout, " total affected rows(insert 1): %lu\n",
        (unsigned long) affected_rows);

if (affected_rows != 1) /* validate affected rows */
{
    fprintf(stderr, " invalid affected rows by MySQL\n");
    exit(0);
}

```

```

/* Specify data values for second row, then re-execute the statement */
int_data= 1000;
strncpy(str_data, "The most popular Open Source database", STRING_SIZE);
str_length= strlen(str_data);
small_data= 1000;          /* smallint */
is_null= 0;                /* reset */

/* Execute the INSERT statement - 2*/
if (mysql_stmt_execute(stmt))
{
    fprintf(stderr, " mysql_stmt_execute, 2 failed\n");
    fprintf(stderr, " %s\n", mysql_stmt_error(stmt));
    exit(0);
}

/* Get the total rows affected */
affected_rows= mysql_stmt_affected_rows(stmt);
fprintf(stdout, " total affected rows(insert 2): %lu\n",
        (unsigned long) affected_rows);

if (affected_rows != 1) /* validate affected rows */
{
    fprintf(stderr, " invalid affected rows by MySQL\n");
    exit(0);
}

/* Close the statement */
if (mysql_stmt_close(stmt))
{
    fprintf(stderr, " failed while closing the statement\n");
    fprintf(stderr, " %s\n", mysql_stmt_error(stmt));
    exit(0);
}

```

**Note:** For complete examples on the use of prepared statement functions, refer to the file [tests/mysql\\_client\\_test.c](#). This file can be obtained from a MySQL source distribution or from the BitKeeper source repository.

### 17.2.7.11. `mysql_stmt_fetch()`

```
int mysql_stmt_fetch(MYSQL_STMT *stmt)
```

#### Description

`mysql_stmt_fetch()` returns the next row in the result set. It can be called only while the result set exists; that is, after a call to `mysql_stmt_execute()` that creates a result set or after `mysql_stmt_store_result()`, which is called after `mysql_stmt_execute()` to buffer the entire result set.

`mysql_stmt_fetch()` returns row data using the buffers bound by `mysql_stmt_bind_result()`. It returns the data in those buffers for all the columns in the current row set and the lengths are returned to the `length` pointer.

All columns must be bound by the application before calling `mysql_stmt_fetch()`.

If a fetched data value is a `NULL` value, the `*is_null` value of the corresponding `MYSQL_BIND` structure contains `TRUE` (1). Otherwise, the data and its length are returned in the `*buffer` and `*length` elements based on the buffer type specified by the application. Each numeric and temporal type has a fixed length, as listed in the following table. The length of the string types depends on the length of the actual data value, as indicated by `data_length`.

Type	Length
<code>MYSQL_TYPE_TINY</code>	1
<code>MYSQL_TYPE_SHORT</code>	2
<code>MYSQL_TYPE_LONG</code>	4

<code>MYSQL_TYPE_LONGLONG</code>	8
<code>MYSQL_TYPE_FLOAT</code>	4
<code>MYSQL_TYPE_DOUBLE</code>	8
<code>MYSQL_TYPE_TIME</code>	<code>sizeof(MYSQL_TIME)</code>
<code>MYSQL_TYPE_DATE</code>	<code>sizeof(MYSQL_TIME)</code>
<code>MYSQL_TYPE_DATETIME</code>	<code>sizeof(MYSQL_TIME)</code>
<code>MYSQL_TYPE_STRING</code>	<code>data_length</code>
<code>MYSQL_TYPE_BLOB</code>	<code>data_length</code>

This function was added in MySQL 4.1.2.

### Return Values

Return Value	Description
0	Successful, the data has been fetched to application data buffers.
1	Error occurred. Error code and message can be obtained by calling <code>mysql_stmt_errno()</code> and <code>mysql_stmt_error()</code> .
<code>MYSQL_NO_DATA</code>	No more rows/data exists
<code>MYSQL_DATA_TRUNCATED</code>	Data truncation occurred

`MYSQL_DATA_TRUNCATED` is not returned unless truncation reporting is enabled with `mysql_options()`. To determine which parameters were truncated when this value is returned, check the `error` members of the `MYSQL_BIND` parameter structures.

### Errors

- `CR_COMMANDS_OUT_OF_SYNC`  
Commands were executed in an improper order.
- `CR_OUT_OF_MEMORY`  
Out of memory.
- `CR_SERVER_GONE_ERROR`  
The MySQL server has gone away.
- `CR_SERVER_LOST`  
The connection to the server was lost during the query.
- `CR_UNKNOWN_ERROR`  
An unknown error occurred.
- `CR_UNSUPPORTED_PARAM_TYPE`  
The buffer type is `MYSQL_TYPE_DATE`, `MYSQL_TYPE_TIME`, `MYSQL_TYPE_DATETIME`, or `MYSQL_TYPE_TIMESTAMP`, but the data type is not `DATE`, `TIME`, `DATETIME`, or `TIMESTAMP`.



- All other unsupported conversion errors are returned from `mysql_stmt_bind_result()`.

### Example

The following example demonstrates how to fetch data from a table using `mysql_stmt_result_metadata()`, `mysql_stmt_bind_result()`, and `mysql_stmt_fetch()`. (This example expects to retrieve the two rows inserted by the example shown in [Section 17.2.7.10](#), “`mysql_stmt_execute()`”.) The `mysql` variable is assumed to be a valid connection handle.

```
#define STRING_SIZE 50

#define SELECT_SAMPLE "SELECT coll, col2, col3, col4 FROM test_table"

MYSQL_STMT *stmt;
MYSQL_BIND bind[4];
MYSQL_RES *prepare_meta_result;
MYSQL_TIME ts;
unsigned long length[4];
int param_count, column_count, row_count;
short small_data;
int int_data;
char str_data[STRING_SIZE];
my_bool is_null[4];

/* Prepare a SELECT query to fetch data from test_table */
stmt = mysql_stmt_init(mysql);
if (!stmt)
{
    fprintf(stderr, " mysql_stmt_init(), out of memory\n");
    exit(0);
}
if (mysql_stmt_prepare(stmt, SELECT_SAMPLE, strlen(SELECT_SAMPLE)))
{
    fprintf(stderr, " mysql_stmt_prepare(), SELECT failed\n");
    fprintf(stderr, " %s\n", mysql_stmt_error(stmt));
    exit(0);
}
fprintf(stdout, " prepare, SELECT successful\n");

/* Get the parameter count from the statement */
param_count= mysql_stmt_param_count(stmt);
fprintf(stdout, " total parameters in SELECT: %d\n", param_count);

if (param_count != 0) /* validate parameter count */
{
    fprintf(stderr, " invalid parameter count returned by MySQL\n");
    exit(0);
}

/* Fetch result set meta information */
prepare_meta_result = mysql_stmt_result_metadata(stmt);
if (!prepare_meta_result)
{
    fprintf(stderr,
        " mysql_stmt_result_metadata(), returned no meta information\n");
    fprintf(stderr, " %s\n", mysql_stmt_error(stmt));
    exit(0);
}

/* Get total columns in the query */
column_count= mysql_num_fields(prepare_meta_result);
fprintf(stdout, " total columns in SELECT statement: %d\n", column_count);

if (column_count != 4) /* validate column count */
{
    fprintf(stderr, " invalid column count returned by MySQL\n");
    exit(0);
}

/* Execute the SELECT query */
if (mysql_stmt_execute(stmt))
{
    fprintf(stderr, " mysql_stmt_execute(), failed\n");
    fprintf(stderr, " %s\n", mysql_stmt_error(stmt));
    exit(0);
}
```

```

}

/* Bind the result buffers for all 4 columns before fetching them */
memset(bind, 0, sizeof(bind));

/* INTEGER COLUMN */
bind[0].buffer_type= MYSQL_TYPE_LONG;
bind[0].buffer= (char *)&int_data;
bind[0].is_null= &is_null[0];
bind[0].length= &length[0];

/* STRING COLUMN */
bind[1].buffer_type= MYSQL_TYPE_STRING;
bind[1].buffer= (char *)str_data;
bind[1].buffer_length= STRING_SIZE;
bind[1].is_null= &is_null[1];
bind[1].length= &length[1];

/* SMALLINT COLUMN */
bind[2].buffer_type= MYSQL_TYPE_SHORT;
bind[2].buffer= (char *)&small_data;
bind[2].is_null= &is_null[2];
bind[2].length= &length[2];

/* TIMESTAMP COLUMN */
bind[3].buffer_type= MYSQL_TYPE_TIMESTAMP;
bind[3].buffer= (char *)&ts;
bind[3].is_null= &is_null[3];
bind[3].length= &length[3];

/* Bind the result buffers */
if (mysql_stmt_bind_result(stmt, bind))
{
    fprintf(stderr, "mysql_stmt_bind_result() failed\n");
    fprintf(stderr, " %s\n", mysql_stmt_error(stmt));
    exit(0);
}

/* Now buffer all results to client */
if (mysql_stmt_store_result(stmt))
{
    fprintf(stderr, "mysql_stmt_store_result() failed\n");
    fprintf(stderr, " %s\n", mysql_stmt_error(stmt));
    exit(0);
}

/* Fetch all rows */
row_count= 0;
fprintf(stdout, "Fetching results ...\n");
while (!mysql_stmt_fetch(stmt))
{
    row_count++;
    fprintf(stdout, " row %d\n", row_count);

    /* column 1 */
    fprintf(stdout, " column1 (integer) : ");
    if (is_null[0])
        fprintf(stdout, " NULL\n");
    else
        fprintf(stdout, " %d(%ld)\n", int_data, length[0]);

    /* column 2 */
    fprintf(stdout, " column2 (string) : ");
    if (is_null[1])
        fprintf(stdout, " NULL\n");
    else
        fprintf(stdout, " %s(%ld)\n", str_data, length[1]);

    /* column 3 */
    fprintf(stdout, " column3 (smallint) : ");
    if (is_null[2])
        fprintf(stdout, " NULL\n");
    else
        fprintf(stdout, " %d(%ld)\n", small_data, length[2]);

    /* column 4 */
    fprintf(stdout, " column4 (timestamp): ");
    if (is_null[3])
        fprintf(stdout, " NULL\n");
}

```

```

else
    fprintf(stdout, " %04d-%02d-%02d %02d:%02d:%02d (%ld)\n",
              ts.year, ts.month, ts.day,
              ts.hour, ts.minute, ts.second,
              length[3]);
    fprintf(stdout, "\n");
}

/* Validate rows fetched */
fprintf(stdout, " total rows fetched: %d\n", row_count);
if (row_count != 2)
{
    fprintf(stderr, " MySQL failed to return all rows\n");
    exit(0);
}

/* Free the prepared result metadata */
mysql_free_result(prepare_meta_result);

/* Close the statement */
if (mysql_stmt_close(stmt))
{
    fprintf(stderr, " failed while closing the statement\n");
    fprintf(stderr, " %s\n", mysql_stmt_error(stmt));
    exit(0);
}

```

### 17.2.7.12. `mysql_stmt_fetch_column()`

```
int mysql_stmt_fetch_column(MYSQL_STMT *stmt, MYSQL_BIND *bind, unsigned int column, unsigned long offset)
```

#### Description

Fetch one column from the current result set row. `bind` provides the buffer where data should be placed. It should be set up the same way as for `mysql_stmt_bind_result()`. `column` indicates which column to fetch. The first column is numbered 0. `offset` is the offset within the data value at which to begin retrieving data. This can be used for fetching the data value in pieces. The beginning of the value is offset 0.

This function was added in MySQL 4.1.2.

#### Return Values

Zero if the value was fetched successfully. Non-zero if an error occurred.

#### Errors

- `CR_INVALID_PARAMETER_NO`  
Invalid column number.
- `CR_NO_DATA`  
The end of the result set has already been reached.

### 17.2.7.13. `mysql_stmt_field_count()`

```
unsigned int mysql_stmt_field_count(MYSQL_STMT *stmt)
```

#### Description

Returns the number of columns for the most recent statement for the statement handler. This value is zero for statements such as `INSERT` or `DELETE` that do not produce result sets.

`mysql_stmt_field_count()` can be called after you have prepared a statement by invoking `mysql_stmt_prepare()`.

This function was added in MySQL 4.1.3.

#### **Return Values**

An unsigned integer representing the number of columns in a result set.

#### **Errors**

None.

### **17.2.7.14. `mysql_stmt_free_result()`**

```
my_bool mysql_stmt_free_result(MYSQL_STMT *stmt)
```

#### **Description**

Releases memory associated with the result set produced by execution of the prepared statement. If there is a cursor open for the statement, `mysql_stmt_free_result()` closes it.

This function was added in MySQL 4.1.1.

#### **Return Values**

Zero if the result set was freed successfully. Non-zero if an error occurred.

#### **Errors**

### **17.2.7.15. `mysql_stmt_init()`**

```
MYSQL_STMT *mysql_stmt_init(MYSQL *mysql)
```

#### **Description**

Create a `MYSQL_STMT` handle. The handle should be freed with `mysql_stmt_close(MYSQL_STMT *)`.

This function was added in MySQL 4.1.2.

#### **Return values**

A pointer to a `MYSQL_STMT` structure in case of success. `NULL` if out of memory.

#### **Errors**

- `CR_OUT_OF_MEMORY`

Out of memory.

### **17.2.7.16. `mysql_stmt_insert_id()`**

```
my_ulonglong mysql_stmt_insert_id(MYSQL_STMT *stmt)
```

**Description**

Returns the value generated for an `AUTO_INCREMENT` column by the prepared `INSERT` or `UPDATE` statement. Use this function after you have executed prepared `INSERT` statement into a table which contains an `AUTO_INCREMENT` field.

See [Section 17.2.3.35](#), “`mysql_insert_id()`”, for more information.

This function was added in MySQL 4.1.2.

**Return Values**

Value for `AUTO_INCREMENT` column which was automatically generated or explicitly set during execution of prepared statement, or value generated by `LAST_INSERT_ID(expr)` function. Return value is undefined if statement does not set `AUTO_INCREMENT` value.

**Errors**

None.

**17.2.7.17. `mysql_stmt_num_rows()`**

```
my_ulonglong mysql_stmt_num_rows(MYSQL_STMT *stmt)
```

**Description**

Returns the number of rows in the result set.

The use of `mysql_stmt_num_rows()` depends on whether you used `mysql_stmt_store_result()` to buffer the entire result set in the statement handle.

If you use `mysql_stmt_store_result()`, `mysql_stmt_num_rows()` may be called immediately.

This function was added in MySQL 4.1.1.

**Return Values**

The number of rows in the result set.

**Errors**

None.

**17.2.7.18. `mysql_stmt_param_count()`**

```
unsigned long mysql_stmt_param_count(MYSQL_STMT *stmt)
```

**Description**

Returns the number of parameter markers present in the prepared statement.

This function was added in MySQL 4.1.2.

**Return Values**

An unsigned long integer representing the number of parameters in a statement.

**Errors**

None.

**Example**

For the usage of `mysql_stmt_param_count()`, refer to the Example from [Section 17.2.7.10](#), “`mysql_stmt_execute()`”.

### 17.2.7.19. `mysql_stmt_param_metadata()`

```
MYSQL_RES *mysql_stmt_param_metadata(MYSQL_STMT *stmt)
```

This function currently does nothing.

This function was added in MySQL 4.1.2.

**Description****Return Values****Errors**

### 17.2.7.20. `mysql_stmt_prepare()`

```
int mysql_stmt_prepare(MYSQL_STMT *stmt, const char *query, unsigned long length)
```

**Description**

Given the statement handle returned by `mysql_stmt_init()`, prepares the SQL statement pointed to by the string `query` and returns a status value. The string length should be given by the `length` argument. The string must consist of a single SQL statement. You should not add a terminating semicolon (`;`) or `\g` to the statement.

The application can include one or more parameter markers in the SQL statement by embedding question mark (`?`) characters into the SQL string at the appropriate positions.

The markers are legal only in certain places in SQL statements. For example, they are allowed in the `VALUES()` list of an `INSERT` statement (to specify column values for a row), or in a comparison with a column in a `WHERE` clause to specify a comparison value. However, they are not allowed for identifiers (such as table or column names), or to specify both operands of a binary operator such as the `=` equal sign. The latter restriction is necessary because it would be impossible to determine the parameter type. In general, parameters are legal only in Data Manipulation Language (DML) statements, and not in Data Definition Language (DDL) statements.

The parameter markers must be bound to application variables using `mysql_stmt_bind_param()` before executing the statement.

This function was added in MySQL 4.1.2.

**Return Values**

Zero if the statement was prepared successfully. Non-zero if an error occurred.

**Errors**

- `CR_COMMANDS_OUT_OF_SYNC`

Commands were executed in an improper order.

- `CR_OUT_OF_MEMORY`  
Out of memory.
- `CR_SERVER_GONE_ERROR`  
The MySQL server has gone away.
- `CR_SERVER_LOST`  
The connection to the server was lost during the query
- `CR_UNKNOWN_ERROR`  
An unknown error occurred.

If the prepare operation was unsuccessful (that is, `mysql_stmt_prepare()` returns non-zero), the error message can be obtained by calling `mysql_stmt_error()`.

#### Example

For the usage of `mysql_stmt_prepare()`, refer to the Example from [Section 17.2.7.10](#), “`mysql_stmt_execute()`”.

### 17.2.7.21. `mysql_stmt_reset()`

```
my_bool mysql_stmt_reset(MYSQL_STMT *stmt)
```

#### Description

Reset the prepared statement on the client and server to state after prepare. This is mainly used to reset data sent with `mysql_stmt_send_long_data()`. Any open cursor for the statement is closed.

To re-prepare the statement with another query, use `mysql_stmt_prepare()`.

This function was added in MySQL 4.1.1.

#### Return Values

Zero if the statement was reset successfully. Non-zero if an error occurred.

#### Errors

- `CR_COMMANDS_OUT_OF_SYNC`  
Commands were executed in an improper order.
- `CR_SERVER_GONE_ERROR`  
The MySQL server has gone away.
- `CR_SERVER_LOST`  
The connection to the server was lost during the query
- `CR_UNKNOWN_ERROR`  
An unknown error occurred.

### 17.2.7.22. `mysql_stmt_result_metadata()`

```
MYSQL_RES *mysql_stmt_result_metadata(MYSQL_STMT *stmt)
```

#### Description

If a statement passed to `mysql_stmt_prepare()` is one that produces a result set, `mysql_stmt_result_metadata()` returns the result set metadata in the form of a pointer to a `MYSQL_RES` structure that can be used to process the meta information such as total number of fields and individual field information. This result set pointer can be passed as an argument to any of the field-based API functions that process result set metadata, such as:

- `mysql_num_fields()`
- `mysql_fetch_field()`
- `mysql_fetch_field_direct()`
- `mysql_fetch_fields()`
- `mysql_field_count()`
- `mysql_field_seek()`
- `mysql_field_tell()`
- `mysql_free_result()`

The result set structure should be freed when you are done with it, which you can do by passing it to `mysql_free_result()`. This is similar to the way you free a result set obtained from a call to `mysql_store_result()`.

The result set returned by `mysql_stmt_result_metadata()` contains only metadata. It does not contain any row results. The rows are obtained by using the statement handle with `mysql_stmt_fetch()`.

This function was added in MySQL 4.1.2.

#### Return Values

A `MYSQL_RES` result structure. `NULL` if no meta information exists for the prepared query.

#### Errors

- `CR_OUT_OF_MEMORY`  
Out of memory.
- `CR_UNKNOWN_ERROR`  
An unknown error occurred.

#### Example

For the usage of `mysql_stmt_result_metadata()`, refer to the Example from [Section 17.2.7.11](#),



`“mysql_stmt_fetch()”`.

### 17.2.7.23. `mysql_stmt_row_seek()`

```
MYSQL_ROW_OFFSET mysql_stmt_row_seek(MYSQL_STMT *stmt,  
MYSQL_ROW_OFFSET offset)
```

#### Description

Sets the row cursor to an arbitrary row in a statement result set. The `offset` value is a row offset that should be a value returned from `mysql_stmt_row_tell()` or from `mysql_stmt_row_seek()`. This value is not a row number; if you want to seek to a row within a result set by number, use `mysql_stmt_data_seek()` instead.

This function requires that the result set structure contains the entire result of the query, so `mysql_stmt_row_seek()` may be used only in conjunction with `mysql_stmt_store_result()`.

This function was added in MySQL 4.1.1.

#### Return Values

The previous value of the row cursor. This value may be passed to a subsequent call to `mysql_stmt_row_seek()`.

#### Errors

None.

### 17.2.7.24. `mysql_stmt_row_tell()`

```
MYSQL_ROW_OFFSET mysql_stmt_row_tell(MYSQL_STMT *stmt)
```

#### Description

Returns the current position of the row cursor for the last `mysql_stmt_fetch()`. This value can be used as an argument to `mysql_stmt_row_seek()`.

You should use `mysql_stmt_row_tell()` only after `mysql_stmt_store_result()`.

This function was added in MySQL 4.1.1.

#### Return Values

The current offset of the row cursor.

#### Errors

None.

### 17.2.7.25. `mysql_stmt_send_long_data()`

```
my_bool mysql_stmt_send_long_data(MYSQL_STMT *stmt, unsigned int para-  
meter_number, const char *data, unsigned long length)
```

#### Description

Allows an application to send parameter data to the server in pieces (or “chunks”). This function can be called multiple times to send the parts of a character or binary data value for a column, which must be

one of the `TEXT` or `BLOB` data types.

`parameter_number` indicates which parameter to associate the data with. Parameters are numbered beginning with 0. `data` is a pointer to a buffer containing data to be sent, and `length` indicates the number of bytes in the buffer.

**Note:** The next `mysql_stmt_execute()` call ignores the bind buffer for all parameters that have been used with `mysql_stmt_send_long_data()` since last `mysql_stmt_execute()` or `mysql_stmt_reset()`.

If you want to reset/forget the sent data, you can do it with `mysql_stmt_reset()`. See [Section 17.2.7.21](#), “`mysql_stmt_reset()`”.

This function was added in MySQL 4.1.2.

### Return Values

Zero if the data is sent successfully to server. Non-zero if an error occurred.

### Errors

- `CR_COMMANDS_OUT_OF_SYNC`  
Commands were executed in an improper order.
- `CR_SERVER_GONE_ERROR`  
The MySQL server has gone away.
- `CR_OUT_OF_MEMORY`  
Out of memory.
- `CR_UNKNOWN_ERROR`  
An unknown error occurred.

### Example

The following example demonstrates how to send the data for a `TEXT` column in chunks. It inserts the data value 'MySQL - The most popular Open Source database' into the `text_column` column. The `mysql` variable is assumed to be a valid connection handle.

```
#define INSERT_QUERY "INSERT INTO test_long_data(text_column) VALUES(?)"

MYSQL_BIND bind[1];
long        length;

stmt = mysql_stmt_init(mysql);
if (!stmt)
{
    fprintf(stderr, "mysql_stmt_init(), out of memory\n");
    exit(0);
}
if (mysql_stmt_prepare(stmt, INSERT_QUERY, strlen(INSERT_QUERY)))
{
    fprintf(stderr, "\nmysql_stmt_prepare(), INSERT failed");
    fprintf(stderr, "\n %s", mysql_stmt_error(stmt));
    exit(0);
}
memset(bind, 0, sizeof(bind));
bind[0].buffer_type= MYSQL_TYPE_STRING;
bind[0].length= &length;
```

```

bind[0].is_null= 0;

/* Bind the buffers */
if (mysql_stmt_bind_param(stmt, bind))
{
    fprintf(stderr, "\n param bind failed");
    fprintf(stderr, "\n %s", mysql_stmt_error(stmt));
    exit(0);
}

/* Supply data in chunks to server */
if (!mysql_stmt_send_long_data(stmt,0,"MySQL",5))
{
    fprintf(stderr, "\n send_long_data failed");
    fprintf(stderr, "\n %s", mysql_stmt_error(stmt));
    exit(0);
}

/* Supply the next piece of data */
if (!mysql_stmt_send_long_data(stmt,0," - The most popular Open Source database",40))
{
    fprintf(stderr, "\n send_long_data failed");
    fprintf(stderr, "\n %s", mysql_stmt_error(stmt));
    exit(0);
}

/* Now, execute the query */
if (mysql_stmt_execute(stmt))
{
    fprintf(stderr, "\n mysql_stmt_execute failed");
    fprintf(stderr, "\n %s", mysql_stmt_error(stmt));
    exit(0);
}

```

### 17.2.7.26. `mysql_stmt_sqlstate()`

```
const char *mysql_stmt_sqlstate(MYSQL_STMT *stmt)
```

#### Description

For the statement specified by `stmt`, `mysql_stmt_sqlstate()` returns a null-terminated string containing the SQLSTATE error code for the most recently invoked prepared statement API function that can succeed or fail. The error code consists of five characters. "00000" means "no error." The values are specified by ANSI SQL and ODBC. For a list of possible values, see [Appendix B, Error Codes and Messages](#).

Note that not all MySQL errors are yet mapped to SQLSTATE codes. The value "HY000" (general error) is used for unmapped errors.

This function was added to MySQL 4.1.1.

#### Return Values

A null-terminated character string containing the SQLSTATE error code.

### 17.2.7.27. `mysql_stmt_store_result()`

```
int mysql_stmt_store_result(MYSQL_STMT *stmt)
```

#### Description

You must call `mysql_stmt_store_result()` for every statement that successfully produces a result set (`SELECT`, `SHOW`, `DESCRIBE`, `EXPLAIN`), and only if you want to buffer the complete result set by the client, so that the subsequent `mysql_stmt_fetch()` call returns buffered data.

It is unnecessary to call `mysql_stmt_store_result()` for other statements, but if you do, it does

not harm or cause any notable performance problem. You can detect whether the statement produced a result set by checking if `mysql_stmt_result_metadata()` returns `NULL`. For more information, refer to [Section 17.2.7.22](#), “`mysql_stmt_result_metadata()`”.

**Note:** MySQL doesn't by default calculate `MYSQL_FIELD->max_length` for all columns in `mysql_stmt_store_result()` because calculating this would slow down `mysql_stmt_store_result()` considerably and most applications doesn't need `max_length`. If you want `max_length` to be updated, you can call `mysql_stmt_attr_set(MYSQL_STMT, STMT_ATTR_UPDATE_MAX_LENGTH, &flag)` to enable this. See [Section 17.2.7.3](#), “`mysql_stmt_attr_set()`”.

This function was added in MySQL 4.1.0.

### Return Values

Zero if the results are buffered successfully. Non-zero if an error occurred.

### Errors

- `CR_COMMANDS_OUT_OF_SYNC`  
Commands were executed in an improper order.
- `CR_OUT_OF_MEMORY`  
Out of memory.
- `CR_SERVER_GONE_ERROR`  
The MySQL server has gone away.
- `CR_SERVER_LOST`  
The connection to the server was lost during the query.
- `CR_UNKNOWN_ERROR`  
An unknown error occurred.

## 17.2.8. C API Prepared statement problems

Here follows a list of the currently known problems with prepared statements:

- `TIME`, `TIMESTAMP`, and `DATETIME` do not support parts of seconds (for example from `DATE_FORMAT()`).
- When converting an integer to string, `ZEROFILL` is honored with prepared statements in some cases where the MySQL server doesn't print the leading zeros. (For example, with `MIN(number-with-zerofill)`).
- When converting a floating point number to a string in the client, the rightmost digits of the converted value may differ slightly from those of the original value.
- *Prepared statements do not use the Query Cache, even in cases where a query does not contain any placeholders. See [Section 5.13.1](#), “How the Query Cache Operates”.*

## 17.2.9. C API Handling of Multiple Statement Execution

From version 4.1, MySQL supports the execution of multiple statements specified in a single query string. To use this capability with a given connection, you must specify the `CLIENT_MULTI_STATEMENTS` option in the flags parameter of `mysql_real_connect()` when opening the connection. You can also set this for an existing connection by calling `mysql_set_server_option(MYSQL_OPTION_MULTI_STATEMENTS_ON)`

By default, `mysql_query()` and `mysql_real_query()` return only the first query status and the subsequent queries status can be processed using `mysql_more_results()` and `mysql_next_result()`.

If you enable multiple-statement support, you should process the results from `mysql_query()` and `mysql_real_query()` within a loop that checks for more results. This is true even for statements such as `DROP TABLE` that return a result but not a result set. Failure to process the result this way may result in a dropped connection to the server.

```
/* Connect to server with option CLIENT_MULTI_STATEMENTS */
mysql_real_connect(..., CLIENT_MULTI_STATEMENTS);

/* Now execute multiple queries */
mysql_query(mysql, "DROP TABLE IF EXISTS test_table;\
                  CREATE TABLE test_table(id INT);\
                  INSERT INTO test_table VALUES(10);\
                  UPDATE test_table SET id=20 WHERE id=10;\
                  SELECT * FROM test_table;\
                  DROP TABLE test_table");

do
{
    /* Process all results */
    ...
    printf("total affected rows: %lld", mysql_affected_rows(mysql));
    ...
    if (!(result= mysql_store_result(mysql)))
    {
        printf(stderr, "Got fatal error processing query\n");
        exit(1);
    }
    process_result_set(result); /* client function */
    mysql_free_result(result);
} while (!mysql_next_result(mysql));
```

The multiple-statement capability can be used with `mysql_query()` or `mysql_real_query()`. It cannot be used with the prepared statement interface. Prepared statement handles are defined to work only with strings that contain a single statement.

## 17.2.10. C API Handling of Date and Time Values

The binary protocol available in MySQL 4.1 and above allows you to send and receive date and time values (`DATE`, `TIME`, `DATETIME`, and `TIMESTAMP`), using the `MYSQL_TIME` structure. The members of this structure are described in [Section 17.2.5, “C API Prepared Statement Data types”](#).

To send temporal data values, create a prepared statement using `mysql_stmt_prepare()`. Then, before calling `mysql_stmt_execute()` to execute the statement, use the following procedure to set up each temporal parameter:

1. In the `MYSQL_BIND` structure associated with the data value, set the `buffer_type` member to the type that indicates what kind of temporal value you're sending. For `DATE`, `TIME`, `DATETIME`, or `TIMESTAMP` values, set `buffer_type` to `MYSQL_TYPE_DATE`, `MYSQL_TYPE_TIME`, `MYSQL_TYPE_DATETIME`, or `MYSQL_TYPE_TIMESTAMP`, respectively.
2. Set the `buffer` member of the `MYSQL_BIND` structure to the address of the `MYSQL_TIME` structure in which you pass the temporal value.

- Fill in the members of the `MYSQL_TIME` structure that are appropriate for the type of temporal value to be passed.

Use `mysql_stmt_bind_param()` to bind the parameter data to the statement. Then you can call `mysql_stmt_execute()`.

To retrieve temporal values, the procedure is similar, except that you set the `buffer_type` member to the type of value you expect to receive, and the `buffer` member to the address of a `MYSQL_TIME` structure into which the returned value should be placed. Use `mysql_bind_results()` to bind the buffers to the statement after calling `mysql_stmt_execute()` and before fetching the results.

Here is a simple example that inserts `DATE`, `TIME`, and `TIMESTAMP` data. The `mysql` variable is assumed to be a valid connection handle.

```

MYSQL_TIME  ts;
MYSQL_BIND  bind[3];
MYSQL_STMT  *stmt;

strmov(query, "INSERT INTO test_table(date_field, time_field,
                                     timestamp_field) VALUES(?,?,?);

stmt = mysql_stmt_init(mysql);
if (!stmt)
{
    fprintf(stderr, " mysql_stmt_init(), out of memory\n");
    exit(0);
}
if (mysql_stmt_prepare(mysql, query, strlen(query)))
{
    fprintf(stderr, "\n mysql_stmt_prepare(), INSERT failed");
    fprintf(stderr, "\n %s", mysql_stmt_error(stmt));
    exit(0);
}

/* set up input buffers for all 3 parameters */
bind[0].buffer_type= MYSQL_TYPE_DATE;
bind[0].buffer= (char *)&ts;
bind[0].is_null= 0;
bind[0].length= 0;
...
bind[1]= bind[2]= bind[0];
...

mysql_stmt_bind_param(stmt, bind);

/* supply the data to be sent in the ts structure */
ts.year= 2002;
ts.month= 02;
ts.day= 03;

ts.hour= 10;
ts.minute= 45;
ts.second= 20;

mysql_stmt_execute(stmt);
..

```

## 17.2.11. C API Threaded Function Descriptions

You need to use the following functions when you want to create a threaded client. See [Section 17.2.15, “How to Make a Threaded Client”](#).

### 17.2.11.1. `my_init()`

```
void my_init(void)
```

**Description**

This function needs to be called once in the program before calling any MySQL function. This initializes some global variables that MySQL needs. If you are using a thread-safe client library, this also calls `mysql_thread_init()` for this thread.

This is automatically called by `mysql_init()`, `mysql_library_init()`, `mysql_server_init()`, and `mysql_connect()`.

**Return Values**

None.

**17.2.11.2. `mysql_thread_init()`**

```
my_bool mysql_thread_init(void)
```

**Description**

This function needs to be called for each created thread to initialize thread-specific variables.

This is automatically called by `my_init()` and `mysql_connect()`.

**Return Values**

Zero if successful. Non-zero if an error occurred.

**17.2.11.3. `mysql_thread_end()`**

```
void mysql_thread_end(void)
```

**Description**

This function needs to be called before calling `pthread_exit()` to free memory allocated by `mysql_thread_init()`.

Note that this function *is not invoked automatically by the client library*. It must be called explicitly to avoid a memory leak.

**Return Values**

None.

**17.2.11.4. `mysql_thread_safe()`**

```
unsigned int mysql_thread_safe(void)
```

**Description**

This function indicates whether the client is compiled as thread-safe.

**Return Values**

1 is the client is thread-safe, 0 otherwise.

**17.2.12. C API Embedded Server Function Descriptions**

If you want to allow your application to be linked against the embedded MySQL server library, you

must use the `mysql_server_init()` and `mysql_server_end()` functions. See [Section 17.1](#), “[libmysqld, the Embedded MySQL Server Library](#)”.

However, to provide improved memory management, even programs that are linked with `-lmysqlclient` rather than `-libmysqld` should include calls to begin and end use of the library. As of MySQL 4.1.10, the `mysql_library_init()` and `mysql_library_end()` functions can be used to do this. These actually are `#define` symbols that make them equivalent to `mysql_server_init()` and `mysql_server_end()`, but the names more clearly indicate that they should be called when beginning and ending use of a MySQL C API library no matter whether the application uses `libmysqlclient` or `libmysqld`. For more information, see [Section 17.2.2](#), “[C API Function Overview](#)”.

### 17.2.12.1. `mysql_server_init()`

```
int mysql_server_init(int argc, char **argv, char **groups)
```

#### Description

This function **must** be called once in the program using the embedded server before calling any other MySQL function. It starts the server and initializes any subsystems (`mysys`, `InnoDB`, and so forth) that the server uses. If this function is not called, the next call to `mysql_init()` executes `mysql_server_init()`.

In a non-multi-threaded environment, the call to `mysql_server_init()` may be omitted, because `mysql_init()` will invoke it automatically as necessary. However, a race condition is possible if `mysql_server_init()` is invoked by `mysql_init()` in a multi-threaded environment: `mysql_server_init()` is not thread-safe, so it should be called prior to any other client library call.

If you are using the `DEBUG` package that comes with MySQL, you should call `mysql_server_init()` after you have called `my_init()`.

The `argc` and `argv` arguments are analogous to the arguments to `main()`. The first element of `argv` is ignored (it typically contains the program name). For convenience, `argc` may be 0 (zero) if there are no command-line arguments for the server. `mysql_server_init()` makes a copy of the arguments so it's safe to destroy `argv` or `groups` after the call.

If you want to connect to an external server without starting the embedded server, you have to specify a negative value for `argc`.

The `NULL`-terminated list of strings in `groups` selects which groups in the option files are active. See [Section 4.3.2](#), “[Using Option Files](#)”. For convenience, `groups` may be `NULL`, in which case the `[server]` and `[embedded]` groups are active.

#### Example

```
#include <mysql.h>
#include <stdlib.h>

static char *server_args[] = {
    "this_program",          /* this string is not used */
    "--datadir=",
    "--key_buffer_size=32M"
};
static char *server_groups[] = {
    "embedded",
    "server",
    "this_program_SERVER",
    (char *)NULL
};

int main(void) {
    if (mysql_server_init(sizeof(server_args) / sizeof(char *),
```



```
        server_args, server_groups))
    exit(1);
/* Use any MySQL API functions here */
mysql_server_end();
return EXIT_SUCCESS;
}
```

**Return Values**

0 if okay, 1 if an error occurred.

### 17.2.12.2. `mysql_server_end()`

```
void mysql_server_end(void)
```

**Description**

This function **must** be called once in the program after all other MySQL functions. It shuts down the embedded server.

**Return Values**

None.

## 17.2.13. Common Questions and Problems When Using the C API

### 17.2.13.1. Why `mysql_store_result()` Sometimes Returns `NULL` After `mysql_query()` Returns Success

It is possible for `mysql_store_result()` to return `NULL` following a successful call to `mysql_query()`. When this happens, it means one of the following conditions occurred:

- There was a `malloc()` failure (for example, if the result set was too large).
- The data couldn't be read (an error occurred on the connection).
- The query returned no data (for example, it was an `INSERT`, `UPDATE`, or `DELETE`).

You can always check whether the statement should have produced a non-empty result by calling `mysql_field_count()`. If `mysql_field_count()` returns zero, the result is empty and the last query was a statement that does not return values (for example, an `INSERT` or a `DELETE`). If `mysql_field_count()` returns a non-zero value, the statement should have produced a non-empty result. See the description of the `mysql_field_count()` function for an example.

You can test for an error by calling `mysql_error()` or `mysql_errno()`.

### 17.2.13.2. What Results You Can Get from a Query

In addition to the result set returned by a query, you can also get the following information:

- `mysql_affected_rows()` returns the number of rows affected by the last query when doing an

[INSERT](#), [UPDATE](#), or [DELETE](#).

In MySQL 3.23, there is an exception when [DELETE](#) is used without a [WHERE](#) clause. In this case, the table is re-created as an empty table and [mysql\\_affected\\_rows\(\)](#) returns zero for the number of records affected. In MySQL 4.0 and later, [DELETE](#) always returns the correct number of rows deleted. For a fast re-create, use [TRUNCATE TABLE](#).

- [mysql\\_num\\_rows\(\)](#) returns the number of rows in a result set. With [mysql\\_store\\_result\(\)](#), [mysql\\_num\\_rows\(\)](#) may be called as soon as [mysql\\_store\\_result\(\)](#) returns. With [mysql\\_use\\_result\(\)](#), [mysql\\_num\\_rows\(\)](#) may be called only after you have fetched all the rows with [mysql\\_fetch\\_row\(\)](#).
- [mysql\\_insert\\_id\(\)](#) returns the ID generated by the last query that inserted a row into a table with an [AUTO\\_INCREMENT](#) index. See [Section 17.2.3.35](#), “[mysql\\_insert\\_id\(\)](#)”.
- Some queries ([LOAD DATA INFILE ...](#), [INSERT INTO ... SELECT ...](#), [UPDATE](#)) return additional information. The result is returned by [mysql\\_info\(\)](#). See the description for [mysql\\_info\(\)](#) for the format of the string that it returns. [mysql\\_info\(\)](#) returns a [NULL](#) pointer if there is no additional information.

### 17.2.13.3. How to Get the Unique ID for the Last Inserted Row

If you insert a record into a table that contains an [AUTO\\_INCREMENT](#) column, you can obtain the value stored into that column by calling the [mysql\\_insert\\_id\(\)](#) function.

You can check from your C applications whether a value was stored in an [AUTO\\_INCREMENT](#) column by executing the following code (which assumes that you've checked that the statement succeeded). It determines whether the query was an [INSERT](#) with an [AUTO\\_INCREMENT](#) index:

```
if ((result = mysql_store_result(&mysql)) == 0 &&
    mysql_field_count(&mysql) == 0 &&
    mysql_insert_id(&mysql) != 0)
{
    used_id = mysql_insert_id(&mysql);
}
```

For more information, see [Section 17.2.3.35](#), “[mysql\\_insert\\_id\(\)](#)”.

When a new [AUTO\\_INCREMENT](#) value has been generated, you can also obtain it by executing a [SELECT LAST\\_INSERT\\_ID\(\)](#) statement with [mysql\\_query\(\)](#) and retrieving the value from the result set returned by the statement.

For [LAST\\_INSERT\\_ID\(\)](#), the most recently generated ID is maintained in the server on a per-connection basis. It is not changed by another client. It is not even changed if you update another [AUTO\\_INCREMENT](#) column with a non-magic value (that is, a value that is not [NULL](#) and not 0).

If you want to use the ID that was generated for one table and insert it into a second table, you can use SQL statements like this:

```
INSERT INTO foo (auto,text)
VALUES(NULL,'text'); # generate ID by inserting NULL
INSERT INTO foo2 (id,text)
VALUES(LAST_INSERT_ID(),'text'); # use ID in second table
```

Note that [mysql\\_insert\\_id\(\)](#) returns the value stored into an [AUTO\\_INCREMENT](#) column, whether that value is automatically generated by storing [NULL](#) or 0 or was specified as an explicit value. [LAST\\_INSERT\\_ID\(\)](#) returns only automatically generated [AUTO\\_INCREMENT](#) values. If you store an explicit value other than [NULL](#) or 0, it does not affect the value returned by [LAST\\_INSERT\\_ID\(\)](#).

### 17.2.13.4. Problems Linking with the C API

When linking with the C API, the following errors may occur on some systems:

```
gcc -g -o client test.o -L/usr/local/lib/mysql -lmysqlclient -lsocket -lnsl
Undefined      first referenced
 symbol        in file
floor          /usr/local/lib/mysql/libmysqlclient.a(password.o)
ld: fatal: Symbol referencing errors. No output written to client
```

If this happens on your system, you must include the math library by adding `-lm` to the end of the compile/link line.

### 17.2.14. Building Client Programs

If you compile MySQL clients that you've written yourself or that you obtain from a third-party, they must be linked using the `-lmysqlclient -lz` options in the link command. You may also need to specify a `-L` option to tell the linker where to find the library. For example, if the library is installed in `/usr/local/mysql/lib`, use `-L/usr/local/mysql/lib -lmysqlclient -lz` in the link command.

For clients that use MySQL header files, you may need to specify an `-I` option when you compile them (for example, `-I/usr/local/mysql/include`), so that the compiler can find the header files.

To make it simpler to compile MySQL programs on Unix, we have provided the `mysql_config` script for you. See [Section 17.9.2, “mysql\\_config — Get Compile Options for Compiling Clients”](#).

You can use it to compile a MySQL client as follows:

```
CFG=/usr/local/mysql/bin/mysql_config
sh -c "gcc -o progname ` $CFG --cflags ` progname.c ` $CFG --libs `"
```

The `sh -c` is needed to get the shell not to treat the output from `mysql_config` as one word.

### 17.2.15. How to Make a Threaded Client

The client library is almost thread-safe. The biggest problem is that the subroutines in `net.c` that read from sockets are not interrupt safe. This was done with the thought that you might want to have your own alarm that can break a long read to a server. If you install interrupt handlers for the `SIGPIPE` interrupt, the socket handling should be thread-safe.

**Note:** Beginning with version 4.0.6, MySQL blocks `SIGPIPE` on the first call to `mysql_server_init()`, `mysql_init()`, or `mysql_connect()`. This is done to avoid aborting the program when a connection terminates. If you want to use your own `SIGPIPE` handler, you should first call `mysql_server_init()` and then install your handler.

In the older binaries we distribute on our Web site (<http://www.mysql.com/>), the client libraries are not normally compiled with the thread-safe option (the Windows binaries are by default compiled to be thread-safe). Newer binary distributions should have both a normal and a thread-safe client library.

To get a threaded client where you can interrupt the client from other threads and set timeouts when talking with the MySQL server, you should use the `-lmysys`, `-lmystrings`, and `-ldbug` libraries and the `net_serv.o` code that the server uses.

If you don't need interrupts or timeouts, you can just compile a thread-safe client library (`mysqlclient_r`) and use this. See [Section 17.2, “MySQL C API”](#). In this case, you don't have to worry about the `net_serv.o` object file or the other MySQL libraries.

When using a threaded client and you want to use timeouts and interrupts, you can make great use of the routines in the `thr_alarm.c` file. If you are using routines from the `mysys` library, the only thing you must remember is to call `my_init()` first! See [Section 17.2.11, “C API Threaded Function Descriptions”](#).

All functions except `mysql_real_connect()` are by default thread-safe. The following notes describe how to compile a thread-safe client library and use it in a thread-safe manner. (The notes below for `mysql_real_connect()` actually apply to `mysql_connect()` as well, but because `mysql_connect()` is deprecated, you should be using `mysql_real_connect()` anyway.)

To make `mysql_real_connect()` thread-safe, you must recompile the client library with this command:

```
shell> ./configure --enable-thread-safe-client
```

This creates a thread-safe client library `libmysqlclient_r`. (Assuming that your OS has a thread-safe `gethostbyname_r()` function.) This library is thread-safe per connection. You can let two threads share the same connection with the following caveats:

- Two threads can't send a query to the MySQL server at the same time on the same connection. In particular, you have to ensure that between a `mysql_query()` and `mysql_store_result()` no other thread is using the same connection.
- Many threads can access different result sets that are retrieved with `mysql_store_result()`.
- If you use `mysql_use_result`, you have to ensure that no other thread is using the same connection until the result set is closed. However, it really is best for threaded clients that share the same connection to use `mysql_store_result()`.
- If you want to use multiple threads on the same connection, you must have a mutex lock around your `mysql_query()` and `mysql_store_result()` call combination. Once `mysql_store_result()` is ready, the lock can be released and other threads may query the same connection.
- If you program with POSIX threads, you can use `pthread_mutex_lock()` and `pthread_mutex_unlock()` to establish and release a mutex lock.

You need to know the following if you have a thread that is calling MySQL functions which did not create the connection to the MySQL database:

When you call `mysql_init()` or `mysql_connect()`, MySQL creates a thread-specific variable for the thread that is used by the debug library (among other things).

If you call a MySQL function, before the thread has called `mysql_init()` or `mysql_connect()`, the thread does not have the necessary thread-specific variables in place and you are likely to end up with a core dump sooner or later.

To get things to work smoothly you have to do the following:

1. Call `my_init()` at the start of your program if it calls any other MySQL function before calling `mysql_real_connect()`.
2. Call `mysql_thread_init()` in the thread handler before calling any MySQL function.
3. In the thread, call `mysql_thread_end()` before calling `pthread_exit()`. This frees the memory used by MySQL thread-specific variables.

You may get some errors because of undefined symbols when linking your client with `libmysqlclient_r`. In most cases this is because you haven't included the thread libraries on the link/compile line.

## 17.3. MySQL PHP API

PHP is a server-side, HTML-embedded scripting language that may be used to create dynamic Web pages. It is available for most operating systems and Web servers, and can access most common databases, including MySQL. PHP may be run as a separate program or compiled as a module for use with the Apache Web server.

PHP actually provides two different MySQL API extensions:

- `mysql`: Available for PHP versions 4 and 5, this extension is intended for use with MySQL versions prior to MySQL 4.1. This extension does not support the improved authentication protocol used in MySQL 5.0, nor does it support prepared statements or multiple statements. If you wish to use this extension with MySQL 4.1, you will likely want to configure the MySQL server to use the `--old-passwords` option (see [Section A.2.3, “Client does not support authentication protocol”](#)). This extension is documented on the PHP Web site at <http://php.net/mysql>.
- `mysqli` - Stands for “MySQL, Improved”; this extension is available only in PHP 5. It is intended for use with MySQL 4.1.1 and later. This extension fully supports the enhanced authentication protocol used beginning with MySQL 4.1.1, as well as the Prepared Statements and Multiple Statements APIs. In addition, this extension provides an advanced, object-oriented programming interface. You can read the documentation for the `mysqli` extension at <http://php.net/mysqli>. A helpful article can be found at <http://www.zend.com/php5/articles/php5-mysqli.php>.

If you're experiencing problems with enabling both the `mysql` and the `mysqli` extension when building PHP on Linux yourself, see [Section 17.3.2, “Enabling Both `mysql` and `mysqli` in PHP”](#).

The PHP distribution and documentation are available from the [PHP Web site](http://www.php.net/) [http://www.php.net/]. MySQL provides the `mysql` and `mysqli` extensions for the Windows operating system for MySQL versions as of 4.1.16 on <http://dev.mysql.com/downloads/connector/php/>. You can find information why you should preferably use the extensions provided by MySQL on that page.

### 17.3.1. Common Problems with MySQL and PHP

- `Error: Maximum Execution Time Exceeded`: This is a PHP limit; go into the `php.ini` file and set the maximum execution time up from 30 seconds to something higher, as needed. It is also not a bad idea to double the RAM allowed per script to 16MB instead of 8MB.
- `Fatal error: Call to unsupported or undefined function mysql_connect() in ...`: This means that your PHP version isn't compiled with MySQL support. You can either compile a dynamic MySQL module and load it into PHP or recompile PHP with built-in MySQL support. This process is described in detail in the PHP manual.
- `Error: Undefined reference to 'uncompress'`: This means that the client library is compiled with support for a compressed client/server protocol. The fix is to add `-lz` last when linking with `-libmysqlclient`.
- `Error: Client does not support authentication protocol`: This is most often encountered when trying to use the older `mysql` extension with MySQL 4.1.1 and later. Possible solutions are: downgrade to MySQL 4.0; switch to PHP 5 and the newer `mysqli` extension; or configure the MySQL server with `--old-passwords`. (See [Section A.2.3, “Client does not support authentication protocol”](#), for more information.)

Those with PHP4 legacy code can make use of a compatibility layer for the old and new MySQL libraries, such as this one: <http://www.coggeshall.org/oss/mysql2i>.

## 17.3.2. Enabling Both `mysql` and `mysqli` in PHP

If you're experiencing problems with enabling both the `mysql` and the `mysqli` extension when building PHP on Linux yourself, you should try the following procedure.

1. Configure PHP like this:

```
./configure --with-mysqli=/usr/bin/mysql_config --with-mysql=/usr
```

2. Edit the `Makefile` and search for a line that starts with `EXTRA_LIBS`. It might look like this (all on one line):

```
EXTRA_LIBS = -lcrypt -lcrypt -lmysqlclient -lz -lresolv -lm -ldl -lnsl
-lxml2 -lz -lm -lxml2 -lz -lm -lmysqlclient -lz -lcrypt -lnsl -lm
-lxml2 -lz -lm -lcrypt -lxml2 -lz -lm -lcrypt
```

Remove all duplicates, so that the line looks like this (all on one line):

```
EXTRA_LIBS = -lcrypt -lcrypt -lmysqlclient -lz -lresolv -lm -ldl -lnsl
-lxml2
```

3. Build and install PHP:

```
make
make install
```

## 17.4. MySQL Perl API

The Perl `DBI` module provides a generic interface for database access. You can write a DBI script that works with many different database engines without change. To use DBI, you must install the `DBI` module, as well as a DataBase Driver (DBD) module for each type of server you want to access. For MySQL, this driver is the `DBD::mysql` module.

Perl DBI is the recommended Perl interface. It replaces an older interface called `mysqlperl`, which should be considered obsolete.

Installation instructions for Perl DBI support are given in [Section 2.14, "Perl Installation Notes"](#).

DBI information is available at the command line, online, or in printed form:

- Once you have the `DBI` and `DBD::mysql` modules installed, you can get information about them at the command line with the `perldoc` command:

```
shell> perldoc DBI
shell> perldoc DBI::FAQ
shell> perldoc DBD::mysql
```

You can also use `pod2man`, `pod2html`, and so forth to translate this information into other formats.

- For online information about Perl DBI, visit the DBI Web site, <http://dbi.perl.org/>. That site hosts a general DBI mailing list. MySQL AB hosts a list specifically about `DBD::mysql`; see [Sec-](#)

tion 1.7.1, “MySQL Mailing Lists”.

- For printed information, the official DBI book is *Programming the Perl DBI* (Alligator Descartes and Tim Bunce, O'Reilly & Associates, 2000). Information about the book is available at the DBI Web site, <http://dbi.perl.org/>.

For information that focuses specifically on using DBI with MySQL, see *MySQL and Perl for the Web* (Paul DuBois, New Riders, 2001). This book's Web site is <http://www.kitebird.com/mysql-perl/>.

## 17.5. MySQL C++ API

`MySQL++` is a MySQL API for C++. Warren Young has taken over this project. More information can be found at <http://www.mysql.com/products/mysql++/>.

### 17.5.1. Borland C++

You can compile the MySQL Windows source with Borland C++ 5.02. (The Windows source includes only projects for Microsoft VC++, for Borland C++ you have to do the project files yourself.)

One known problem with Borland C++ is that it uses a different structure alignment than VC++. This means that you run into problems if you try to use the default `libmysql.dll` libraries (that were compiled using VC++) with Borland C++. To avoid this problem, only call `mysql_init()` with `NULL` as an argument, not a pre-allocated `MYSQL` structure.

## 17.6. MySQL Python API

`MySQLdb` provides MySQL support for Python, compliant with the Python DB API version 2.0. It can be found at <http://sourceforge.net/projects/mysql-python/>.

## 17.7. MySQL Tcl API

`MySQLtcl` is a simple API for accessing a MySQL database server from the Tcl programming language. It can be found at <http://www.xdobry.de/mysqltcl/>.

## 17.8. MySQL Eiffel Wrapper

Eiffel MySQL is an interface to the MySQL database server using the Eiffel programming language, written by Michael Ravits. It can be found at <http://efsa.sourceforge.net/archive/ravits/mysql.htm>.

## 17.9. MySQL Program Development Utilities

This section describes some utilities that you may find useful when developing MySQL programs.

- `msql2mysql`

A shell script that converts `mSQL` programs to MySQL. It doesn't handle every case, but it gives a good start when converting.

- `mysql_config`

A shell script that produces the option values needed when compiling MySQL programs.

## 17.9.1. `msql2mysql` — Convert mSQL Programs for Use with MySQL

Initially, the MySQL C API was developed to be very similar to that for the mSQL database system. Because of this, mSQL programs often can be converted relatively easily for use with MySQL by changing the names of the C API functions.

The `msql2mysql` utility performs the conversion of mSQL C API function calls to their MySQL equivalents. `msql2mysql` converts the input file in place, so make a copy of the original before converting it. For example, use `msql2mysql` like this:

```
shell> cp client-prog.c client-prog.c.orig
shell> msql2mysql client-prog.c
client-prog.c converted
```

Then examine `client-prog.c` and make any post-conversion revisions that may be necessary.

`msql2mysql` uses the `replace` utility to make the function name substitutions. See [Section 8.18](#), “`replace` — A String-Replacement Utility”.

## 17.9.2. `mysql_config` — Get Compile Options for Compiling Clients

`mysql_config` provides you with useful information for compiling your MySQL client and connecting it to MySQL.

`mysql_config` supports the following options:

- `--cflags`

Compiler flags to find include files and critical compiler flags and defines used when compiling the `libmysqlclient` library.

- `--include`

Compiler options to find MySQL include files. (Note that normally you would use `--cflags` instead of this option.)

- `--libmysqld-libs, --embedded`

Libraries and options required to link with the MySQL embedded server.

- `--libs`

Libraries and options required to link with the MySQL client library.

- `--libs_r`

Libraries and options required to link with the thread-safe MySQL client library.

- `--port`

The default TCP/IP port number, defined when configuring MySQL.

- `--socket`

The default Unix socket file, defined when configuring MySQL.



- `--version`

Version number for the MySQL distribution.

If you invoke `mysql_config` with no options, it displays a list of all options that it supports, and their values:

```
shell> mysql_config
Usage: /usr/local/mysql/bin/mysql_config [options]
Options:
--cflags      [-I/usr/local/mysql/include/mysql -mcpu=pentiumpro]
--include    [-I/usr/local/mysql/include/mysql]
--libs       [-L/usr/local/mysql/lib/mysql -lmysqlclient -lz
             -lcrypt -lnsl -lm -L/usr/lib -lssl -lcrypto]
--libs_r     [-L/usr/local/mysql/lib/mysql -lmysqlclient_r
             -lpthread -lz -lcrypt -lnsl -lm -lpthread]
--socket     [/tmp/mysql.sock]
--port       [3306]
--version    [4.0.16]
--libmysqld-libs [-L/usr/local/mysql/lib/mysql -lmysqld -lpthread -lz
                -lcrypt -lnsl -lm -lpthread -lrt]
```

You can use `mysql_config` within a command line to include the value that it displays for a particular option. For example, to compile a MySQL client program, use `mysql_config` as follows:

```
shell> CFG=/usr/local/mysql/bin/mysql_config
shell> sh -c "gcc -o progname ` $CFG --cflags ` progname.c ` $CFG --libs `"
```

When you use `mysql_config` this way, be sure to invoke it within backtick (``) characters. That tells the shell to execute it and substitute its output into the surrounding command.

---

# Chapter 18. Connectors

This chapter describes MySQL Connectors, drivers that provide connectivity to the MySQL server for client programs. There are currently five MySQL Connectors:

- Connector/ODBC provides driver support for connecting to a MySQL server using the Open Database Connectivity (ODBC) API. Support is available for ODBC connectivity from Windows, Unix and Mac OS X platforms.
- Connector/NET enables developers to create .NET applications that use data stored in a MySQL database. Connector/NET implement a fully-functional ADO.NET interface and provides support for use with ADO.NET aware tools. Applications that want to use Connector/NET can be written in any of the supported .NET languages.
- Connector/J provides driver support for connecting to MySQL from a Java application using the standard Java Database Connectivity (JDBC) API.
- Connector/MXJ is a tool that enables easy deployment and management of MySQL server and database through your Java application.
- Connector/PHP is a Windows-only connector for PHP that provides the `mysql` and `mysqli` extensions for use with MySQL 5.0.18 and later.

For information on connecting to a MySQL server using other languages and interfaces than those detailed above, including Perl, Python and PHP for other platforms and environments, please refer to the [Chapter 17, APIs and Libraries](#) chapter.

## 18.1. MySQL Connector/ODBC

The MySQL Connector/ODBC is the name for the family of MySQL ODBC drivers (also called MyODBC drivers) that provide access to a MySQL database using the industry standard Open Database Connectivity (ODBC) API. This reference covers Connector/ODBC 3.51, a version of the API that provides ODBC 3.5x compliant access to a MySQL database.

The manual for versions of MyODBC older than 3.51 can be located in the corresponding binary or source distribution.

For more information on the ODBC API standard and how to use it, refer to <http://www.microsoft.com/data/>.

The application development part of this reference assumes a good working knowledge of C, general DBMS knowledge, and finally, but not least, familiarity with MySQL. For more information about MySQL functionality and its syntax, refer to <http://dev.mysql.com/doc/>.

Typically, you need to install MyODBC only on Windows machines. For Unix and Mac OS X you can use the native MySQL network or named pipe to communicate with your MySQL database. You may need MyODBC for Unix or Mac OS X if you have an application that requires an ODBC interface to communicate with database.. Applications that require ODBC to communicate with MySQL include ColdFusion, Microsoft Office, and Filemaker Pro.

If you want to install the MyODBC connector on a Unix host, then you must also install an ODBC manager.

If you have questions that are not answered in this document, please send a mail message to [<myodbc@lists.mysql.com>](mailto:myodbc@lists.mysql.com).

## 18.1.1. Introduction to MyODBC

ODBC (Open Database Connectivity) provides a way for client programs to access a wide range of databases or data sources. ODBC is a standardized API that allows connections to SQL database servers. It was developed according to the specifications of the SQL Access Group and defines a set of function calls, error codes, and data types that can be used to develop database-independent applications. ODBC usually is used when database independence or simultaneous access to different data sources is required.

For more information about ODBC, refer to <http://www.microsoft.com/data/>.

### 18.1.1.1. MyODBC Versions

There are currently two version of MyODBC available:

- MyODBC 5.0, currently in beta status, has been designed to extend the functionality of the MyODBC 3.51 driver and incorporate full support for the functionality in the MySQL 5.0 server release, including stored procedures and views. Applications using MyODBC 3.51 will be compatible with MyODBC 5.0, while being able to take advantage of the new features. Features and functionality of the MyODBC 5.0 driver are not currently included in this guide.
- MyODBC 3.51 is the current release of the 32-bit ODBC driver, also known as the MySQL ODBC 3.51 driver. This version is enhanced compared to the older MyODBC 2.50 driver. It has support for ODBC 3.5x specification level 1 (complete core API + level 2 features) in order to continue to provide all functionality of ODBC for accessing MySQL.
- MyODBC 2.50 is the previous version of the 32-bit ODBC driver from MySQL AB that is based on ODBC 2.50 specification level 0 (with level 1 and 2 features). Information about the MyODBC 2.50 driver is included in this guide for the purposes of comparison only.

Note: From this section onward, the primary focus of this guide is the MyODBC 3.51 driver. More information about the MyODBC 2.50 driver in the documentation included in the installation packages for that version. If there is a specific issue (error or known problem) that only affects the 2.50 version, it may be included here for reference.

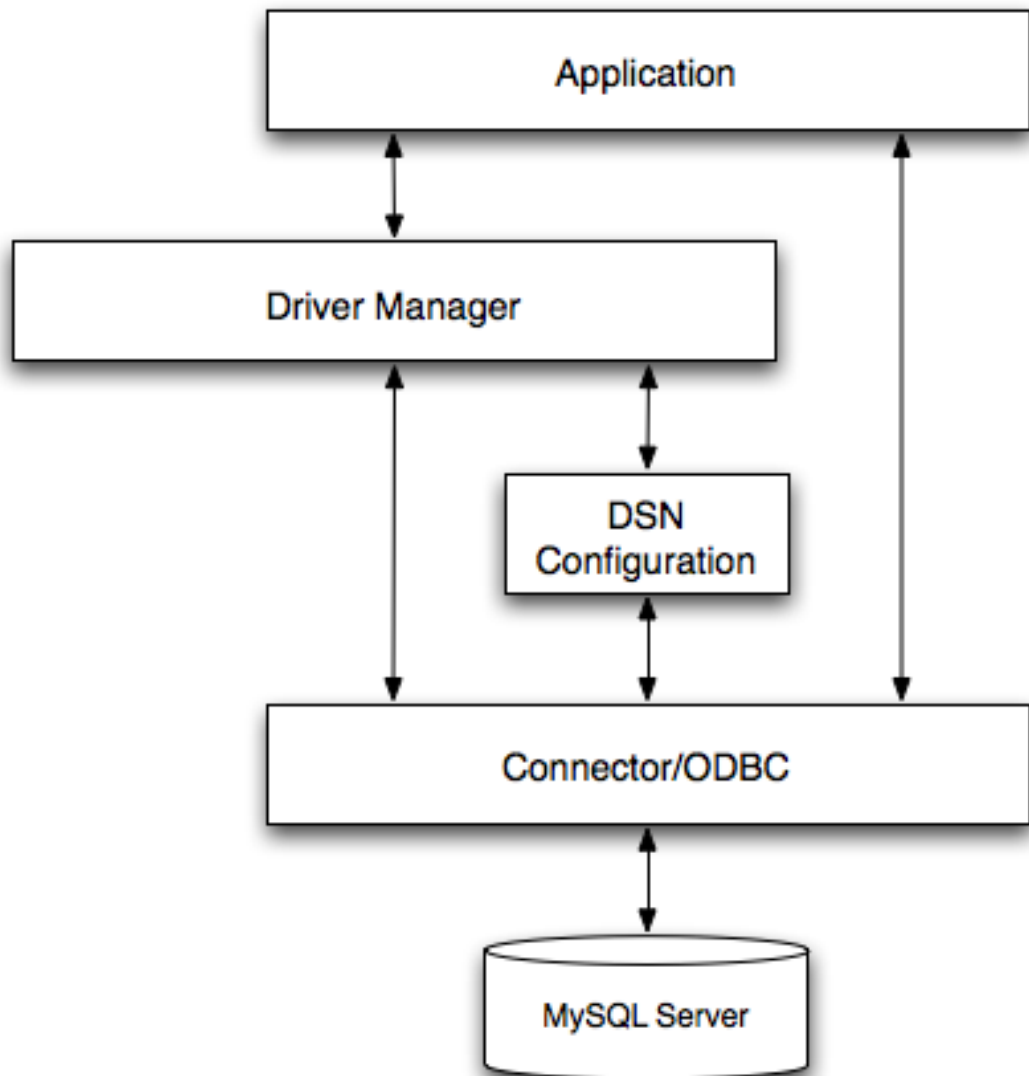
### 18.1.1.2. General Information About ODBC and MyODBC

Open Database Connectivity (ODBC) is a widely accepted application-programming interface (API) for database access. It is based on the Call-Level Interface (CLI) specifications from X/Open and ISO/IEC for database APIs and uses Structured Query Language (SQL) as its database access language.

A survey of ODBC functions supported by MyODBC is given at [Section 18.1.5.1, “MyODBC API Reference”](#). For general information about ODBC, see <http://www.microsoft.com/data/>.

#### 18.1.1.2.1. MyODBC Architecture

The MyODBC architecture is based on five components, as shown in the following diagram:



- **Application:**

The Application uses the ODBC API to access the data from the MySQL server. The ODBC API in turn uses the communicates with the Driver Manager. The Application communicates with the Driver Manager using the standard ODBC calls. The Application does not care where the data is stored, how it is stored, or even how the system is configured to access the data. It needs to know only the Data Source Name (DSN).

A number of tasks are common to all applications, no matter how they use ODBC. These tasks are:

- Selecting the MySQL server and connecting to it
- Submitting SQL statements for execution
- Retrieving results (if any)
- Processing errors

- Committing or rolling back the transaction enclosing the SQL statement
- Disconnecting from the MySQL server

Because most data access work is done with SQL, the primary tasks for applications that use ODBC are submitting SQL statements and retrieving any results generated by those statements.

- **Driver manager:**

The Driver Manager is a library that manages communication between application and driver or drivers. It performs the following tasks:

- Resolves Data Source Names (DSN). The DSN is a configuration string that identifies a given database driver, database, database host and optionally authentication information that enables an ODBC application to connect to a database using a standardized reference.

Because the database connectivity information is identified by the DSN, any ODBC compliant application can connect to the data source using the same DSN reference. This eliminates the need to separately configure each application that needs access to a given database; instead you instruct the application to use a pre-configured DSN.

- Loading and unloading of the driver required to access a specific database as defined within the DSN. For example, if you have configured a DSN that connects to a MySQL database then the driver manager will load the MyODBC driver to enable the ODBC API to communicate with the MySQL host.
- Processes ODBC function calls or passes them to the driver for processing.

- **MyODBC Driver:**

The MyODBC driver is a library that implements the functions supported by the ODBC API. It processes ODBC function calls, submits SQL requests to MySQL server, and returns results back to the application. If necessary, the driver modifies an application's request so that the request conforms to syntax supported by MySQL.

- **DSN Configuration:**

The ODBC configuration file stores the driver and database information required to connect to the server. It is used by the Driver Manager to determine which driver to be loaded according to the definition in the DSN. The driver uses this to read connection parameters based on the DSN specified. For more information, [Section 18.1.3, “MyODBC Configuration”](#).

- **MySQL Server:**

The MySQL database where the information is stored. The database is used as the source of the data (during queries) and the destination for data (during inserts and updates).

### 18.1.1.2.2. ODBC Driver Managers

An ODBC Driver Manager is a library that manages communication between the ODBC-aware application and any drivers. Its main functionality includes:

- Resolving Data Source Names (DSN).
- Driver loading and unloading.

- Processing ODBC function calls or passing them to the driver.

Both Windows and Mac OS X include ODBC driver managers with the operating system. Most ODBC Driver Manager implementations also include an administration application that makes the configuration of DSN and drivers easier. Examples and information on these managers, including Unix ODBC driver managers are listed below:

- Microsoft Windows ODBC Driver Manager (`odbc32.dll`), <http://www.microsoft.com/data/>.
- Mac OS X includes `ODBC Administrator`, a GUI application that provides a simpler configuration mechanism for the Unix iODBC Driver Manager. You can configure DSN and driver information either through ODBC Administrator or through the iODBC configuration files. This also means that you can test ODBC Administrator configurations using the `iodbctest` command. <http://www.apple.com>.
- `unixODBC` Driver Manager for Unix (`libodbc.so`). See <http://www.unixodbc.org>, for more information. The `unixODBC` Driver Manager includes the MyODBC driver 3.51 in the installation package, starting with version `unixODBC 2.1.2`.
- `iODBC` ODBC Driver Manager for Unix (`libiodbc.so`), see <http://www.iodbc.org>, for more information.

## 18.1.2. How to Install MyODBC

You can install the MyODBC drivers using two different methods, a binary installation and a source installation. The binary installation is the easiest and most straightforward method of installation. Using the source installation methods should only be necessary on platforms where a binary installation package is not available, or in situations where you want to customize or modify the installation process or MyODBC drivers before installation.

### 18.1.2.1. Where to Get MyODBC

MySQL AB distributes all its products under the General Public License (GPL). You can get a copy of the latest version of MyODBC binaries and sources from the MySQL AB Web site <http://dev.mysql.com/downloads/>.

For more information about MyODBC, visit <http://www.mysql.com/products/myodbc/>.

For more information about licensing, visit <http://www.mysql.com/company/legal/licensing/>.

### 18.1.2.2. Supported Platforms

MyODBC can be used on all major platforms supported by MySQL. You can install it on:

- Windows 95, 98, Me, NT, 2000, XP, and 2003
- All Unix-like Operating Systems, including: AIX, Amiga, BSDI, DEC, FreeBSD, HP-UX 10/11, Linux, NetBSD, OpenBSD, OS/2, SGI Irix, Solaris, SunOS, SCO OpenServer, SCO UnixWare, Tru64 Unix
- Mac OS X and Mac OS X Server

If a binary distribution is not available for a particular platform, see [Section 18.1.2.4, “Installing MyOD-](#)

[BC from a source distribution](#)", to build the driver from the original source code. You can contribute the binaries you create to MySQL by sending a mail message to [<myodbc@lists.mysql.com>](mailto:myodbc@lists.mysql.com), so that it becomes available for other users.

### 18.1.2.3. Installing MyODBC from a binary distribution

Using a binary distribution offers the most straightforward method for installing MyODBC. If you want more control over the driver, the installation location and or to customize elements of the driver you will need to build and install from the source. See the [Section 18.1.2.4, "Installing MyODBC from a source distribution"](#).

#### 18.1.2.3.1. Installing MyODBC from a Binary Distribution on Windows

Before installing the MyODBC drivers on Windows you should ensure that your Microsoft Data Access Components (MDAC) are up to date. You can obtain the latest version from the [Microsoft Data Access and Storage](http://www.microsoft.com/data/) [http://www.microsoft.com/data/] website.

There are three available distribution types to use when installing for Windows. The contents in each case are identical, it is only the installation method which is different.


- Zipped installer consists of a Zipped package containing a standalone installation application. To install from this package, you must unzip the installer, and then run the installation application. See [Section 18.1.2.3.1.1, "Installing the Windows MyODBC Driver using an installer"](#) to complete the installation.
- MSI installer, an installation file that can be used with the installer included in Windows 2000, Windows XP and Windows Server 2003. See [Section 18.1.2.3.1.1, "Installing the Windows MyODBC Driver using an installer"](#) to complete the installation.
- Zipped DLL package, containing the DLL files that need must be manually installed. See [Section 18.1.2.3.1.2, "Installing the Windows MyODBC Driver using the Zipped DLL package"](#) to complete the installation.

##### 18.1.2.3.1.1. Installing the Windows MyODBC Driver using an installer

The installer packages offer a very simple method for installing the MyODBC drivers. If you have downloaded the zipped installer then you must extract the installer application. The basic installation process is identical for both installers.

You should follow these steps to complete the installation:

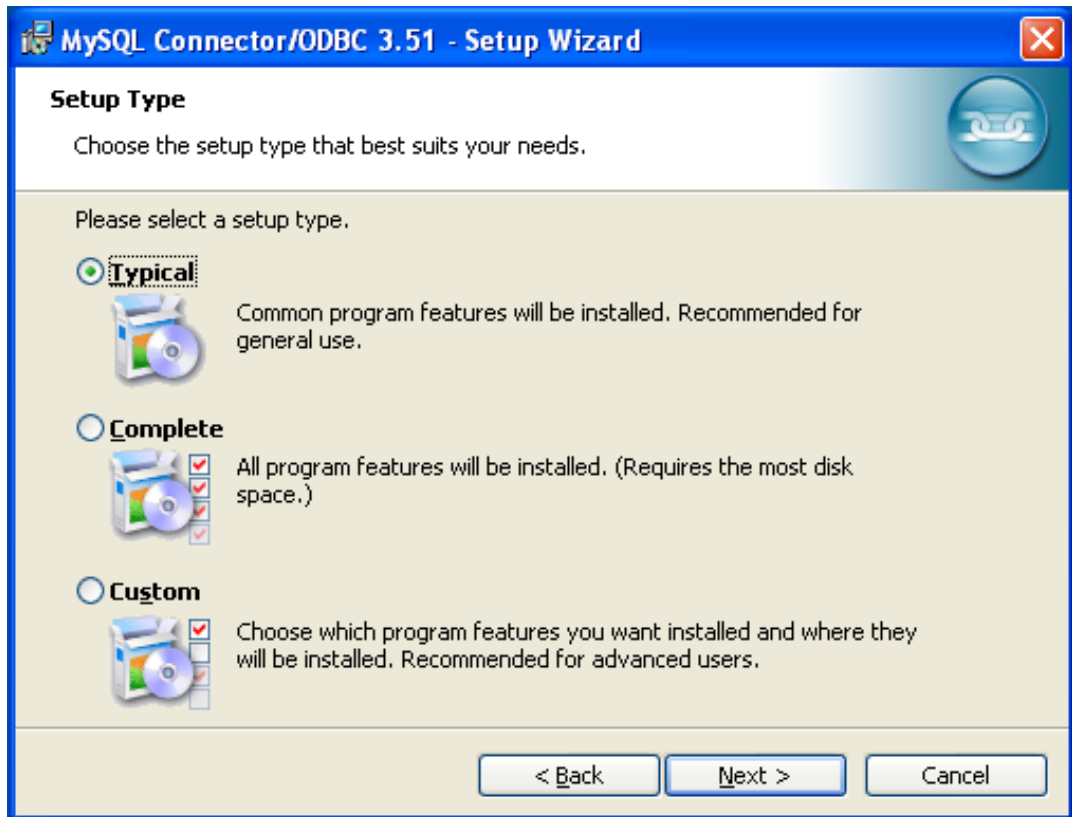
1. Double click on the standalone installer that you extracted, or the MSI file you downloaded.
2. The MySQL Connector/ODBC 3.51 - Setup Wizard will start. Click the [Next](#) button to begin the installation process.



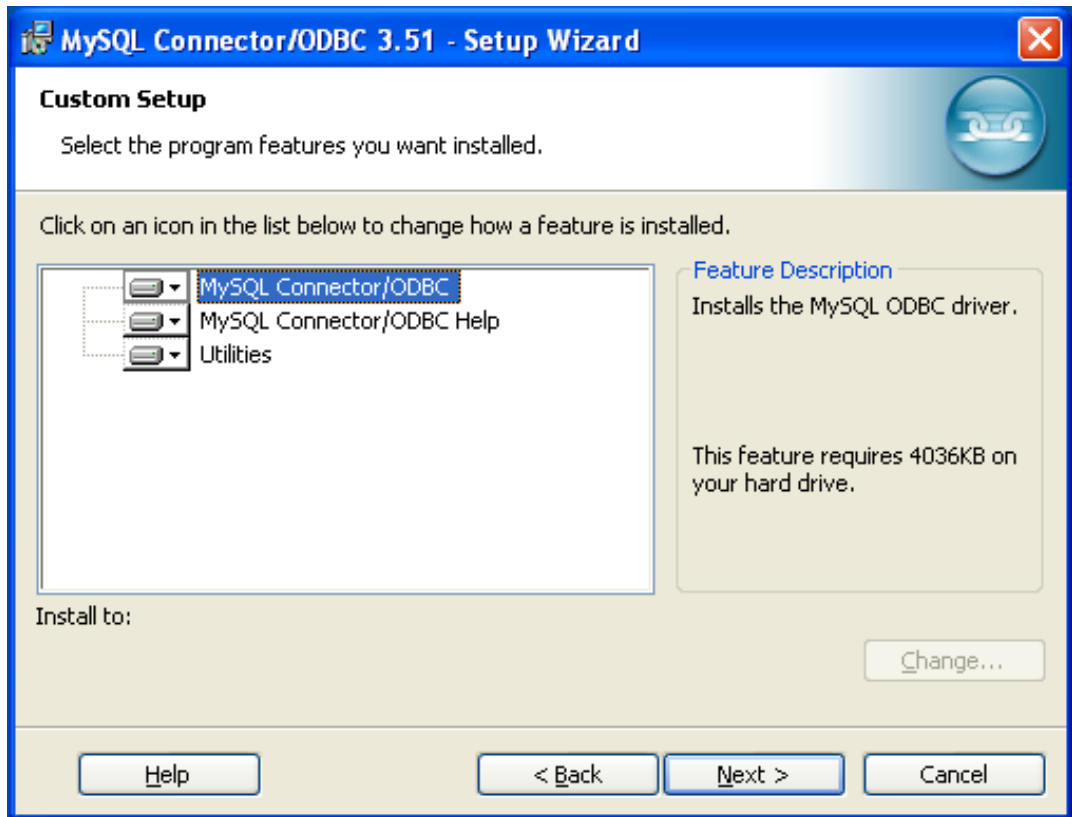
3. You will need to choose the installation type. The Typical installation provides the standard files you will need to connect to a MySQL database using ODBC. The Complete option installs all the available files, including debug and utility components. It is recommended you choose one of these two options to complete the installation. If choose one of these methods, click [Next](#) and then proceed to step 5.

You may also choose a Custom installation, which enables you to select the individual components that you want to install. You have chosen this method, click [Next](#) and then proceed to step 4.






4. If you have chosen a custom installation, use the popups to select which components to install and then click [Next](#) to install the necessary files.



5. Once the files have copied to your machine, the installation is complete. Click **Finish** to exit the installer.



Now the installation is complete, you can continue to configure your ODBC connections using [Section 18.1.3, “MyODBC Configuration”](#).

#### 18.1.2.3.1.2. Installing the Windows MyODBC Driver using the Zipped DLL package

If you have downloaded the Zipped DLL package then you must install the individual files required for MyODBC operation manually. Once you have unzipped the installation files, you can either perform this operation by hand, executing each statement individually, or you can use the included Batch file to perform an installation to the default locations.

To install using the Batch file:

1. Unzip the MyODBC Zipped DLL package.
2. Open a Command Prompt.
3. Change to the directory created when you unzipped the MyODBC Zipped DLL package.
4. Run `Install.bat`:

```
C:\> Install.bat
```

This will copy the necessary files into the default location, and then register the MyODBC driver with the Windows ODBC manager.

If you want to copy the files to an alternative location - for example, to run or test different versions of

the MyODBC driver on the same machine, then you must copy the files by hand. It is however not recommended to install these files in a non-standard location. To copy the files by hand to the default installation location use the following steps:

1. Unzip the MyODBC Zipped DLL package.
2. Open a Command Prompt.
3. Change to the directory created when you unzipped the MyODBC Zipped DLL package.
4. Copy the library files to a suitable directory. The default is to copy them into the default Windows system directory `\Windows\System32`:

```
C:\> copy lib\myodbc3s.dll \Windows\System32
C:\> copy lib\myodbc3s.lib \Windows\System32
C:\> copy lib\myodbc3.dll \Windows\System32
C:\> copy lib\myodbc3.lib \Windows\System32
```

5. Copy the MyODBC tools. These must be placed into a directory that is in the system `PATH`. The default is to install these into the Windows system directory `\Windows\System32`:

```
C:\> copy bin\myodbc3i.exe \Windows\System32
C:\> copy bin\myodbc3m.exe \Windows\System32
C:\> copy bin\myodbc3c.exe \Windows\System32
```

6. Optionally copy the help files. For these files to be accessible through the help system, they must be installed in the Windows system directory:

```
C:\> copy doc\*.hlp \Windows\System32
```

7. Finally, you must register the MyODBC driver with the ODBC manager:

```
C:\> myodbc3i -a -d -t"MySQL ODBC 3.51 Driver;\
DRIVER=myodbc3.dll;SETUP=myodbc3s.dll"
```

You must change the references to the DLL files and command location in the above statement if you have not installed these files into the default location.

### 18.1.2.3.1.3. Handling Installation Errors

On Windows, you may get the following error when trying to install the older MyODBC 2.50 driver:

```
An error occurred while copying C:\WINDOWS\SYSTEM\MFC30.DLL.
Restart Windows and try installing again (before running any
applications which use ODBC)
```

The reason for the error is that another application is currently using the ODBC system. Windows may not allow you to complete the installation. In most cases, you can continue by pressing `Ignore` to copy the rest of the MyODBC files and the final installation should still work. If it doesn't, the solution is to re-boot your computer in "safe mode." Choose safe mode by pressing `F8` just before your machine starts Windows during re-booting, install the MyODBC drivers, and re-boot to normal mode.

### 18.1.2.3.2. Installing MyODBC from a Binary Distribution on Unix

There are two methods available for installing MyODBC on Unix from a binary distribution. For most Unix environments you will need to use the tarball distribution. For Linux systems, there is also an RPM distribution available.

### 18.1.2.3.2.1. Installing MyODBC from a Binary Tarball Distribution

To install the driver from a tarball distribution (`.tar.gz` file), download the latest version of the driver for your operating system and follow these steps that demonstrate the process using the Linux version of the tarball:

```
shell> su root
shell> gunzip MyODBC-3.51.11-i686-pc-linux.tar.gz
shell> tar xvf MyODBC-3.51.11-i686-pc-linux.tar
shell> cd MyODBC-3.51.11-i686-pc-linux
```

Read the installation instructions in the `INSTALL-BINARY` file and execute these commands.

```
shell> cp libmyodbc* /usr/local/lib
shell> cp odbc.ini /usr/local/etc
shell> export ODBCINI=/usr/local/etc/odbc.ini
```

Then proceed on to [Section 18.1.3.4, “Configuring a MyODBC DSN on Unix”](#), to configure the DSN for MyODBC. For more information, refer to the `INSTALL-BINARY` file that comes with your distribution.

### 18.1.2.3.2.2. Installing MyODBC from an RPM Distribution

To install or upgrade MyODBC from an RPM distribution on Linux, simply download the RPM distribution of the latest version of MyODBC and follow the instructions below. Use `su root` to become `root`, then install the RPM file.

If you are installing for the first time:

```
shell> su root
shell> rpm -ivh MyODBC-3.51.12.i386.rpm
```

If the driver exists, upgrade it like this:

```
shell> su root
shell> rpm -Uvh MyODBC-3.51.12.i386.rpm
```

If there is any dependency error for MySQL client library, `libmysqlclient`, simply ignore it by supplying the `--nodeps` option, and then make sure the MySQL client shared library is in the path or set through `LD_LIBRARY_PATH`.

This installs the driver libraries and related documents to `/usr/local/lib` and `/usr/share/doc/MyODBC`, respectively. Proceed onto [Section 18.1.3.4, “Configuring a MyODBC DSN on Unix”](#).

To **uninstall** the driver, become `root` and execute an `rpm` command:

```
shell> su root
shell> rpm -e MyODBC
```

### 18.1.2.3.3. Installing MyODBC on Mac OS X


Mac OS X is based on the FreeBSD operating system, and you can normally use the MySQL network port for connecting to MySQL servers on other hosts. Installing the MyODBC driver enables you to connect to MySQL databases on any platform through the ODBC interface. You should only need to install the MyODBC driver when your application requires an ODBC interface. Applications that require or can use ODBC (and therefore the MyODBC driver) include ColdFusion, Filemaker Pro, 4th Dimension and many other applications.

Mac OS X includes its own ODBC manager, based on the [iODBC](#) manager. Mac OS X includes an administration tool that provides easier administration of ODBC drivers and configuration, updating the underlying [iODBC](#) configuration files.


#### 18.1.2.3.3.1. Installing the MyODBC Driver

You can install MyODBC on a Mac OS X or Mac OS X Server computer by using the binary distribution. The package is available as a compressed disk image ([.dmg](#)) file. To install MyODBC on your computer using this method, follow these steps:

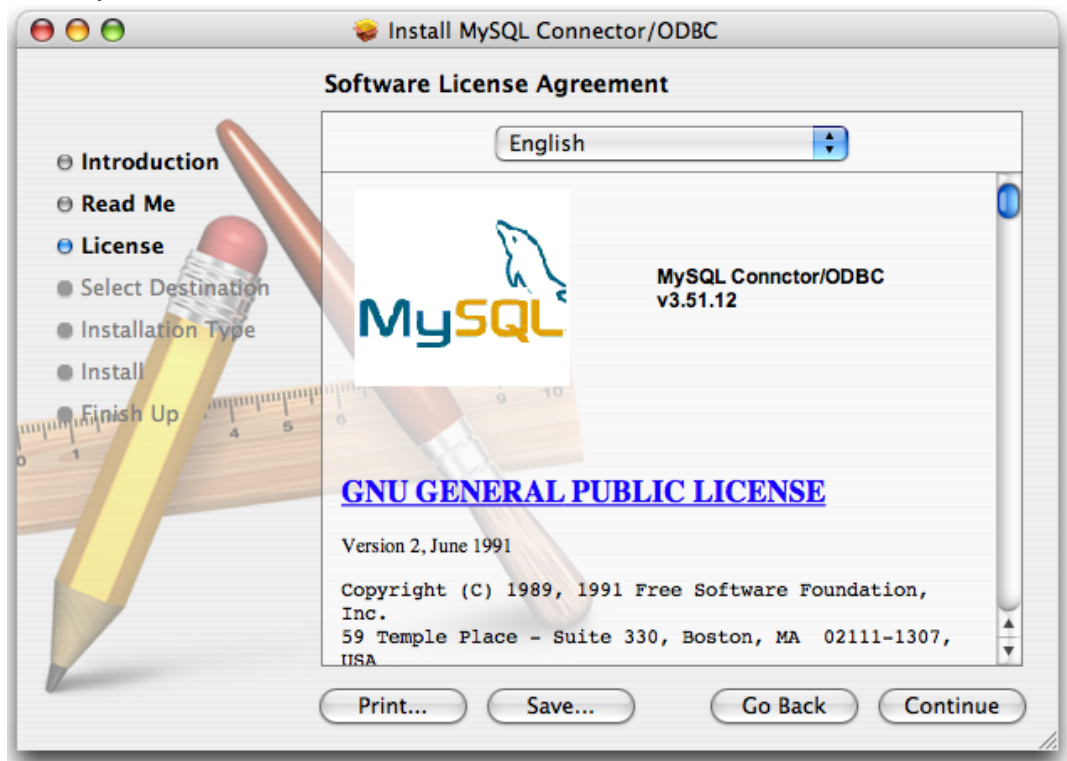
1. Download the file to your computer and double-click on the downloaded image file.
2. Within the disk image you will find an installer package (with the [.pkg](#) extension). Double click on this file to start the Mac OS X installer.
3. You will be presented with the installer welcome message. Click the [Continue](#) button to begin the installation process.




4. Please take the time to read the Important Information as it contains guidance on how to complete the installation process. Once you have read the notice and collected the necessary information, click [Continue](#).



5. MyODBC drivers are made available under the GNU General Public License. Please read the license if you are not familiar with it before continuing installation. Click [Continue](#) to approve the license (you will be asked to confirm that decision) and continue the installation.

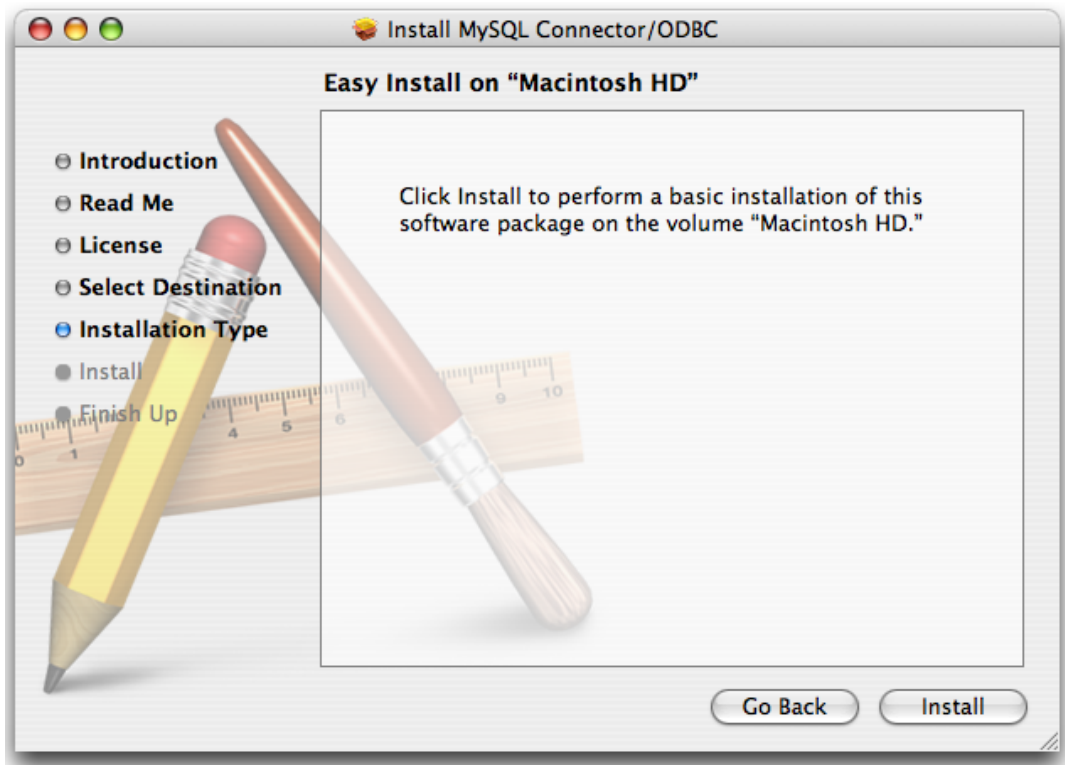


6. Choose a location to install the MyODBC drivers and the ODBC Administrator application. You must install the files onto a drive with an operating system and you may be limited in the choices available. Select the drive you want to use, and then click [Continue](#).

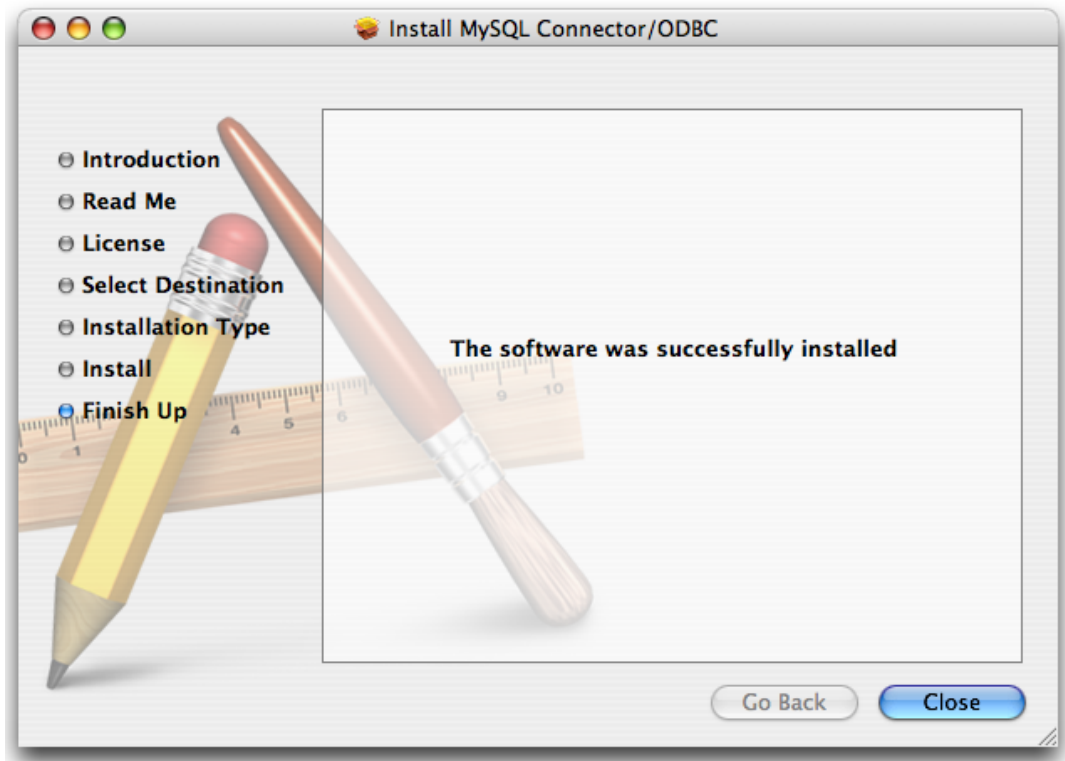


7. The installer will automatically select the files that need to be installed on your machine. Click [Install](#) to continue. The installer will copy the necessary files to your machine. A progress bar will be shown indicating the installation progress.





8. When installation has been completed you will get a window like the one shown below. Click **Close** to close and quit the installer.



## 18.1.2.4. Installing MyODBC from a source distribution

Installing MyODBC from a source distribution gives you greater flexibility in the contents and installation location of the MyODBC components. It also enables you to build and install MyODBC on platforms where a pre-compiled binary is not available.

MyODBC sources are available either as a downloadable package, or through the revision control system used by the MyODBC developers.

### 18.1.2.4.1. Installing MyODBC from a Source Distribution on Windows

You should only need to install MyODBC from source on Windows if you want to change or modify the source or installation. If you are unsure whether to install from source, please use the binary installation detailed in [Section 18.1.2.3.1, “Installing MyODBC from a Binary Distribution on Windows”](#).

Installing MyODBC from source on Windows requires a number of different tools and packages:

- MDAC, Microsoft Data Access SDK from <http://www.microsoft.com/data/>.
- Suitable C compiler, such as Microsoft Visual C++ or the C compiler included with Microsoft Visual Studio.
- Compatible `make` tool. Microsoft's `nmake` is used in the examples in this section.
- MySQL client libraries and include files from MySQL 4.0.0 or higher. (Preferably MySQL 4.0.16 or higher). This is required because MyODBC uses new calls and structures that exist only starting from this version of the library. To get the client libraries and include files, visit <http://dev.mysql.com/downloads/>.

#### 18.1.2.4.1.1. Building MyODBC 3.51

MyODBC source distributions include `Makefiles` that require the `nmake` or other `make` utility. In the distribution, you can find `Makefile` for building the release version and `Makefile_debug` for building debugging versions of the driver libraries and DLLs.

To build the driver, use this procedure:

1. Download and extract the sources to a folder, then change directory into that folder. The following command assumes the folder is named `myodbc3-src`:

```
C:\> cd myodbc3-src
```

2. Edit `Makefile` to specify the correct path for the MySQL client libraries and header files. Then use the following commands to build and install the release version:

```
C:\> nmake -f Makefile
C:\> nmake -f Makefile install
```

`nmake -f Makefile` builds the release version of the driver and places the binaries in subdirectory called `Release`.

`nmake -f Makefile install` installs (copies) the driver DLLs and libraries (`myodbc3.dll`, `myodbc3.lib`) to your system directory.

3. To build the debug version, use `Makefile_Debug` rather than `Makefile`, as shown below:

```
C:\> nmake -f Makefile_debug
C:\> nmake -f Makefile_debug install
```

4. You can clean and rebuild the driver by using:

```
C:\> nmake -f Makefile_clean
C:\> nmake -f Makefile_install
```

**Note:**

- Make sure to specify the correct MySQL client libraries and header files path in the Makefiles (set the `MYSQL_LIB_PATH` and `MYSQL_INCLUDE_PATH` variables). The default header file path is assumed to be `C:\mysql\include`. The default library path is assumed to be `C:\mysql\lib\opt` for release DLLs and `C:\mysql\lib\debug` for debug versions.
- For the complete usage of `nmake`, visit [http://msdn.microsoft.com/library/default.asp?url=/library/en-us/dv\\_vcce4/html/evgrfRunningNMAKE.asp](http://msdn.microsoft.com/library/default.asp?url=/library/en-us/dv_vcce4/html/evgrfRunningNMAKE.asp).
- If you are using the Subversion tree for compiling, all Windows-specific `Makefiles` are named as `Win_Makefile*`.

#### 18.1.2.4.1.2. Testing

After the driver libraries are copied/installed to the system directory, you can test whether the libraries are properly built by using the samples provided in the `samples` subdirectory:

```
C:\> cd samples
C:\> nmake -f Makefile_all
```

#### 18.1.2.4.1.3. Building MyODBC 2.50

The MyODBC 2.50 source distribution includes VC workspace files. You can build the driver using these files (`.dsp` and `.dsw`) directly by loading them from Microsoft Visual Studio 6.0 or higher.

#### 18.1.2.4.2. Installing MyODBC from a Source Distribution on Unix

You need the following tools to build MySQL from source on Unix:

- A working ANSI C++ compiler. `gcc` 2.95.2 or later, `egcs` 1.0.2 or later or `egcs 2.91.66`, SGI C++, and SunPro C++ are some of the compilers that are known to work.
- A good `make` program. GNU `make` is always recommended and is sometimes required.
- MySQL client libraries and include files from MySQL 4.0.0 or higher. (Preferably MySQL 4.0.16 or higher). This is required because MyODBC uses new calls and structures that exist only starting from this version of the library. To get the client libraries and include files, visit <http://dev.mysql.com/downloads/>.

If you have built your own MySQL server and/or client libraries from source then you must have used the `--enable-thread-safe-client` option to `configure` when the libraries were built.

You should also ensure that the `libmysqlclient` library were built and installed as a shared lib-

rary.

- A compatible ODBC manager must be installed. MyODBC is known to work with the `iODBC` and `unixODBC` managers. See [Section 18.1.1.2.2, “ODBC Driver Managers”](#), for more information.
- If you are using a character set that isn't compiled into the MySQL client library then you need to install the MySQL character definitions from the `charsets` directory into `SHAREDIR` (by default, `/usr/local/mysql/share/mysql/charsets`). These should be in place if you have installed the MySQL server on the same machine. See [Chapter 10, \*Character Set Support\*](#), for more information on character set support.

Once you have all the required files, unpack the source files to a separate directory, you then have to run `configure` and build the library using `make`.

#### 18.1.2.4.2.1. Typical `configure` Options

The `configure` script gives you a great deal of control over how you configure your MyODBC build. Typically you do this using options on the `configure` command line. You can also affect `configure` using certain environment variables. For a list of options and environment variables supported by `configure`, run this command:

```
shell> ./configure --help
```

Some of the more commonly used `configure` options are described here:

1. To compile MyODBC, you need to supply the MySQL client include and library files path using the `--with-mysql-path=DIR` option, where `DIR` is the directory where MySQL is installed.

MySQL compile options can be determined by running `DIR/bin/mysql_config`.

2. Supply the standard header and library files path for your ODBC Driver Manager (`iODBC` or `unixODBC`).
  - If you are using `iODBC` and `iODBC` is not installed in its default location (`/usr/local`), you might have to use the `--with-iodbc=DIR` option, where `DIR` is the directory where `iODBC` is installed.

If the `iODBC` headers do not reside in `DIR/include`, you can use the `--with-iodbc-includes=INCDIR` option to specify their location.

The applies to libraries. If they are not in `DIR/lib`, you can use the `--with-iodbc-libs=LIBDIR` option.

- If you are using `unixODBC`, use the `--with-unixODBC=DIR` option (case sensitive) to make `configure` look for `unixODBC` instead of `iODBC` by default, `DIR` is the directory where `unixODBC` is installed.

If the `unixODBC` headers and libraries aren't located in `DIR/include` and `DIR/lib`, use the `--with-unixODBC-includes=INCDIR` and `--with-unixODBC-libs=LIBDIR` options.

3. You might want to specify an installation prefix other than `/usr/local`. For example, to install the MyODBC drivers in `/usr/local/odbc/lib`, use the `--prefix=/usr/local/odbc` option.

The final configuration command looks something like this:

```
shell> ./configure --prefix=/usr/local \
--with-iodbc=/usr/local \
--with-mysql-path=/usr/local/mysql
```

#### 18.1.2.4.2.2. Additional configure Options

There are a number of other options that you need, or want, to set when configuring the MyODBC driver before it is built.

- To link the driver with MySQL thread safe client libraries `libmysqlclient_r.so` or `libmysqlclient_r.a`, you must specify the following `configure` option:

```
--enable-thread-safe
```

and can be disabled (default) using

```
--disable-thread-safe
```

This option enables the building of the driver thread-safe library `libmyodbc3_r.so` from by linking with MySQL thread-safe client library `libmysqlclient_r.so` (The extensions are OS dependent).

If the compilation with the thread-safe option fails, it may be because the correct thread-libraries on the system could not be located. You should set the value of `LIBS` to point to the correct thread library for your system.

```
LIBS="-lpthread" ./configure ..
```

- You can enable or disable the shared and static versions of MyODBC using these options:

```
--enable-shared[=yes/no]
--disable-shared
--enable-static[=yes/no]
--disable-static
```

- By default, all the binary distributions are built as non-debugging versions (configured with `--without-debug`).

To enable debugging information, build the driver from source distribution and use the `--with-debug` option when you run `configure`.

- This option is available only for source trees that have been obtained from the Subversion repository. This option does not apply to the packaged source distributions.

By default, the driver is built with the `--without-docs` option. If you would like the documentation to be built, then execute `configure` with:

```
--with-docs
```

#### 18.1.2.4.2.3. Building and Compilation

To build the driver libraries, you have to just execute `make`.

```
shell> make
```

If any errors occur, correct them and continue the build process. If you aren't able to build, then send a

detailed email to [<myodbc@lists.mysql.com>](mailto:myodbc@lists.mysql.com) for further assistance.

#### 18.1.2.4.2.4. Building Shared Libraries

On most platforms, MySQL does not build or support `.so` (shared) client libraries by default. This is based on our experience of problems when building shared libraries.

In cases like this, you have to download the MySQL distribution and configure it with these options:

```
--without-server --enable-shared
```

To build shared driver libraries, you must specify the `--enable-shared` option for `configure`. By default, `configure` does not enable this option.

If you have configured with the `--disable-shared` option, you can build the `.so` file from the static libraries using the following commands:

```
shell> cd MyODBC-3.51.01
shell> make
shell> cd driver
shell> CC=/usr/bin/gcc \
  $CC -bundle -flat_namespace -undefined error \
  -o .libs/libmyodbc3-3.51.01.so \
  catalog.o connect.o cursor.o dll.o error.o execute.o \
  handle.o info.o misc.o myodbc3.o options.o prepare.o \
  results.o transact.o utility.o \
  -L/usr/local/mysql/lib/mysql/ \
  -L/usr/local/iodbc/lib/ \
  -lz -lc -lmysqlclient -liodbcinst
```

Make sure to change `-liodbcinst` to `-lodbcinst` if you are using `unixODBC` instead of `iodbc`, and configure the library paths accordingly.

This builds and places the `libmyodbc3-3.51.01.so` file in the `.libs` directory. Copy this file to the MyODBC library installation directory (`/usr/local/lib` (or the `lib` directory under the installation directory that you supplied with the `--prefix`).

```
shell> cd .libs
shell> cp libmyodbc3-3.51.01.so /usr/local/lib
shell> cd /usr/local/lib
shell> ln -s libmyodbc3-3.51.01.so libmyodbc3.so
```

To build the thread-safe driver library:

```
shell> CC=/usr/bin/gcc \
  $CC -bundle -flat_namespace -undefined error \
  -o .libs/libmyodbc3_r-3.51.01.so \
  catalog.o connect.o cursor.o dll.o error.o execute.o \
  handle.o info.o misc.o myodbc3.o options.o prepare.o \
  results.o transact.o utility.o \
  -L/usr/local/mysql/lib/mysql/ \
  -L/usr/local/iodbc/lib/ \
  -lz -lc -lmysqlclient_r -liodbcinst
```

#### 18.1.2.4.2.5. Installing Driver Libraries

To install the driver libraries, execute the following command:

```
shell> make install
```

That command installs one of the following sets of libraries:

For MyODBC 3.51:

- `libmyodbc3.so`
- `libmyodbc3-3.51.01.so`, where 3.51.01 is the version of the driver
- `libmyodbc3.a`

For thread-safe MyODBC 3.51:

- `libmyodbc3_r.so`
- `libmyodbc3-3_r.51.01.so`
- `libmyodbc3_r.a`

For MyODBC 2.5.0:

- `libmyodbc.so`
- `libmyodbc-2.50.39.so`, where 2.50.39 is the version of the driver
- `libmyodbc.a`

For more information on build process, refer to the [INSTALL](#) file that comes with the source distribution. Note that if you are trying to use the `make` from Sun, you may end up with errors. On the other hand, GNU `gmake` should work fine on all platforms.

#### 18.1.2.4.2.6. Testing MyODBC on Unix

To run the basic samples provided in the distribution with the libraries that you built, use the following command:

```
shell> make test
```

Before running the tests, create the DSN 'myodbc3' in `odbc.ini` and set the environment variable `ODBCINI` to the correct `odbc.ini` file; and MySQL server is running. You can find a sample `odbc.ini` with the driver distribution.

You can even modify the `samples/run-samples` script to pass the desired DSN, UID, and PASSWORD values as the command-line arguments to each sample.

#### 18.1.2.4.2.7. Building MyODBC from Source on Mac OS X

To build the driver on Mac OS X (Darwin), make use of the following `configure` example:

```
shell> ./configure --prefix=/usr/local
--with-unixODBC=/usr/local
--with-mysql-path=/usr/local/mysql
--disable-shared
--enable-gui=no
--host=powerpc-apple
```

The command assumes that the `unixODBC` and MySQL are installed in the default locations. If not, configure accordingly.

On Mac OS X, `--enable-shared` builds `.dylib` files by default. You can build `.so` files like this:

```
shell> make
shell> cd driver
shell> CC=/usr/bin/gcc \
    $CC -bundle -flat_namespace -undefined error
    -o .libs/libmyodbc3-3.51.01.so *.o
    -L/usr/local/mysql/lib/
    -L/usr/local/iodbc/lib
    -liodbcinst -lmysqlclient -lz -lc
```

To build the thread-safe driver library:

```
shell> CC=/usr/bin/gcc \
    $CC -bundle -flat_namespace -undefined error
    -o .libs/libmyodbc3-3.51.01.so *.o
    -L/usr/local/mysql/lib/
    -L/usr/local/iodbc/lib
    -liodbcinst -lmysqlclient_r -lz -lc -lpthread
```

Make sure to change the `-liodbcinst` to `-lodbcinst` in case of using `unixODBC` instead of `iODBC` and configure the libraries path accordingly.

In Apple's version of GCC, both `cc` and `gcc` are actually symbolic links to `gcc3`.

Copy this library to the `$prefix/lib` directory and symlink to `libmyodbc3.so`.

You can cross-check the output shared-library properties using this command:

```
shell> otool -LD .libs/libmyodbc3-3.51.01.so
```

#### 18.1.2.4.2.8. Building MyODBC from Source on HP-UX

To build the driver on HP-UX 10.x or 11.x, make use of the following `configure` example:

If using `cc`:

```
shell> CC="cc" \
    CFLAGS="+z" \
    LDFLAGS="-Wl,+b:-Wl,+s" \
    ./configure --prefix=/usr/local
    --with-unixodbc=/usr/local
    --with-mysql-path=/usr/local/mysql/lib/mysql
    --enable-shared
    --enable-thread-safe
```

If using `gcc`:

```
shell> CC="gcc" \
    LDFLAGS="-Wl,+b:-Wl,+s" \
    ./configure --prefix=/usr/local
    --with-unixodbc=/usr/local
    --with-mysql-path=/usr/local/mysql
    --enable-shared
    --enable-thread-safe
```

Once the driver is built, cross-check its attributes using `chatr .libs/libmyodbc3.sl` to determine whether you need to have set the MySQL client library path using the `SHLIB_PATH` environment variable. For static versions, ignore all shared-library options and run `configure` with the `--disable-shared` option.

#### 18.1.2.4.2.9. Building MyODBC from Source on AIX

To build the driver on AIX, make use of the following `configure` example:

```
shell> ./configure --prefix=/usr/local
```



```
--with-unixodbc=/usr/local
--with-mysql-path=/usr/local/mysql
--disable-shared
--enable-thread-safe
```

**NOTE:** For more information about how to build and set up the static and shared libraries across the different platforms refer to ' [Using static and shared libraries across platforms](http://www.fortran-2000.com/ArnaudRecipes/sharedlib.html) [<http://www.fortran-2000.com/ArnaudRecipes/sharedlib.html>]'.

### 18.1.2.4.3. Installing MyODBC from the Development Source Tree

**Caution:** You should read this section only if you are interested in helping us test our new code. If you just want to get MySQL Connector/ODBC up and running on your system, you should use a standard release distribution.

To be able to access the MyODBC source tree, you must have Subversion installed. Subversion is freely available from <http://subversion.tigris.org/>.

To build from the source trees, you need the following tools:

- autoconf 2.52 (or newer)
- automake 1.4 (or newer)
- libtool 1.4 (or newer)
- m4

The most recent development source tree is available from our public Subversion trees at <http://dev.mysql.com/tech-resources/sources.html>.

To checkout out the Connector/ODBC sources, change to the directory where you want the copy of the MyODBC tree to be stored, then use the following command:

```
shell> svn co http://svn.mysql.com/svnpublic/connector-odbc3
```

You should now have a copy of the entire MyODBC source tree in the directory `connector-odbc3`. To build from this source tree on Unix or Linux follow these steps:

```
shell> cd connector-odbc3
shell> aclocal
shell> autoheader
shell> autoconf
shell> automake;
shell> ./configure # Add your favorite options here
shell> make
```

For more information on how to build, refer to the `INSTALL` file located in the same directory. For more information on options to `configure`, see [Section 18.1.2.4.2.1, “Typical configure Options”](#)

When the build is done, run `make install` to install the MyODBC 3.51 driver on your system.

If you have gotten to the `make` stage and the distribution does not compile, please report it to [<myodbc@lists.mysql.com>](mailto:myodbc@lists.mysql.com).

On Windows, make use of Windows Makefiles `WIN-Makefile` and `WIN-Makefile_debug` in building the driver. For more information, see [Section 18.1.2.4.1, “Installing MyODBC from a Source Distribution on Windows”](#).

After the initial checkout operation to get the source tree, you should run `svn update` periodically update your source according to the latest version.

## 18.1.3. MyODBC Configuration

Before you connect to a MySQL database using the MyODBC driver you must configure an ODBC *Data Source Name*. The DSN associates the various configuration parameters required to communicate with a database to a specific name. You use the DSN in an application to communicate with the database, rather than specifying individual parameters within the application itself. DSN information can be user specific, system specific, or provided in a special file. ODBC data source names are configured in different ways, depending on your platform and ODBC driver.

### 18.1.3.1. Data Source Names

A Data Source Name associates the configuration parameters for communicating with a specific database. Generally a DSN consists of the following parameters:

- Name
- Hostname
- Database Name
- Login
- Password

In addition, different ODBC drivers, including MyODBC, may accept additional driver-specific options and parameters.

There are three types of DSN:

- A *System DSN* is a global DSN definition that is available to any user and application on a particular system. A System DSN can normally only be configured by a systems administrator, or by a user who has specific permissions that let them create System DSNs.
- A *User DSN* is specific to an individual user, and can be used to store database connectivity information that the user regularly uses.
- A *File DSN* uses a simple file to define the DSN configuration. File DSNs can be shared between users and machines and are therefore more practical when installing or deploying DSN information as part of an application across many machines.

DSN information is stored in different locations depending on your platform and environment.

### 18.1.3.2. Configuring a MyODBC DSN on Windows

The [ODBC Data Source Administrator](#) within Windows enables you to create DSNs, check driver installation and configure ODBC systems such as tracing (used for debugging) and connection pooling.

Different editions and versions of Windows store the [ODBC Data Source Administrator](#) in different locations depending on the version of Windows that you are using.

To open the [ODBC Data Source Administrator](#) in Windows Server 2003:

1. On the [Start](#) menu, choose [Administrative Tools](#), and then click [Data Sources \(ODBC\)](#).

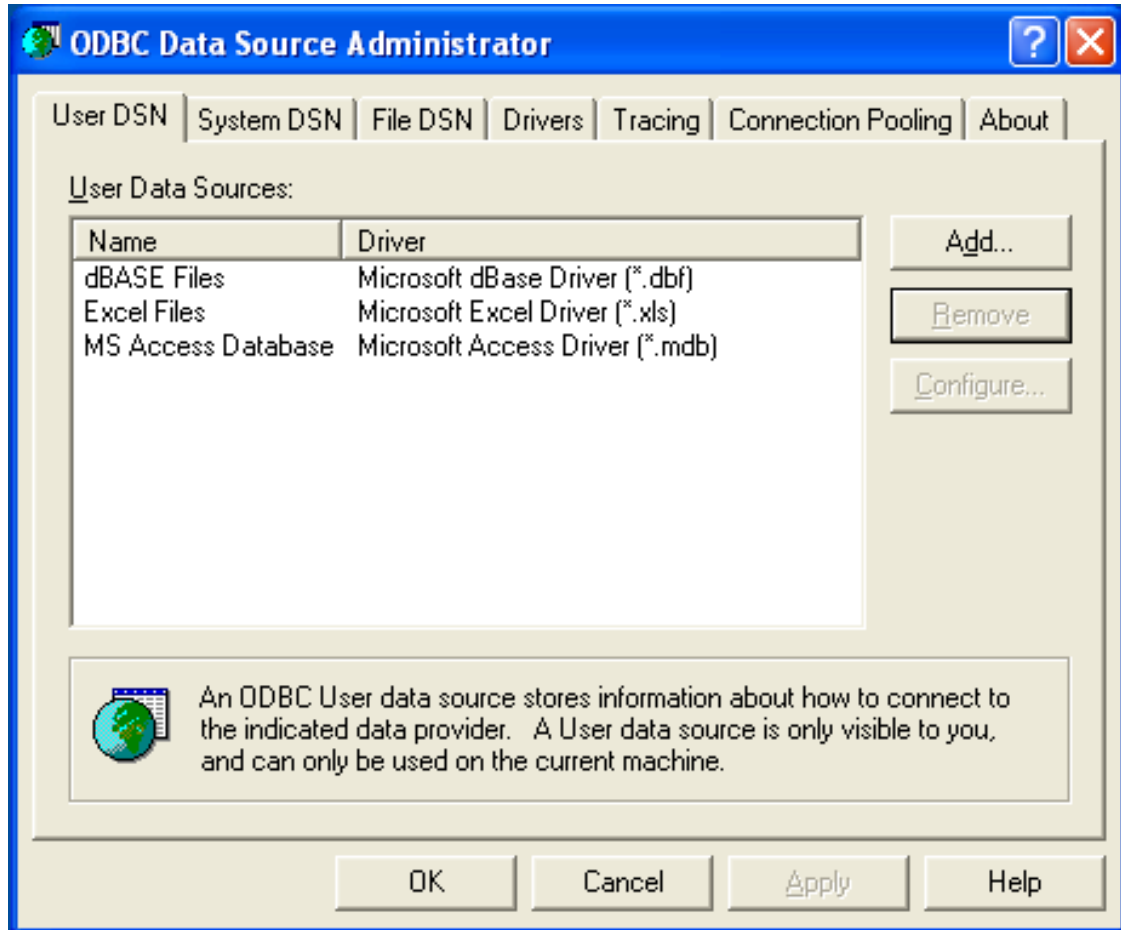
To open the [ODBC Data Source Administrator](#) in Windows 2000 Server or Windows 2000 Professional:

1. On the [Start](#) menu, choose [Settings](#), and then click [Control Panel](#).
2. In [Control Panel](#), click [Administrative Tools](#).
3. In [Administrative Tools](#), click [Data Sources \(ODBC\)](#).

To open the [ODBC Data Source Administrator](#) on Windows XP:

1. On the [Start](#) menu, click [Control Panel](#).
2. In the [Control Panel](#) when in [Category View](#) click [Performance and Maintenance](#) and then click [Administrative Tools](#). If you are viewing the [Control Panel](#) in [Classic View](#), click [Administrative Tools](#).
3. In [Administrative Tools](#), click [Data Sources \(ODBC\)](#).

Irrespective of your Windows version, you should be presented the [ODBC Data Source Administrator](#) window:



Within Windows XP, you can add the [Administrative Tools](#) folder to your [Start](#) menu to make it easier to locate the ODBC Data Source Administrator. To do this:

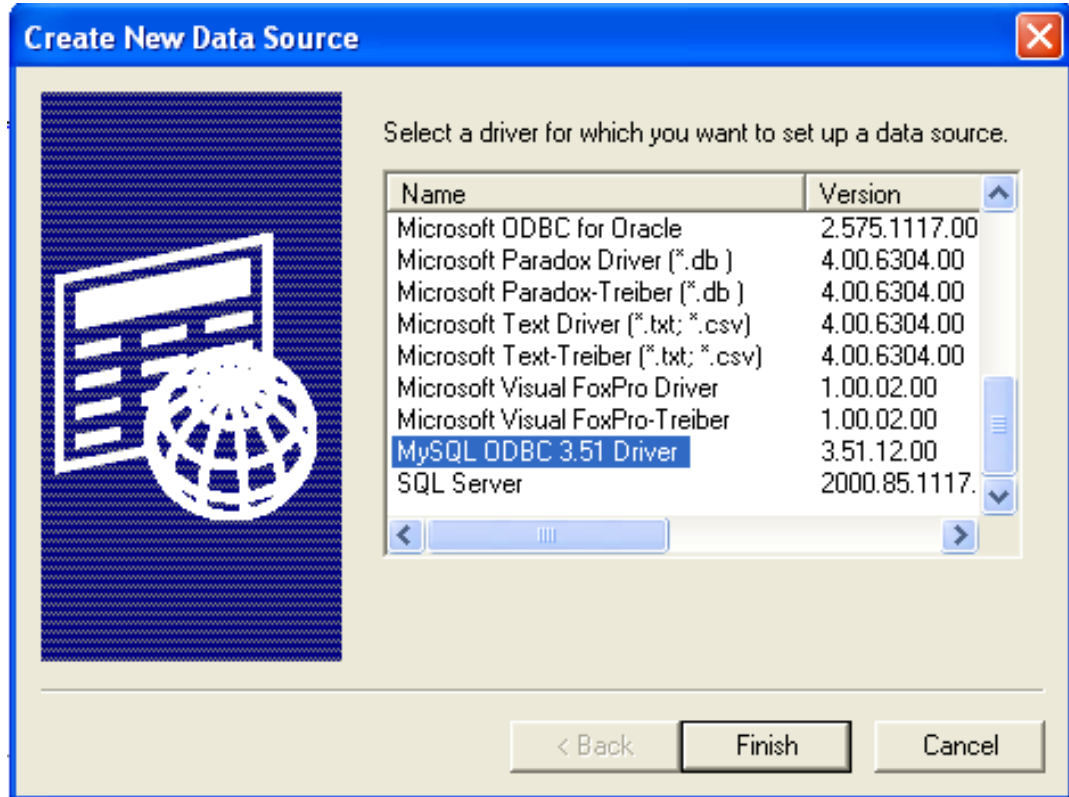
1. Right click on the [Start](#) menu.
2. Select [Properties](#).
3. Click [Customize...](#)
4. Select the [Advanced](#) tab.
5. Within [Start menu items](#), within the [System Administrative Tools](#) section, select [Display on the All Programs menu](#).

Within both Windows Server 2003 and Windows XP you may want to permanently add the [ODBC Data Source Administrator](#) to your [Start](#) menu. To do this, locate the [Data Sources \(ODBC\)](#) icon using the methods shown, then right-click on the icon and then choose [Pin to Start Menu](#).

### 18.1.3.2.1. Adding a MyODBC DSN on Windows


To add and configure a new MyODBC data source on Windows, use the [ODBC Data Source Administrator](#):

1. Open the [ODBC Data Source Administrator](#).
2. To create a System DSN (which will be available to all users) , select the [System DSN](#) tab. To create a User DSN, which will be unique only to the current user, click the [Add. .](#) button.
3. You will need to select the ODBC driver for this DSN.



Select [MySQL ODBC 3.51 Driver](#), then click [Finish](#).

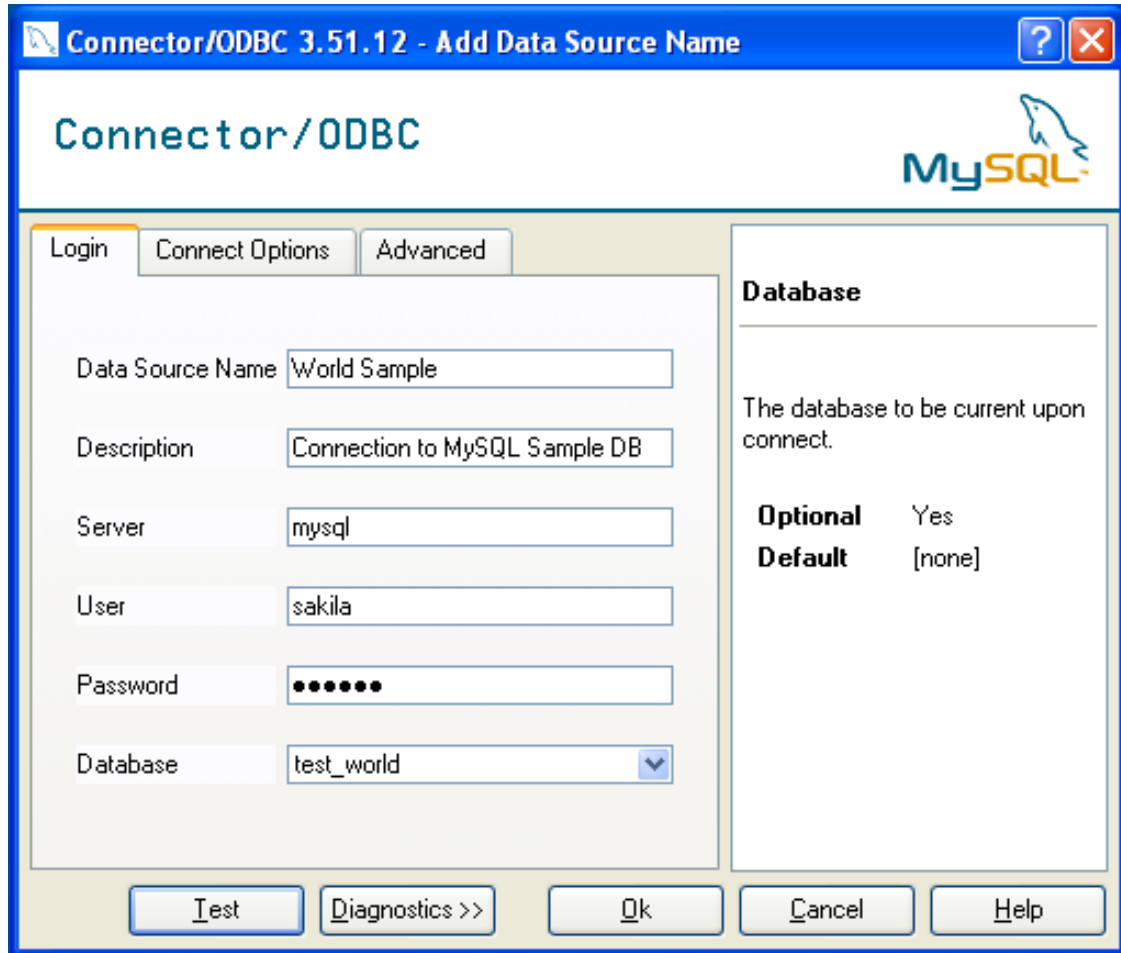
4. You now need to configure the specific fields for the DSN you are creating through the [Add Data Source Name](#) dialog.



In the **Data Source Name** box, enter the name of the data source you want to access. It can be any valid name that you choose.

5. In the **Description** box, enter some text to help identify the connection.
6. In the **Server** field, enter the name of the MySQL server host that you want to access. By default, it is `localhost`.
7. In the **User** field, enter the user name to use for this connection.
8. In the **Password** field, enter the corresponding password for this connection.
9. The **Database** popup should automatically populate with the list of databases that the user has permissions to access.
10. Click **OK** to save the DSN.

A completed DSN configuration may look like this:



### 18.1.3.2.2. Checking MyODBC DSN Configuration on Windows


You can verify the connection using the parameters you have entered by clicking the [Test](#) button. If the connection could be made successfully, you will be notified with a [Success; connection was made!](#) dialog.

If the connection failed, you can obtain more information on the test and why it may have failed by clicking the [Diagnostics...](#) button to show additional error messages.

### 18.1.3.2.3. MyODBC DSN Configuration Options

You can configure a number of options for a specific DSN by using either the [Connect Options](#) or [Advanced](#) tabs in the DSN configuration dialog.

The [Connection Options](#) dialog can be seen below.




The three options you can configure are:

- **Port** sets the TCP/IP port number to use when communicating with MySQL. Communication with MySQL uses port 3306 by default. If your server is configured to use a different TCP/IP port, you must specify that port number here.
- **Socket** sets the name or location of a specific socket or Windows pipe to use when communicating with MySQL.
- **Initial Statement** defines an SQL statement that will be executed when the connection to MySQL is opened. You can use this to set MySQL options for your connection, such as setting the default character set or database to use during your connection.

The **Advanced** tab enables you to configure MyODBC connection parameters. Refer to [Section 18.1.3.5, “MyODBC Connection Parameters”](#), for information about the meaning of these options.





#### 18.1.3.2.4. Errors and Debugging

This section answers MyODBC connection-related questions.

- **While configuring a MyODBC DSN, a `Could Not Load Translator or Setup Library` error occurs**

For more information, refer to [MS KnowledgeBase Article\(Q260558\)](http://support.microsoft.com/default.aspx?scid=kb;EN-US;q260558)

[<http://support.microsoft.com/default.aspx?scid=kb;EN-US;q260558>]. Also, make sure you have the latest valid `ct13d32.dll` in your system directory.

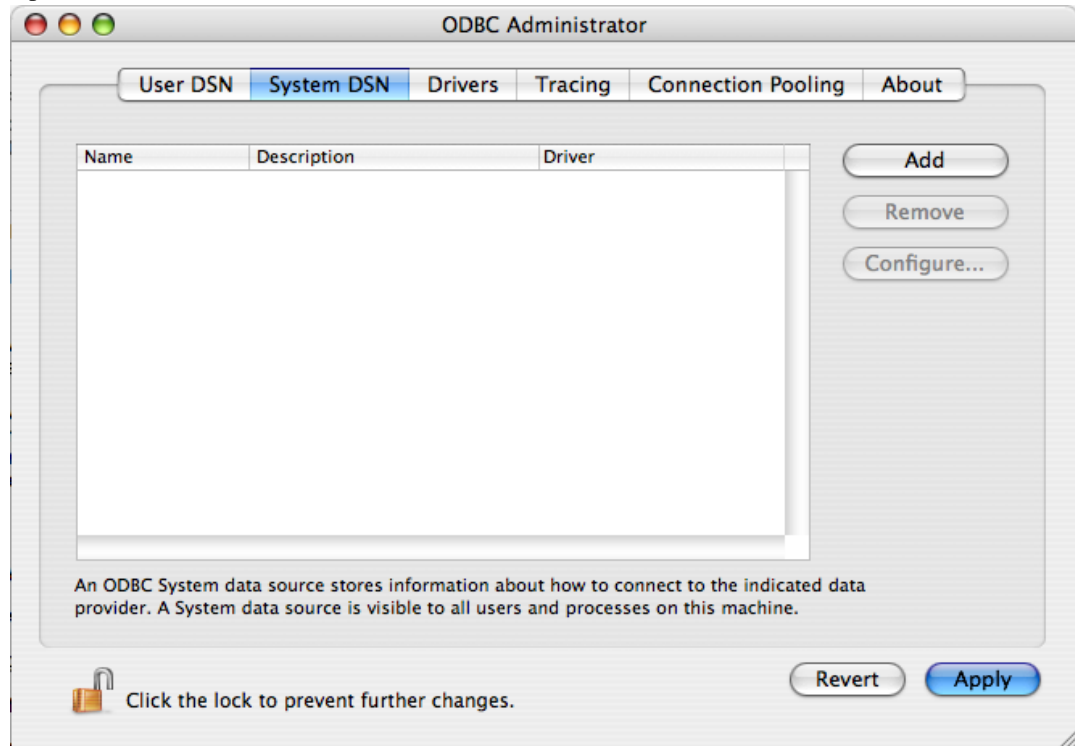
- On Windows, the default `myodbc3.dll` is compiled for optimal performance. If you want to debug MyODBC 3.51 (for example, to enable tracing), you should instead use `myodbc3d.dll`. To install this file, copy `myodbc3d.dll` over the installed `myodbc3.dll` file. Make sure to revert back to the release version of the driver DLL once you are done with the debugging because the debug version may cause performance issues. Note that the `myodbc3d.dll` isn't included in MyODBC 3.51.07 through 3.51.11. If you are using one of these versions, you should copy that DLL from a previous version (for example, 3.51.06).

For MyODBC 2.50, `myodbc.dll` and `myodbcd.dll` are used instead.


#### 18.1.3.3. Configuring a MyODBC DSN on Mac OS X

To configure a DSN on Mac OS X you should use the ODBC Administrator. If you have Mac OS X 10.2 or earlier, refer to [Section 18.1.3.4, “Configuring a MyODBC DSN on Unix”](#). Select whether you want to create a User DSN or a System DSN. If you want to add a System DSN, you may need to authenticate with the system. You must click the padlock and enter a user and password with administrator privileges.

1. Open the ODBC Administrator from the [Utilities](#) folder in the [Applications](#) folder.




2. On the User DSN or System DSN panel, click [Add](#).
3. Select the MyODBC driver and click [OK](#).
4. You will be presented with the [Data Source Name](#) dialog. Enter The [Data Source Name](#) and an optional [Description](#) for the DSN.



5. Click **Add** to add a new keyword/value pair to the panel. You should configure at least four pairs to specify the `server`, `username`, `password` and `database` connection parameters. See [Section 18.1.3.5, “MyODBC Connection Parameters”](#).
6. Click **OK** to add the DSN to the list of configured data source names.

A completed DSN configuration may look like this:



You can configure additional ODBC options to your DSN by adding further keyword/value pairs and

setting the corresponding values. See [Section 18.1.3.5, “MyODBC Connection Parameters”](#).

### 18.1.3.4. Configuring a MyODBC DSN on Unix

On [Unix](#), you configure DSN entries directly in the `odbc.ini` file. Here is a typical `odbc.ini` file that configures `myodbc` and `myodbc3` as the DSN names for MyODBC 2.50 and MyODBC 3.51, respectively:

```

;
; odbc.ini configuration for MyODBC and MyODBC 3.51 drivers
;

[ODBC Data Sources]
myodbc      = MyODBC 2.50 Driver DSN
myodbc3     = MyODBC 3.51 Driver DSN

[myodbc]
Driver      = /usr/local/lib/libmyodbc.so
Description = MyODBC 2.50 Driver DSN
SERVER      = localhost
PORT        =
USER        = root
Password    =
Database    = test
OPTION      = 3
SOCKET      =

[myodbc3]
Driver      = /usr/local/lib/libmyodbc3.so
Description = MyODBC 3.51 Driver DSN
SERVER      = localhost
PORT        =
USER        = root
Password    =
Database    = test
OPTION      = 3
SOCKET      =

[Default]
Driver      = /usr/local/lib/libmyodbc3.so
Description = MyODBC 3.51 Driver DSN
SERVER      = localhost
PORT        =
USER        = root
Password    =
Database    = test
OPTION      = 3
SOCKET      =

```

Refer to the [Section 18.1.3.5, “MyODBC Connection Parameters”](#), for the list of connection parameters that can be supplied.

**Note:** If you are using `unixODBC`, you can use the following tools to set up the DSN:

- ODBCConfig GUI tool([HOWTO: ODBCConfig](http://www.unixodbc.org/config.html) [<http://www.unixodbc.org/config.html>])
- `odbcinst`

In some cases when using `unixODBC`, you might get this error:

```
Data source name not found and no default driver specified
```

If this happens, make sure the `ODBCINI` and `ODBCSYSINI` environment variables are pointing to the right `odbc.ini` file. For example, if your `odbc.ini` file is located in `/usr/local/etc`, set the environment variables like this:

```
export ODBCINI=/usr/local/etc/odbc.ini
```

```
export ODBCYSINI=/usr/local/etc
```

### 18.1.3.5. MyODBC Connection Parameters

You can specify the parameters in the following tables for MyODBC when configuring a DSN. Users on Windows can use the Options and Advanced panels when configuring a DSN to set these parameters; see the table for information on which options relate to which fields and checkboxes. On Unix and Mac OS X, use the parameter name and value as the keyword/value pair in the DSN configuration. Alternatively, you can set these parameters within the `InConnectionString` argument in the `SQLDriverConnect()` call.

Parameter	Default Value	Comment
<code>user</code>	ODBC (on Windows)	The username used to connect to MySQL.
<code>server</code>	<code>localhost</code>	The hostname of the MySQL server.
<code>database</code>		The default database.
<code>option</code>	0	Options that specify how MyODBC should work. See below.
<code>port</code>	3306	The TCP/IP port to use if <code>server</code> is not <code>localhost</code> .
<code>stmt</code>		A statement to execute when connecting to MySQL.
<code>password</code>		The password for the <code>user</code> account on <code>server</code> .
<code>socket</code>		The Unix socket file or Windows named pipe to connect to if <code>server</code> is <code>localhost</code> .

The `option` argument is used to tell MyODBC that the client isn't 100% ODBC compliant. On Windows, you normally select options by toggling the checkboxes in the connection screen, but you can also select them in the `option` argument. The following options are listed in the order in which they appear in the MyODBC connect screen:

Value	Windows Checkbox	Description
1	Don't Optimized Column Width	The client can't handle that MyODBC returns the real width of a column.
2	Return Matching Rows	The client can't handle that MySQL returns the true value of affected rows. If this flag is set, MySQL returns "found rows" instead. You must have MySQL 3.21.14 or newer to get this to work.
4	Trace Driver Calls To myodbc.log	Make a debug log in <code>C:\myodbc.log</code> on Windows, or <code>/tmp/myodbc.log</code> on Unix variants.
8	Allow Big Results	Don't set any packet limit for results and parameters.
16	Don't Prompt Upon Connect	Don't prompt for questions even if driver would like to prompt.
32	Enable Dynamic Cursor	Enable or disable the dynamic cursor support. (Not allowed in MyODBC 2.50.)
64	Ignore # in Table Name	Ignore use of database name in <code>db_name.tbl_name.col_name</code> .
128	User Manager Cursors	Force use of ODBC manager cursors (experimental).
256	Don't Use Set Locale	Disable the use of extended fetch (experimental).
512	Pad Char To Full Length	Pad <code>CHAR</code> columns to full column length.
1024	Return Table Names for SQL-	<code>SQLDescribeCol()</code> returns fully qualified column

	DescribeCol	names.
2048	Use Compressed Protocol	Use the compressed client/server protocol.
4096	Ignore Space After Function Names	Tell server to ignore space after function name and before '(' (needed by PowerBuilder). This makes all function names keywords.
8192	Force Use of Named Pipes	Connect with named pipes to a <code>mysqld</code> server running on NT.
16384	Change BIGINT Columns to Int	Change <code>BIGINT</code> columns to <code>INT</code> columns (some applications can't handle <code>BIGINT</code> ).
32768	No Catalog (exp)	Return 'user' as <code>Table_qualifier</code> and <code>Table_owner</code> from <code>SQLTables</code> (experimental).
65536	Read Options From <code>my.cnf</code>	Read parameters from the <code>[client]</code> and <code>[odbc]</code> groups from <code>my.cnf</code> .
131072	Safe	Add some extra safety checks (should not be needed but...).
262144	Disable transaction	Disable transactions.
524288	Save queries to <code>myodbc.sql</code>	Enable query logging to <code>c:\myodbc.sql(/tmp/myodbc.sql)</code> file. (Enabled only in debug mode.)
1048576	Don't Cache Result (forward only cursors)	Do not cache the results locally in the driver, instead read from server ( <code>mysql_use_result()</code> ). This works only for forward-only cursors. This option is very important in dealing with large tables when you don't want the driver to cache the entire result set.
2097152	Force Use Of Forward Only Cursors	Force the use of <code>Forward-only</code> cursor type. In case of applications setting the default static/dynamic cursor type, and one wants the driver to use non-cache result sets, then this option ensures the forward-only cursor behavior.

To select multiple options, add together their values. For example, setting `option` to 12 (4+8) gives you debugging without packet limits.

The following table shows some recommended `option` values for various configurations:

Configuration	Option Value
Microsoft Access, Visual Basic	3
Driver trace generation (Debug mode)	4
Microsoft Access (with improved DELETE queries)	35
Large tables with too many rows	2049
Sybase PowerBuilder	135168
Query log generation (Debug mode)	524288
Generate driver trace as well as query log (Debug mode)	524292
Large tables with no-cache results	3145731

### 18.1.3.6. Connecting Without a Predefined DSN

You can connect to the MySQL server using `SQLDriverConnect`, by specifying the `DRIVER` name field. Here are the connection strings for MyODBC using DSN-Less connections:

**For MyODBC 2.50:**

```
ConnectionString = "DRIVER={MySQL};\  
SERVER=localhost;\  
DATABASE=test;\  
USER=venu;\  
PASSWORD=venu;\  
OPTION=3; "
```

**For MyODBC 3.51:**

```
ConnectionString = "DRIVER={MySQL ODBC 3.51 Driver};\  
SERVER=localhost;\  
DATABASE=test;\  
USER=venu;\  
PASSWORD=venu;\  
OPTION=3; "
```

If your programming language converts backslash followed by whitespace to a space, it is preferable to specify the connection string as a single long string, or to use a concatenation of multiple strings that does not add spaces in between. For example:

```
ConnectionString = "DRIVER={MySQL ODBC 3.51 Driver};"  
"SERVER=localhost;"  
"DATABASE=test;"  
"USER=venu;"  
"PASSWORD=venu;"  
"OPTION=3;"
```

**Note.** Note that on Mac OS X you may need to specify the full path to the MyODBC driver library.

Refer to the [Section 18.1.3.5, “MyODBC Connection Parameters”](#), for the list of connection parameters that can be supplied.

### 18.1.3.7. ODBC Connection Pooling

Connection pooling enables the ODBC driver to re-use existing connections to a given database from a pool of connections, instead of opening a new connection each time the database is accessed. By enabling connection pooling you can improve the overall performance of your application by lowering the time taken to open a connection to a database in the connection pool.

For more information about connection pooling: <http://support.microsoft.com/default.aspx?scid=kb;EN-US;q169470>.


### 18.1.3.8. Getting an ODBC Trace File

If you encounter difficulties or problems with MyODBC, you should start by making a log file from the [ODBC Manager](#) and MyODBC. This is called *tracing*, and is enabled through the ODBC Manager. The procedure for this differs for Windows, Mac OS X and Unix.

#### 18.1.3.8.1. Enabling ODBC Tracing on Windows

To enable the trace option on Windows:

1. The [Tracing](#) tab of the ODBC Data Source Administrator dialog box enables you to configure the way ODBC function calls are traced.



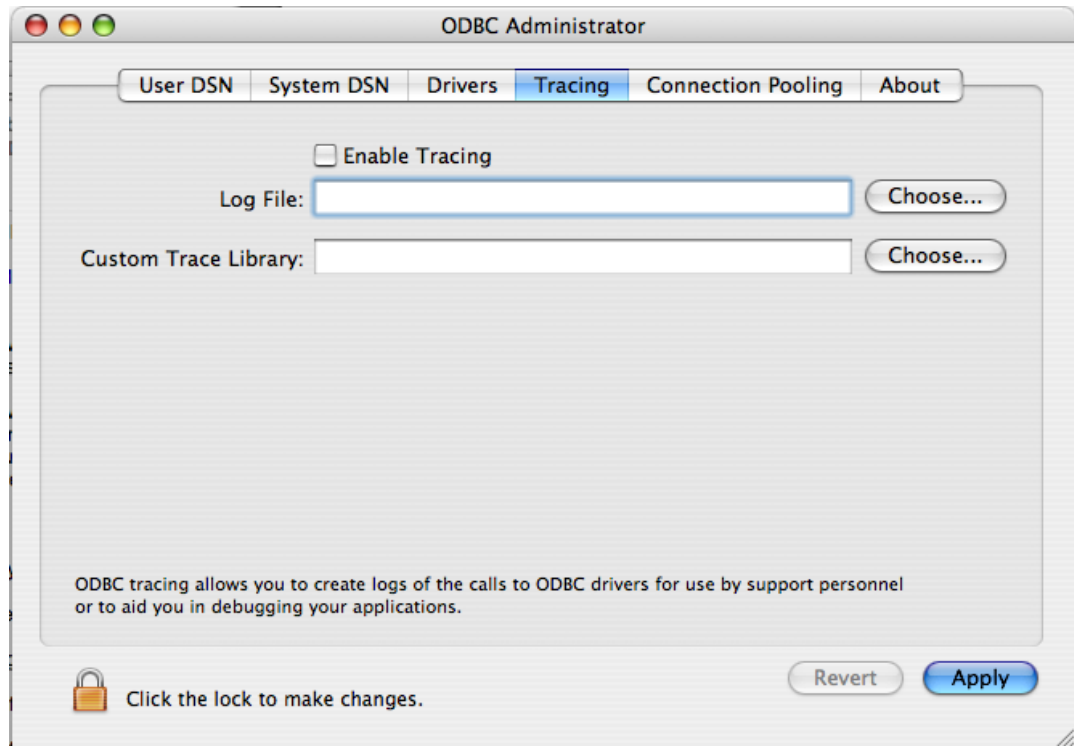
2. When you activate tracing from the [Tracing](#) tab, the [Driver Manager](#) logs all ODBC function calls for all subsequently run applications.
3. ODBC function calls from applications running before tracing is activated are not logged. ODBC function calls are recorded in a log file you specify.
4. Tracing ceases only after you click [Stop Tracing Now](#). Remember that while tracing is on, the log file continues to increase in size and that tracing affects the performance of all your ODBC applications.

### 18.1.3.8.2. Enabling ODBC Tracing on Mac OS X

To enable the trace option on Mac OS X 10.3 or later you should use the [Tracing](#) tab within ODBC Administrator .

1. Open the ODBC Administrator.
2. Select the [Tracing](#) tab.





3. Select the `Enable Tracing` checkbox.
4. Enter the location where you want to save the Tracing log. If you want to append information to an existing log file, click the `Choose...` button.

### 18.1.3.8.3. Enabling ODBC Tracing on Unix

To enable the trace option on Mac OS X 10.2 (or earlier) or Unix you must add the `trace` option to the ODBC configuration:

1. On Unix, you need to explicitly set the `Trace` option in the `ODBC.INI` file.

Set the tracing `ON` or `OFF` by using `TraceFile` and `Trace` parameters in `odbc.ini` as shown below:

```
TraceFile = /tmp/odbc.trace
Trace     = 1
```

`TraceFile` specifies the name and full path of the trace file and `Trace` is set to `ON` or `OFF`. You can also use `1` or `YES` for `ON` and `0` or `NO` for `OFF`. If you are using `ODBCConfig` from `unixODBC`, then follow the instructions for tracing `unixODBC` calls at [HOWTO-ODBCConfig](http://www.unixodbc.org/config.html) [<http://www.unixodbc.org/config.html>].

### 18.1.3.8.4. Enabling a MyODBC Log

To generate a MyODBC log, do the following:

1. Within Windows, enable the `Trace MyODBC` option flag in the MyODBC connect/configure

screen. The log is written to file `C:\myodbc.log`. If the trace option is not remembered when you are going back to the above screen, it means that you are not using the `myodbc.dll` driver, see [Section 18.1.3.2.4, “Errors and Debugging”](#).

On Mac OS X, Unix, or if you are using DSN-Less connection, then you need to supply `OPTIION=4` in the connection string or set the corresponding keyword/value pair in the DSN.

2. Start your application and try to get it to fail. Then check the MyODBC trace file to find out what could be wrong.

If you need help determining what is wrong, see [Section 18.1.7.1, “MyODBC Community Support”](#).

## 18.1.4. MyODBC Examples


Once you have configured a DSN to provide access to a database, how you access and use that connection is dependent on the application or programming language. As ODBC is a standardized interface, any application or language that supports ODBC can use the DSN and connect to the configured database.

### 18.1.4.1. Basic MyODBC Application Steps

Interacting with a MySQL server from an applications using the MyODBC typically involves the following operations:

- Configure the MyODBC DSN
- Connect to MySQL server
- Initialization operations
- Execute SQL statements
- Retrieve results
- Perform Transactions
- Disconnect from the server

Most applications use some variation of these steps. The basic application steps are shown in the following diagram:



### 18.1.4.2. Step-by-step Guide to Connecting to a MySQL Database through MyODBC

A typical installation situation where you would install MyODBC is when you want to access a database on a Linux or Unix host from a Windows machine.

As an example of the process required to set up access between two machines, the steps below take you through the basic steps. These instructions assume that you want to connect to system ALPHA from system BETA with a username and password of `myuser` and `mypassword`.

On system ALPHA (the MySQL server) follow these steps:

1. Start the MySQL server.
2. Use `GRANT` to set up an account with a username of `myuser` that can connect from system BETA using a password of `myuser` to the database `test`:

```
GRANT ALL ON test.* to 'myuser'@'BETA' IDENTIFIED BY 'mypassword';
```

For more information about MySQL privileges, refer to [Section 5.8, “MySQL User Account Management”](#).

On system BETA (the MyODBC client), follow these steps:

1. Configure a MyODBC DSN using parameters that match the server, database and authentication information that you have just configured on system ALPHA.

Parameter	Value	Comment
DSN	remote_test	A name to identify the connection.
SERVER	ALPHA	The address of the remote server.
DATABASE	test	The name of the default database.
USER	myuser	The username configured for access to this database.
PASSWORD	mypassword	The password for <code>myuser</code> .

2. Using an ODBC-capable application, such as Microsoft Office, connect to the MySQL server using the DSN you have just created. If the connection fails, use tracing to examine the connection process. See [Section 18.1.3.8, “Getting an ODBC Trace File”](#), for more information.

### 18.1.4.3. MyODBC and Third-Party ODBC Tools

Once you have configured your MyODBC DSN, you can access your MySQL database through any application that supports the ODBC interface, including programming languages and third-party applications. This section contains guides and help on using MyODBC with various ODBC-compatible tools and applications, including Microsoft Word, Microsoft Excel and Adobe/Macromedia ColdFusion.

#### 18.1.4.3.1. Applications Tested with MyODBC

MyODBC has been tested with the following applications:

Publisher	Application	Notes
Adobe	ColdFusion	Formerly Macromedia ColdFusion
Borland	C++ Builder	
	Builder 4	
	Delphi	
Business Objects	Crystal Reports	
Claris	Filemaker Pro	
Corel	Paradox	

Computer Associates	Visual Objects	Also known as CAVO
	AllFusion ERwin Data Modeler	
Gupta	Team Developer	Previously known as Centura Team Developer; Gupta SQL/Windows
Gensym	G2-ODBC Bridge	
Inline	iHTML	
Lotus	Notes	Versions 4.5 and 4.6
Microsoft	Access	
	Excel	
	Visio Enterprise	
	Visual C++	
	Visual Basic	
	ODBC.NET	Using C#, Visual Basic, C++
	FoxPro	
	Visual Interdev	
OpenOffice.org	OpenOffice.org	
Perl	DBD::ODBC	
Pervasive Software	DataJunction	
Sambar Technologies	Sambar Server	
SPSS	SPSS	
SoftVelocity	Clarion	
SQLExpress	SQLExpress for Xbase++	
Sun	StarOffice	
SunSystems	Vision	
Sybase	PowerBuilder	
	PowerDesigner	
theKompany.com	Data Architect	

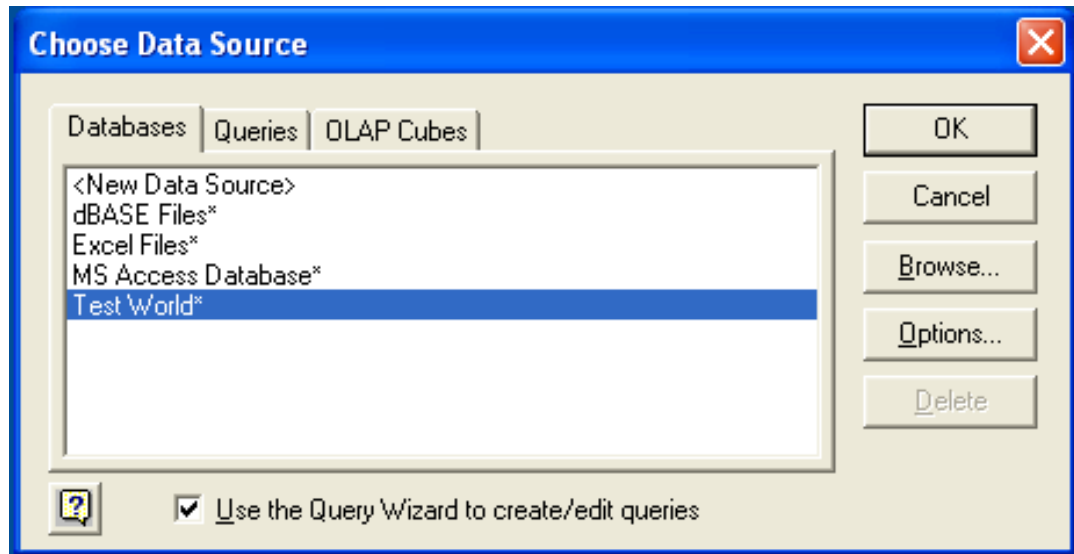
If you know of any other applications that work with MyODBC, please send mail to [myodbc@lists.mysql.com](mailto:myodbc@lists.mysql.com) about them.

#### 18.1.4.3.2. Using MyODBC with Microsoft Word or Excel

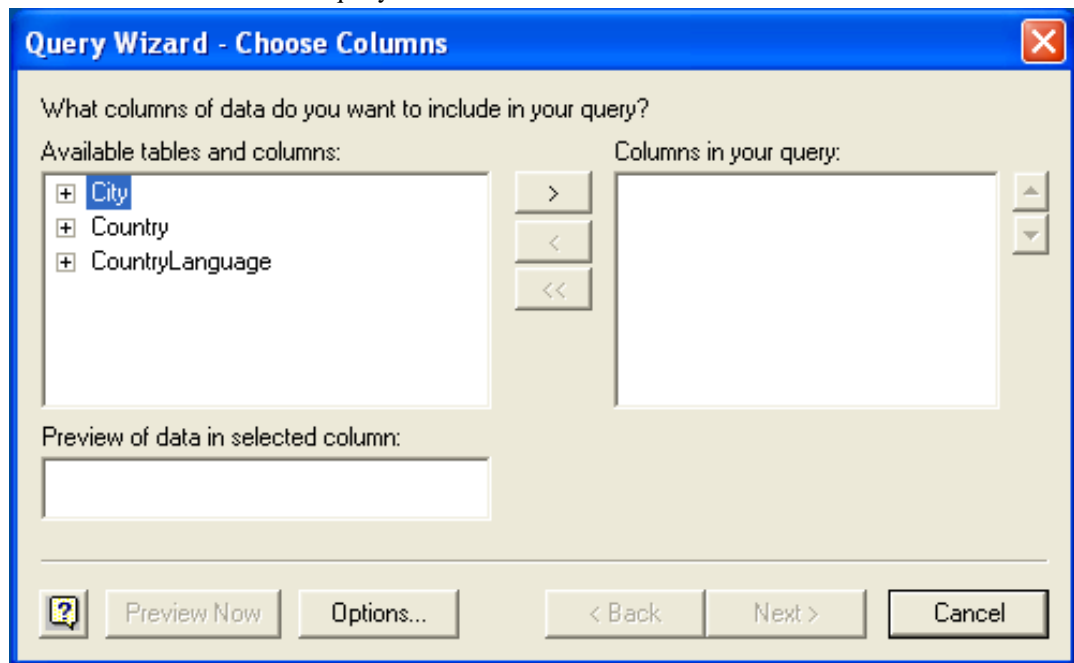
You can use Microsoft Word and Microsoft Excel to access information from a MySQL database using MyODBC. Within Microsoft Word, this facility is most useful when importing data for mailmerge, or for tables and data to be included in reports. Within Microsoft Excel, you can execute queries on your MySQL server and import the data directly into an Excel Worksheet, presenting the data as a series of rows and columns.

With both applications, data is accessed and imported into the application using Microsoft Query , which enables you to execute a query though an ODBC source. You use Microsoft Query to build the SQL statement to be executed, selecting the tables, fields, selection criteria and sort order. For example, to insert information from a table in the World test database into an Excel spreadsheet, using the DSN samples shown in [Section 18.1.3, “MyODBC Configuration”](#):

1. Create a new Worksheet.
2. From the **Data** menu, choose **Import External Data**, and then select **New Database Query**.
3. Microsoft Query will start. First, you need to choose the data source, by selecting an existing Data Source Name.




4. Within the **Query Wizard**, you must choose the columns that you want to import. The list of tables available to the user configured through the DSN is shown on the left, the columns that will be added to your query are shown on the right. The columns you choose are equivalent to those in the first section of a **SELECT** query. Click **Next** to continue.



5. You can filter rows from the query (the equivalent of a **WHERE** clause) using the **Filter Data** dialog. Click **Next** to continue.

6. Select an (optional) sort order for the data. This is equivalent to using a `ORDER BY` clause in your SQL query. You can select up to three fields for sorting the information returned by the query. Click [Next](#) to continue.

7. Select the destination for your query. You can select to return the data Microsoft Excel, where you can choose a worksheet and cell where the data will be inserted; you can continue to view the query and results within Microsoft Query, where you can edit the SQL query and further filter and sort the information returned; or you can create an OLAP Cube from the query, which can then be used directly within Microsoft Excel. Click [Finish](#).



The same process can be used to import data into a Word document, where the data will be inserted as a table. This can be used for mail merge purposes (where the field data is read from a Word table), or where you want to include data and reports within a report or other document.

### 18.1.4.3.3. Using MyODBC and Microsoft Access

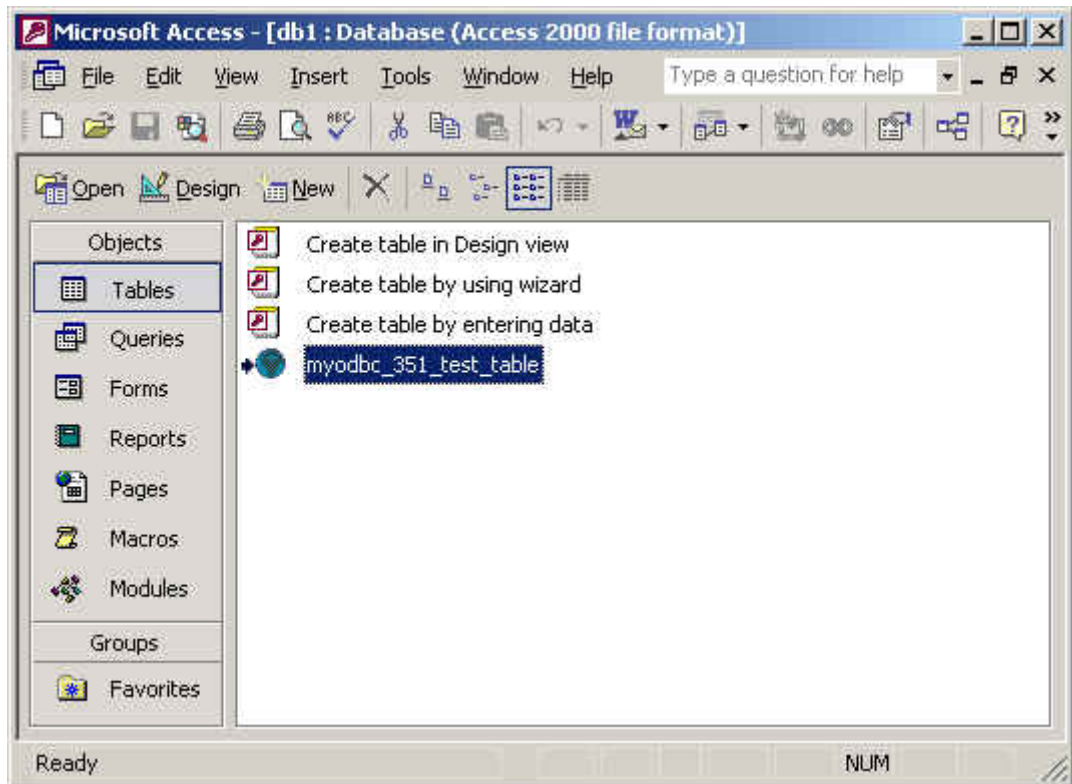
You can use MySQL database with Microsoft Access using MyODBC. The MySQL database can be used as an import source, an export source, or as a linked table for direct use within an Access application, so you can use Access as the front-end interface to a MySQL database.

#### 18.1.4.3.3.1. Exporting Access Data to MySQL

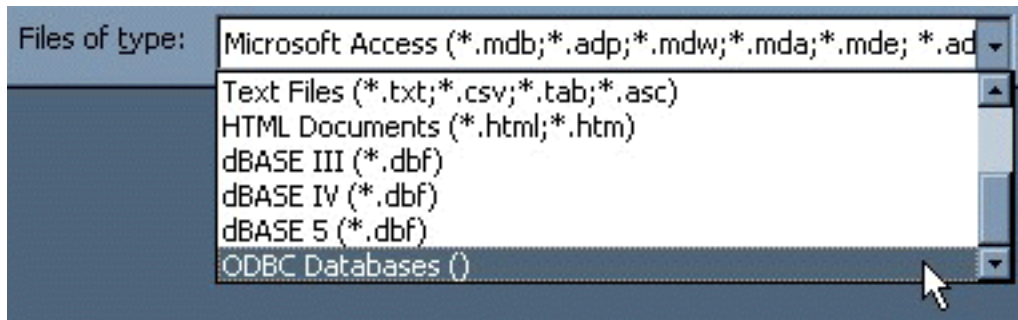
To export a table of data from an Access database to MySQL, follow these instructions:

1. When you open an Access database or an Access project, a Database window appears. It displays shortcuts for creating new database objects and opening existing objects.





2. Click the name of the [table](#) or [query](#) you want to export, and then in the [File](#) menu, select [Export](#).
3. In the [Export Object Type Object name To](#) dialog box, in the [Save As Type](#) box, select [ODBC Databases \(\)](#) as shown here:



4. In the [Export](#) dialog box, enter a name for the file (or use the suggested name), and then select [OK](#).
5. The [Select Data Source](#) dialog box is displayed; it lists the defined data sources for any ODBC drivers installed on your computer. Click either the [File Data Source](#) or [Machine Data Source](#) tab, and then double-click the [MyODBC](#) or [MyODBC 3.51](#) data source that you want to export to. To define a new data source for [MyODBC](#), please [Section 18.1.3.2, "Configuring a MyODBC DSN on Windows"](#).

Microsoft Access connects to the MySQL Server through this data source and exports new tables and or data.

#### 18.1.4.3.3.2. Importing MySQL Data to Access

To import or link a table or tables from MySQL to Access, follow these instructions:

1. Open a database, or switch to the Database window for the open database.
2. To import tables, on the **File** menu, point to **Get External Data**, and then click **Import**. To link tables, on the **File** menu, point to **Get External Data**, and then click **Link Tables**.
3. In the **Import** (or **Link**) dialog box, in the **Files Of Type** box, select **ODBC Databases ( )**. The **Select Data Source** dialog box lists the defined data sources. The **Select Data Source** dialog box is displayed; it lists the defined data source names.
4. If the ODBC data source that you selected requires you to log on, enter your login ID and password (additional information might also be required), and then click **OK**.
5. Microsoft Access connects to the MySQL server through **ODBC data source** and displays the list of tables that you can **import** or **link**.
6. Click each table that you want to **import** or **link**, and then click **OK**. If you're linking a table and it doesn't have an index that uniquely identifies each record, Microsoft Access displays a list of the fields in the linked table. Click a field or a combination of fields that uniquely identifies each record, and then click **OK**.

#### 18.1.4.3.3.3. Linking MySQL Data to Access Tables

Use the following procedure to view or to refresh links when the structure or location of a linked table has changed. The **Linked Table Manager** lists the paths to all currently linked tables.

**To view or refresh links:**

1. Open the database that contains links to tables.
2. On the **Tools** menu, point to **Add-ins (Database Utilities** in Access 2000 or newer), and then click **Linked Table Manager**.
3. Select the check box for the tables whose links you want to refresh.
4. Click **OK** to refresh the links.

Microsoft Access confirms a successful refresh or, if the table wasn't found, displays the **Select New Location of <table name>** dialog box in which you can specify its the table's new location. If several selected tables have moved to the new location that you specify, the **Linked Table Manager** searches that location for all selected tables, and updates all links in one step.

**To change the path for a set of linked tables:**

1. Open the database that contains links to tables.
2. On the **Tools** menu, point to **Add-ins (Database Utilities** in Access 2000 or newer), and then click **Linked Table Manager**.
3. Select the **Always Prompt For A New Location** check box.
4. Select the check box for the tables whose links you want to change, and then click **OK**.

5. In the **Select New Location of** <table name> dialog box, specify the new location, click **Open**, and then click **OK**.

#### 18.1.4.4. MyODBC Programming Examples

With a suitable ODBC Manager and the my MyODBC driver installed, any programming language or environment that can support ODBC should be able to connect to a MySQL database through MyODBC.

This includes, but is certainly not limited to, Microsoft support languages (including Visual Basic, C# and interfaces such as ODBC.NET), Perl (through the DBI module, and the DBD::ODBC driver).

##### 18.1.4.4.1. Using MyODBC with Visual Basic Using ADO, DAO and RDO

This section contains simple examples of the use of MySQL ODBC 3.51 Driver with ADO, DAO and RDO.

###### 18.1.4.4.1.1. ADO: `rs.addNew`, `rs.delete`, and `rs.update`

The following ADO (ActiveX Data Objects) example creates a table `my_ado` and demonstrates the use of `rs.addNew`, `rs.delete`, and `rs.update`.

```
Private Sub myodbc_ado_Click()

Dim conn As ADODB.Connection
Dim rs As ADODB.Recordset
Dim fld As ADODB.Field
Dim sql As String

'connect to MySQL server using MySQL ODBC 3.51 Driver
Set conn = New ADODB.Connection
conn.ConnectionString = "DRIVER={MySQL ODBC 3.51 Driver};"_
& "SERVER=localhost;"_
& " DATABASE=test;"_
& "UID=venu;PWD=venu; OPTION=3"

conn.Open

'create table
conn.Execute "DROP TABLE IF EXISTS my_ado"
conn.Execute "CREATE TABLE my_ado(id int not null primary key, name varchar(20)," _
& "txt text, dt date, tm time, ts timestamp)"

'direct insert
conn.Execute "INSERT INTO my_ado(id,name,txt) values(1,100,'venu')"
conn.Execute "INSERT INTO my_ado(id,name,txt) values(2,200,'MySQL')"
conn.Execute "INSERT INTO my_ado(id,name,txt) values(3,300,'Delete')"

Set rs = New ADODB.Recordset
rs.CursorLocation = adUseServer

'fetch the initial table ..
rs.Open "SELECT * FROM my_ado", conn
Debug.Print rs.RecordCount
rs.MoveFirst
Debug.Print String(50, "-") & "Initial my_ado Result Set " & String(50, "-")
For Each fld In rs.Fields
Debug.Print fld.Name,
Next
Debug.Print

Do Until rs.EOF
For Each fld In rs.Fields
Debug.Print fld.Value,
Next
rs.MoveNext
Debug.Print
Loop
rs.Close

'rs insert
```

```

rs.Open "select * from my_ado", conn, adOpenDynamic, adLockOptimistic
rs.AddNew
rs!Name = "Monty"
rs!txt = "Insert row"
rs.Update
rs.Close

'rs update
rs.Open "SELECT * FROM my_ado"
rs!Name = "update"
rs!txt = "updated-row"
rs.Update
rs.Close

'rs update second time..
rs.Open "SELECT * FROM my_ado"
rs!Name = "update"
rs!txt = "updated-second-time"
rs.Update
rs.Close

'rs delete
rs.Open "SELECT * FROM my_ado"
rs.MoveNext
rs.MoveNext
rs.Delete
rs.Close

'fetch the updated table ..
rs.Open "SELECT * FROM my_ado", conn
Debug.Print rs.RecordCount
rs.MoveFirst
Debug.Print String(50, "-") & "Updated my_ado Result Set " & String(50, "-")
For Each fld In rs.Fields
Debug.Print fld.Name,
Next
Debug.Print

Do Until rs.EOF
For Each fld In rs.Fields
Debug.Print fld.Value,
Next
rs.MoveNext
Debug.Print
Loop
rs.Close
conn.Close
End Sub

```

#### 18.1.4.4.1.2. DAO: `rs.addNew`, `rs.update`, and Scrolling

The following DAO (Data Access Objects) example creates a table `my_dao` and demonstrates the use of `rs.addNew`, `rs.update`, and result set scrolling.

```

Private Sub myodbc_dao_Click()

Dim ws As Workspace
Dim conn As Connection
Dim queryDef As queryDef
Dim str As String

'connect to MySQL using MySQL ODBC 3.51 Driver
Set ws = DBEngine.CreateWorkspace("", "venu", "venu", dbUseODBC)
str = "odbc;DRIVER={MySQL ODBC 3.51 Driver};"_
& "SERVER=localhost;"_
& " DATABASE=test;"_
& "UID=venu;PWD=venu; OPTION=3"
Set conn = ws.OpenConnection("test", dbDriverNoPrompt, False, str)

'Create table my_dao
Set queryDef = conn.CreateQueryDef("", "drop table if exists my_dao")
queryDef.Execute

Set queryDef = conn.CreateQueryDef("", "create table my_dao(Id INT AUTO_INCREMENT PRIMARY KEY, " _
& "Ts TIMESTAMP(14) NOT NULL, Name varchar(20), Id2 INT)")
queryDef.Execute

```

```

'Insert new records using rs.addNew
Set rs = conn.OpenRecordset("my_dao")
Dim i As Integer

For i = 10 To 15
rs.AddNew
rs!Name = "insert record" & i
rs!Id2 = i
rs.Update
Next i
rs.Close

'rs update..
Set rs = conn.OpenRecordset("my_dao")
rs.Edit
rs!Name = "updated-string"
rs.Update
rs.Close

'fetch the table back...
Set rs = conn.OpenRecordset("my_dao", dbOpenDynamic)
str = "Results:"
rs.MoveFirst
While Not rs.EOF
str = " " & rs!Id & " , " & rs!Name & " , " & rs!Ts & " , " & rs!Id2
Debug.Print "DATA:" & str
rs.MoveNext
Wend

'rs Scrolling
rs.MoveFirst
str = " FIRST ROW: " & rs!Id & " , " & rs!Name & " , " & rs!Ts & " , " & rs!Id2
Debug.Print str

rs.MoveLast
str = " LAST ROW: " & rs!Id & " , " & rs!Name & " , " & rs!Ts & " , " & rs!Id2
Debug.Print str

rs.MovePrevious
str = " LAST-1 ROW: " & rs!Id & " , " & rs!Name & " , " & rs!Ts & " , " & rs!Id2
Debug.Print str

'free all resources
rs.Close
queryDef.Close
conn.Close
ws.Close

End Sub

```

#### 18.1.4.4.1.3. RDO: `rs.addNew` and `rs.update`

The following RDO (Remote Data Objects) example creates a table `my_rdo` and demonstrates the use of `rs.addNew` and `rs.update`.

```

Dim rs As rdoResultset
Dim cn As New rdoConnection
Dim cl As rdoColumn
Dim SQL As String

'cn.Connect = "DSN=test;"
cn.Connect = "DRIVER={MySQL ODBC 3.51 Driver};"_
& "SERVER=localhost;"_
& " DATABASE=test;"_
& "UID=venu;PWD=venu; OPTION=3"

cn.CursorDriver = rdUseOdbc
cn.EstablishConnection rdDriverPrompt

'drop table my_rdo
SQL = "drop table if exists my_rdo"
cn.Execute SQL, rdExecDirect

'create table my_rdo
SQL = "create table my_rdo(id int, name varchar(20))"
cn.Execute SQL, rdExecDirect

```

```

'insert - direct
SQL = "insert into my_rdo values (100,'venu')"
cn.Execute SQL, rdExecDirect

SQL = "insert into my_rdo values (200,'MySQL')"
cn.Execute SQL, rdExecDirect

'rs insert
SQL = "select * from my_rdo"
Set rs = cn.OpenResultset(SQL, rdOpenStatic, rdConcurRowVer, rdExecDirect)
rs.AddNew
rs!id = 300
rs!Name = "Insert1"
rs.Update
rs.Close

'rs insert
SQL = "select * from my_rdo"
Set rs = cn.OpenResultset(SQL, rdOpenStatic, rdConcurRowVer, rdExecDirect)
rs.AddNew
rs!id = 400
rs!Name = "Insert 2"
rs.Update
rs.Close

'rs update
SQL = "select * from my_rdo"
Set rs = cn.OpenResultset(SQL, rdOpenStatic, rdConcurRowVer, rdExecDirect)
rs.Edit
rs!id = 999
rs!Name = "updated"
rs.Update
rs.Close

'fetch back...
SQL = "select * from my_rdo"
Set rs = cn.OpenResultset(SQL, rdOpenStatic, rdConcurRowVer, rdExecDirect)
Do Until rs.EOF
For Each cl In rs.rdoColumns
Debug.Print cl.Value,
Next
rs.MoveNext
Debug.Print
Loop
Debug.Print "Row count="; rs.RowCount

'close
rs.Close
cn.Close

End Sub

```

#### 18.1.4.4.2. Using MyODBC with .NET

This section contains simple examples that demonstrate the use of MyODBC drivers with ODBC.NET.

##### 18.1.4.4.2.1. Using MyODBC with ODBC.NET and C# (C sharp)

The following sample creates a table `my_odbc_net` and demonstrates its use in C#.

```

/**
 * @sample      : mycon.cs
 * @purpose     : Demo sample for ODBC.NET using MyODBC
 * @author      : Venu, <myodbc@lists.mysql.com>
 *
 * (C) Copyright MySQL AB, 1995-2006
 *
 **/

/* build command
 *
 * csc /t:exe
 *      /out:mycon.exe mycon.cs
 *      /r:Microsoft.Data.Odbc.dll
 */

```

```
using Console = System.Console;
using Microsoft.Data.Odbc;

namespace myodbc3
{
    class mycon
    {
        static void Main(string[] args)
        {
            try
            {
                //Connection string for MyODBC 2.50
                /*string MyConString = "DRIVER={MySQL};" +
                    "SERVER=localhost;" +
                    "DATABASE=test;" +
                    "UID=venu;" +
                    "PASSWORD=venu;" +
                    "OPTION=3";
                */
                //Connection string for MyODBC 3.51
                string MyConString = "DRIVER={MySQL ODBC 3.51 Driver};" +
                    "SERVER=localhost;" +
                    "DATABASE=test;" +
                    "UID=venu;" +
                    "PASSWORD=venu;" +
                    "OPTION=3";

                //Connect to MySQL using MyODBC
                OdbcConnection MyConnection = new OdbcConnection(MyConString);
                MyConnection.Open();

                Console.WriteLine("\n !!! success, connected successfully !!!\n");

                //Display connection information
                Console.WriteLine("Connection Information:");
                Console.WriteLine("\tConnection String:" +
                    MyConnection.ConnectionString);
                Console.WriteLine("\tConnection Timeout:" +
                    MyConnection.ConnectionTimeout);
                Console.WriteLine("\tDatabase:" +
                    MyConnection.Database);
                Console.WriteLine("\tDataSource:" +
                    MyConnection.DataSource);
                Console.WriteLine("\tDriver:" +
                    MyConnection.Driver);
                Console.WriteLine("\tServerVersion:" +
                    MyConnection.ServerVersion);

                //Create a sample table
                OdbcCommand MyCommand =
                    new OdbcCommand("DROP TABLE IF EXISTS my_odbc_net",
                        MyConnection);
                MyCommand.ExecuteNonQuery();
                MyCommand.CommandText =
                    "CREATE TABLE my_odbc_net(id int, name varchar(20), idb bigint)";
                MyCommand.ExecuteNonQuery();

                //Insert
                MyCommand.CommandText =
                    "INSERT INTO my_odbc_net VALUES(10,'venu', 300)";
                Console.WriteLine("INSERT, Total rows affected:" +
                    MyCommand.ExecuteNonQuery());

                //Insert
                MyCommand.CommandText =
                    "INSERT INTO my_odbc_net VALUES(20,'mysql',400)";
                Console.WriteLine("INSERT, Total rows affected:" +
                    MyCommand.ExecuteNonQuery());

                //Insert
                MyCommand.CommandText =
                    "INSERT INTO my_odbc_net VALUES(20,'mysql',500)";
                Console.WriteLine("INSERT, Total rows affected:" +
                    MyCommand.ExecuteNonQuery());

                //Update
                MyCommand.CommandText =
                    "UPDATE my_odbc_net SET id=999 WHERE id=20";
                Console.WriteLine("Update, Total rows affected:" +
                    MyCommand.ExecuteNonQuery());
            }
        }
    }
}
```

```

//COUNT(*)
MyCommand.CommandText =
    "SELECT COUNT(*) as TRows FROM my_odbc_net";
Console.WriteLine("Total Rows:" +
    MyCommand.ExecuteScalar());

//Fetch
MyCommand.CommandText = "SELECT * FROM my_odbc_net";
OdbcDataReader MyDataReader;
MyDataReader = MyCommand.ExecuteReader();
while (MyDataReader.Read())
{
    if(string.Compare(MyConnection.Driver,"myodbc3.dll") == 0) {
        //Supported only by MyODBC 3.51
        Console.WriteLine("Data:" + MyDataReader.GetInt32(0) + " " +
            MyDataReader.GetString(1) + " " +
            MyDataReader.GetInt64(2));
    }
    else {
        //BIGINTs not supported by MyODBC
        Console.WriteLine("Data:" + MyDataReader.GetInt32(0) + " " +
            MyDataReader.GetString(1) + " " +
            MyDataReader.GetInt32(2));
    }
}

//Close all resources
MyDataReader.Close();
MyConnection.Close();
}
catch (OdbcException MyOdbcException) //Catch any ODBC exception ..
{
    for (int i=0; i < MyOdbcException.Errors.Count; i++)
    {
        Console.Write("ERROR #" + i + "\n" +
            "Message: " +
            MyOdbcException.Errors[i].Message + "\n" +
            "Native: " +
            MyOdbcException.Errors[i].NativeError.ToString() + "\n" +
            "Source: " +
            MyOdbcException.Errors[i].Source + "\n" +
            "SQL: " +
            MyOdbcException.Errors[i].SQLState + "\n");
    }
}
}
}
}
}

```

#### 18.1.4.4.2.2. Using MyODBC with ODBC.NET and Visual Basic

The following sample creates a table `my_vb_net` and demonstrates the use in VB.

```

' @sample      : myvb.vb
' @purpose     : Demo sample for ODBC.NET using MyODBC
' @author      : Venu, <myodbc@lists.mysql.com>
'
' (C) Copyright MySQL AB, 1995-2006
'
'
'
' build command
'
' vbc /target:exe
'     /out:myvb.exe
'     /r:Microsoft.Data.Odbc.dll
'     /r:System.dll
'     /r:System.Data.dll
'
Imports Microsoft.Data.Odbc
Imports System

Module myvb
    Sub Main()
        Try

```



```

'MyODBC 3.51 connection string
Dim MyConString As String = "DRIVER={MySQL ODBC 3.51 Driver};" & _
"SERVER=localhost;" & _
"DATABASE=test;" & _
"UID=venu;" & _
"PASSWORD=venu;" & _
"OPTION=3;"

'Connection
Dim MyConnection As New OdbcConnection(MyConString)
MyConnection.Open()

Console.WriteLine("Connection State::" & MyConnection.State.ToString)

'Drop
Console.WriteLine("Dropping table")
Dim MyCommand As New OdbcCommand()
MyCommand.Connection = MyConnection
MyCommand.CommandText = "DROP TABLE IF EXISTS my_vb_net"
MyCommand.ExecuteNonQuery()

'Create
Console.WriteLine("Creating...")
MyCommand.CommandText = "CREATE TABLE my_vb_net(id int, name varchar(30))"
MyCommand.ExecuteNonQuery()

'Insert
MyCommand.CommandText = "INSERT INTO my_vb_net VALUES(10,'venu')"
Console.WriteLine("INSERT, Total rows affected:" & _
MyCommand.ExecuteNonQuery())

'Insert
MyCommand.CommandText = "INSERT INTO my_vb_net VALUES(20,'mysql')"
Console.WriteLine("INSERT, Total rows affected:" & _
MyCommand.ExecuteNonQuery())

'Insert
MyCommand.CommandText = "INSERT INTO my_vb_net VALUES(20,'mysql')"
Console.WriteLine("INSERT, Total rows affected:" & _
MyCommand.ExecuteNonQuery())

'Insert
MyCommand.CommandText = "INSERT INTO my_vb_net(id) VALUES(30)"
Console.WriteLine("INSERT, Total rows affected:" & _
MyCommand.ExecuteNonQuery())

'Update
MyCommand.CommandText = "UPDATE my_vb_net SET id=999 WHERE id=20"
Console.WriteLine("Update, Total rows affected:" & _
MyCommand.ExecuteNonQuery())

'COUNT(*)
MyCommand.CommandText = "SELECT COUNT(*) as TRows FROM my_vb_net"
Console.WriteLine("Total Rows:" & MyCommand.ExecuteScalar())

'Select
Console.WriteLine("Select * FROM my_vb_net")
MyCommand.CommandText = "SELECT * FROM my_vb_net"
Dim MyDataReader As OdbcDataReader
MyDataReader = MyCommand.ExecuteReader
While MyDataReader.Read
    If MyDataReader("name") Is DBNull.Value Then
        Console.WriteLine("id = " & _
CStr(MyDataReader("id")) & " name = " & _
"NULL")
    Else
        Console.WriteLine("id = " & _
CStr(MyDataReader("id")) & " name = " & _
CStr(MyDataReader("name")))
    End If
End While

'Catch ODBC Exception
Catch MyOdbcException As OdbcException
Dim i As Integer
Console.WriteLine(MyOdbcException.ToString)

'Catch program exception
Catch MyException As Exception
Console.WriteLine(MyException.ToString)
End Try

```

## 18.1.5. MyODBC Reference

This section provides reference material for the MyODBC API, showing supported functions and methods, supported MySQL column types and the corresponding native type in MyODBC, and the error codes returned by MyODBC when a fault occurs.

### 18.1.5.1. MyODBC API Reference

This section summarizes ODBC routines, categorized by functionality.

For the complete ODBC API reference, please refer to the ODBC Programmer's Reference at [http://msdn.microsoft.com/library/en-us/odbc/hdm/odbcabout\\_this\\_manual.asp](http://msdn.microsoft.com/library/en-us/odbc/hdm/odbcabout_this_manual.asp).

An application can call `SQLGetInfo` function to obtain conformance information about MyODBC. To obtain information about support for a specific function in the driver, an application can call `SQLGetFunctions`.

Note: For backward compatibility, the MyODBC 3.51 driver supports all deprecated functions.

The following tables list MyODBC API calls grouped by task:

#### Connecting to a data source:

Function name	MyODBC		Standard	Purpose
	2.50	3.51		
<code>SQLAllocHandle</code>	No	Yes	ISO 92	Obtains an environment, connection, statement, or descriptor handle.
<code>SQLConnect</code>	Yes	Yes	ISO 92	Connects to a specific driver by data source name, user ID, and password.
<code>SQLDriverConnect</code>	Yes	Yes	ODBC	Connects to a specific driver by connection string or requests that the Driver Manager and driver display connection dialog boxes for the user.
<code>SQLAllocEnv</code>	Yes	Yes	Deprecated	Obtains an environment handle allocated from driver.
<code>SQLAllocConnect</code>	Yes	Yes	Deprecated	Obtains a connection handle

#### Obtaining information about a driver and data source:

Function name	MyODBC		Standard	Purpose
	2.50	3.51		
<code>SQLDataSources</code>	No	No	ISO 92	Returns the list of available data sources, handled by the Driver Manager
<code>SQLDrivers</code>	No	No	ODBC	Returns the list of installed drivers and their attributes, handles by Driver Manager
<code>SQLGetInfo</code>	Yes	Yes	ISO 92	Returns information about a specific driver and data source.
<code>SQLGetFunctions</code>	Yes	Yes	ISO 92	Returns supported driver functions.
<code>SQLGetTypeInfo</code>	Yes	Yes	ISO 92	Returns information about supported data types.

**Setting and retrieving driver attributes:**

Function name	MyODBC		Standard	Purpose
	2.50	3.51		
<a href="#">SQLSetConnectAttr</a>	No	Yes	ISO 92	Sets a connection attribute.
<a href="#">SQLGetConnectAttr</a>	No	Yes	ISO 92	Returns the value of a connection attribute.
<a href="#">SQLSetConnectOption</a>	Yes	Yes	Deprecated	Sets a connection option
<a href="#">SQLGetConnectOption</a>	Yes	Yes	Deprecated	Returns the value of a connection option
<a href="#">SQLSetEnvAttr</a>	No	Yes	ISO 92	Sets an environment attribute.
<a href="#">SQLGetEnvAttr</a>	No	Yes	ISO 92	Returns the value of an environment attribute.
<a href="#">SQLSetStmtAttr</a>	No	Yes	ISO 92	Sets a statement attribute.
<a href="#">SQLGetStmtAttr</a>	No	Yes	ISO 92	Returns the value of a statement attribute.
<a href="#">SQLSetStmtOption</a>	Yes	Yes	Deprecated	Sets a statement option
<a href="#">SQLGetStmtOption</a>	Yes	Yes	Deprecated	Returns the value of a statement option

**Preparing SQL requests:**

Function name	MyODBC		Standard	Purpose
	2.50	3.51		
<a href="#">SQLAllocStmt</a>	Yes	Yes	Deprecated	Allocates a statement handle
<a href="#">SQLPrepare</a>	Yes	Yes	ISO 92	Prepares an SQL statement for later execution.
<a href="#">SQLBindParameter</a>	Yes	Yes	ODBC	Assigns storage for a parameter in an SQL statement.
<a href="#">SQLGetCursorName</a>	Yes	Yes	ISO 92	Returns the cursor name associated with a statement handle.
<a href="#">SQLSetCursorName</a>	Yes	Yes	ISO 92	Specifies a cursor name.
<a href="#">SQLSetScrollOptions</a>	Yes	Yes	ODBC	Sets options that control cursor behavior.

**Submitting requests:**

Function name	MyODBC		Standard	Purpose
	2.50	3.51		
<a href="#">SQLExecute</a>	Yes	Yes	ISO 92	Executes a prepared statement.
<a href="#">SQLExecDirect</a>	Yes	Yes	ISO 92	Executes a statement
<a href="#">SQLNativeSql</a>	Yes	Yes	ODBC	Returns the text of an SQL statement as translated by the driver.
<a href="#">SQLDescribeParam</a>	Yes	Yes	ODBC	Returns the description for a specific parameter in a statement.
<a href="#">SQLNumParams</a>	Yes	Yes	ISO 92	Returns the number of parameters in a statement.
<a href="#">SQLParamData</a>	Yes	Yes	ISO 92	Used in conjunction with <a href="#">SQLPutData</a> to supply parameter data at execution time. (Useful for long data values.)

SQLPutData	Yes	Yes	ISO 92	Sends part or all of a data value for a parameter. (Useful for long data values.)
------------	-----	-----	--------	---

**Retrieving results and information about results:**

Function name	MyODBC		Standard	Purpose
	2.50	3.51		
SQLRowCount	Yes	Yes	ISO 92	Returns the number of rows affected by an insert, update, or delete request.
SQLNumResultCols	Yes	Yes	ISO 92	Returns the number of columns in the result set.
SQLDescribeCol	Yes	Yes	ISO 92	Describes a column in the result set.
SQLColAttribute	No	Yes	ISO 92	Describes attributes of a column in the result set.
SQLColAttributes	Yes	Yes	Deprecated	Describes attributes of a column in the result set.
SQLFetch	Yes	Yes	ISO 92	Returns multiple result rows.
SQLFetchScroll	No	Yes	ISO 92	Returns scrollable result rows.
SQLExtendedFetch	Yes	Yes	Deprecated	Returns scrollable result rows.
SQLSetPos	Yes	Yes	ODBC	Positions a cursor within a fetched block of data and allows an application to refresh data in the rowset or to update or delete data in the result set.
SQLBulkOperations	No	Yes	ODBC	Performs bulk insertions and bulk bookmark operations, including update, delete, and fetch by bookmark.

**Retrieving error or diagnostic information:**

Function name	MyODBC		Standard	Purpose
	2.50	3.51		
SQLError	Yes	Yes	Deprecated	Returns additional error or status information
SQLGetDiagField	Yes	Yes	ISO 92	Returns additional diagnostic information (a single field of the diagnostic data structure).
SQLGetDiagRec	Yes	Yes	ISO 92	Returns additional diagnostic information (multiple fields of the diagnostic data structure).

**Obtaining information about the data source's system tables (catalog functions) item:**

Function name	MyODBC		Standard	Purpose
	2.50	3.51		
SQLColumnPrivileges	Yes	Yes	ODBC	Returns a list of columns and associated privileges for one or more tables.
SQLColumns	Yes	Yes	X/Open	Returns the list of column names in specified tables.
SQLForeignKeys	Yes	Yes	ODBC	Returns a list of column names that make up foreign keys, if they exist for a specified table.
SQLPrimaryKeys	Yes	Yes	ODBC	Returns the list of column names that make up the primary key for a table.
SQLSpecialColumns	Yes	Yes	X/Open	Returns information about the optimal set of

				columns that uniquely identifies a row in a specified table, or the columns that are automatically updated when any value in the row is updated by a transaction.
<code>SQLStatistics</code>	Yes	Yes	ISO 92	Returns statistics about a single table and the list of indexes associated with the table.
<code>SQLTablePrivileges</code>	Yes	Yes	ODBC	Returns a list of tables and the privileges associated with each table.
<code>SQLTables</code>	Yes	Yes	X/Open	Returns the list of table names stored in a specific data source.

**Performing transactions:**

Function name	MyODBC		Standard	Purpose
	2.50	3.51		
<code>SQLTransact</code>	Yes	Yes	Deprecated	Commits or rolls back a transaction
<code>SQLEndTran</code>	No	Yes	ISO 92	Commits or rolls back a transaction.

**Terminating a statement:**

Function name	MyODBC		Standard	Purpose
	2.50	3.51		
<code>SQLFreeStmt</code>	Yes	Yes	ISO 92	Ends statement processing, discards pending results, and, optionally, frees all resources associated with the statement handle.
<code>SQLCloseCursor</code>	Yes	Yes	ISO 92	Closes a cursor that has been opened on a statement handle.
<code>SQLCancel</code>	Yes	Yes	ISO 92	Cancels an SQL statement.

**Terminating a connection:**

Function name	MyODBC		Standard	Purpose
	2.50	3.51		
<code>SQLDisconnect</code>	Yes	Yes	ISO 92	Closes the connection.
<code>SQLFreeHandle</code>	No	Yes	ISO 92	Releases an environment, connection, statement, or descriptor handle.
<code>SQLFreeConnect</code>	Yes	Yes	Deprecated	Releases connection handle
<code>SQLFreeEnv</code>	Yes	Yes	Deprecated	Releases an environment handle

## 18.1.5.2. MyODBC Data Types

The following table illustrates how driver maps the server data types to default SQL and C data types:

Native Value	SQL Type	C Type
<code>bit</code>	<code>SQL_BIT</code>	<code>SQL_C_BIT</code>
<code>tinyint</code>	<code>SQL_TINYINT</code>	<code>SQL_C_STINYINT</code>

## Connectors

tinyint unsigned	SQL_TINYINT	SQL_C_UTINYINT
bigint	SQL_BIGINT	SQL_C_SBIGINT
bigint unsigned	SQL_BIGINT	SQL_C_UBIGINT
long varbinary	SQL_LONGVARBINARY	SQL_C_BINARY
blob	SQL_LONGVARBINARY	SQL_C_BINARY
longblob	SQL_LONGVARBINARY	SQL_C_BINARY
tinyblob	SQL_LONGVARBINARY	SQL_C_BINARY
mediumblob	SQL_LONGVARBINARY	SQL_C_BINARY
long varchar	SQL_LONGVARCHAR	SQL_C_CHAR
text	SQL_LONGVARCHAR	SQL_C_CHAR
mediumtext	SQL_LONGVARCHAR	SQL_C_CHAR
char	SQL_CHAR	SQL_C_CHAR
numeric	SQL_NUMERIC	SQL_C_CHAR
decimal	SQL_DECIMAL	SQL_C_CHAR
integer	SQL_INTEGER	SQL_C_SLONG
integer unsigned	SQL_INTEGER	SQL_C_ULONG
int	SQL_INTEGER	SQL_C_SLONG
int unsigned	SQL_INTEGER	SQL_C_ULONG
mediumint	SQL_INTEGER	SQL_C_SLONG
mediumint unsigned	SQL_INTEGER	SQL_C_ULONG
smallint	SQL_SMALLINT	SQL_C_SSHORT
smallint unsigned	SQL_SMALLINT	SQL_C_USHORT
real	SQL_FLOAT	SQL_C_DOUBLE
double	SQL_FLOAT	SQL_C_DOUBLE
float	SQL_REAL	SQL_C_FLOAT
double precision	SQL_DOUBLE	SQL_C_DOUBLE
date	SQL_DATE	SQL_C_DATE
time	SQL_TIME	SQL_C_TIME
year	SQL_SMALLINT	SQL_C_SHORT
datetime	SQL_TIMESTAMP	SQL_C_TIMESTAMP
timestamp	SQL_TIMESTAMP	SQL_C_TIMESTAMP
text	SQL_VARCHAR	SQL_C_CHAR
varchar	SQL_VARCHAR	SQL_C_CHAR
enum	SQL_VARCHAR	SQL_C_CHAR
set	SQL_VARCHAR	SQL_C_CHAR
bit	SQL_CHAR	SQL_C_CHAR
bool	SQL_CHAR	SQL_C_CHAR

### 18.1.5.3. MyODBC Error Codes

The following tables lists the error codes returned by the driver apart from the server errors.

## Connectors

Native Code	SQLSTATE 2	SQLSTATE 3	Error Message
500	01000	01000	General warning
501	01004	01004	String data, right truncated
502	01S02	01S02	Option value changed
503	01S03	01S03	No rows updated/deleted
504	01S04	01S04	More than one row updated/deleted
505	01S06	01S06	Attempt to fetch before the result set returned the first row set
506	07001	07002	<a href="#">SQLBindParameter</a> not used for all parameters
507	07005	07005	Prepared statement not a cursor-specification
508	07009	07009	Invalid descriptor index
509	08002	08002	Connection name in use
510	08003	08003	Connection does not exist
511	24000	24000	Invalid cursor state
512	25000	25000	Invalid transaction state
513	25S01	25S01	Transaction state unknown
514	34000	34000	Invalid cursor name
515	S1000	HY000	General driver defined error
516	S1001	HY001	Memory allocation error
517	S1002	HY002	Invalid column number
518	S1003	HY003	Invalid application buffer type
519	S1004	HY004	Invalid SQL data type
520	S1009	HY009	Invalid use of null pointer
521	S1010	HY010	Function sequence error
522	S1011	HY011	Attribute can not be set now
523	S1012	HY012	Invalid transaction operation code
524	S1013	HY013	Memory management error
525	S1015	HY015	No cursor name available
526	S1024	HY024	Invalid attribute value
527	S1090	HY090	Invalid string or buffer length
528	S1091	HY091	Invalid descriptor field identifier
529	S1092	HY092	Invalid attribute/option identifier
530	S1093	HY093	Invalid parameter number
531	S1095	HY095	Function type out of range
532	S1106	HY106	Fetch type out of range
533	S1117	HY117	Row value out of range
534	S1109	HY109	Invalid cursor position
535	S1C00	HYC00	Optional feature not implemented
0	21S01	21S01	Column count does not match value count
0	23000	23000	Integrity constraint violation
0	42000	42000	Syntax error or access violation

0	42S02	42S02	Base table or view not found
0	42S12	42S12	Index not found
0	42S21	42S21	Column already exists
0	42S22	42S22	Column not found
0	08S01	08S01	Communication link failure

## 18.1.6. MyODBC Notes and Tips

Here are some common notes and tips for using MyODBC within different environments, applications and tools. The notes provided here are based on the experiences of MyODBC developers and users.

### 18.1.6.1. MyODBC General Functionality

This section provides help with common queries and areas of functionality in MySQL and how to use them with MyODBC.

#### 18.1.6.1.1. Obtaining Auto-Increment Values

Obtaining the value of column that uses `AUTO_INCREMENT` after an `INSERT` statement can be achieved in a number of different ways. To obtain the value immediately after an `INSERT`, use a `SELECT` query with the `LAST_INSERT_ID()` function.

For example, using MyODBC you would execute two separate statements, the `INSERT` statement and the `SELECT` query to obtain the auto-increment value.

```
INSERT INTO tbl (auto,text) VALUES(NULL,'text');
SELECT LAST_INSERT_ID();
```

If you do not require the value within your application, but do require the value as part of another `INSERT`, the entire process can be handled by executing the following statements:

```
INSERT INTO tbl (auto,text) VALUES(NULL,'text');
INSERT INTO tbl2 (id,text) VALUES(LAST_INSERT_ID(),'text');
```

Certain ODBC applications (including Delphi and Access) may have trouble obtaining the auto-increment value using the previous examples. In this case, try the following statement as an alternative:

```
SELECT * FROM tbl WHERE auto IS NULL;
```

See [Section 17.2.13.3, “How to Get the Unique ID for the Last Inserted Row”](#).

#### 18.1.6.1.2. Dynamic Cursor Support

Support for the `dynamic cursor` is provided in MyODBC 3.51, but dynamic cursors are not enabled by default. You can enable this function within Windows by selecting the `Enable Dynamic Cursor` checkbox within the ODBC Data Source Administrator.

On other platforms, you can enable the dynamic cursor by adding `32` to the `OPTION` value when creating the DSN.

#### 18.1.6.1.3. MyODBC Performance

The MyODBC driver has been optimized to provide very fast performance. If you experience problems with the performance of MyODBC, or notice a large amount of disk activity for simple queries, there are



a number of aspects you should check:

- Ensure that **ODBC Tracing** is not enabled. With tracing enabled, a lot of information is recorded in the tracing file by the ODBC Manager. You can check, and disable, tracing within Windows using the **Tracing** panel of the ODBC Data Source Administrator. Within Mac OS X, check the **Tracing** panel of ODBC Administrator. See [Section 18.1.3.8, “Getting an ODBC Trace File”](#).
- Make sure you are using the standard version of the driver, and not the debug version. The debug version includes additional checks and reporting measures.
- Disable the MyODBC driver trace and query logs. These options are enabled for each DSN, so make sure to examine only the DSN that you are using in your application. Within Windows, you can disable the MyODBC and query logs by modifying the DSN configuration. Within Mac OS X and Unix, ensure that the driver trace (option value 4) and query logging (option value 524288) are not enabled.

#### 18.1.6.1.4. Setting ODBC Query Timeout in Windows

For more information on how to set the query timeout on Microsoft Windows when executing queries through an ODBC connection, read the Microsoft knowledgebase document at <http://support.microsoft.com/default.aspx?scid=kb%3Ben-us%3B153756>.

#### 18.1.6.2. MyODBC Application Specific Tips

Most programs should work with MyODBC, but for each of those listed here, there are specific notes and tips to improve or enhance the way you work with MyODBC and these applications.

With all applications you should ensure that you are using the latest MyODBC drivers, ODBC Manager and any supporting libraries and interfaces used by your application. For example, on Windows, using the latest version of Microsoft Data Access Components (MDAC) will improve the compatibility with ODBC in general, and with the MyODBC driver.

##### 18.1.6.2.1. Using MyODBC with Microsoft Applications

The majority of Microsoft applications have been tested with MyODBC, including Microsoft Office, Microsoft Access and the various programming languages supported within ASP and Microsoft Visual Studio.

If you have problem with MyODBC and your program also works with OLEDB, you should try the OLEDB driver.

###### 18.1.6.2.1.1. Microsoft Access

To improve the integration between Microsoft Access and MySQL through MyODBC:

- For all versions of Access, you should enable the MyODBC **Return matching rows** option. For Access 2.0, you should additionally enable the **Simulate ODBC 1.0** option.
- You should have a **TIMESTAMP** column in all tables that you want to be able to update. For maximum portability, don't use a length specification in the column declaration (which is unsupported within MySQL in versions earlier than 4.1).
- You should have a primary key in each MySQL table you want to use with Access. If not, new or updated rows may show up as **#DELETED#**.

- Use only `DOUBLE` float fields. Access fails when comparing with single-precision floats. The symptom usually is that new or updated rows may show up as `#DELETED#` or that you can't find or update rows.
- If you are using MyODBC to link to a table that has a `BIGINT` column, the results are displayed as `#DELETED#`. The work around solution is:
  - Have one more dummy column with `TIMESTAMP` as the data type.
  - Select the `Change BIGINT columns to INT` option in the connection dialog in ODBC DSN Administrator.
  - Delete the table link from Access and re-create it.

Old records may still display as `#DELETED#`, but newly added/updated records are displayed properly.

- If you still get the error `Another user has changed your data` after adding a `TIMESTAMP` column, the following trick may help you:

Don't use a `table` data sheet view. Instead, create a form with the fields you want, and use that `form` data sheet view. You should set the `DefaultValue` property for the `TIMESTAMP` column to `NOW( )`. It may be a good idea to hide the `TIMESTAMP` column from view so your users are not confused.

- In some cases, Access may generate SQL statements that MySQL can't understand. You can fix this by selecting `"Query | SQLSpecific | Pass-Through"` from the Access menu.
- On Windows NT, Access reports `BLOB` columns as `OLE OBJECTS`. If you want to have `MEMO` columns instead, you should change `BLOB` columns to `TEXT` with `ALTER TABLE`.
- Access can't always handle the MySQL `DATE` column properly. If you have a problem with these, change the columns to `DATETIME`.
- If you have in Access a column defined as `BYTE`, Access tries to export this as `TINYINT` instead of `TINYINT UNSIGNED`. This gives you problems if you have values larger than 127 in the column.
- If you have very large (long) tables in Access, it might take a very long time to open them. Or you might run low on virtual memory and eventually get an `ODBC Query Failed` error and the table cannot open. To deal with this, select the following options:
  - Return Matching Rows (2)
  - Allow BIG Results (8).

These add up to a value of 10 (`OPTION=10`).

Some external articles and tips that may be useful when using Access, ODBC and MyODBC:

- Read [How to Trap ODBC Login Error Messages in Access](http://support.microsoft.com/support/kb/articles/Q124/9/01.asp?LN=EN-US&SD=gn&FR=0%3CP%3E)  
[http://support.microsoft.com/support/kb/articles/Q124/9/01.asp?LN=EN-US&SD=gn&FR=0%3CP%3E]
- [Optimizing Access ODBC Applications](http://support.microsoft.com/default.aspx?scid=kb;en-us;128808)
  - [Optimizing for Client/Server Performance](http://support.microsoft.com/default.aspx?scid=kb;en-us;128808)  
[http://support.microsoft.com/default.aspx?scid=kb;en-us;128808]

- [Tips for Converting Applications to Using ODBCDirect](http://support.microsoft.com/default.aspx?scid=kb;en-us;164481)  
[http://support.microsoft.com/default.aspx?scid=kb;en-us;164481]
- [Tips for Optimizing Queries on Attached SQL Tables](http://support.microsoft.com/default.aspx?scid=kb;EN-US;q99321)  
[http://support.microsoft.com/default.aspx?scid=kb;EN-US;q99321]
- For a list of tools that can be used with Access and ODBC data sources, refer to [converters](#)  
[http://www.mysql.com/portal/software/convertors/] section for list of available tools.

#### 18.1.6.2.1.2. Microsoft Excel and Column Types

If you have problems importing data into Microsoft Excel, particularly numerical, date, and time values, this is probably because of a bug in Excel, where the column type of the source data is used to determine the data type when that data is inserted into a cell within the worksheet. The result is that Excel incorrectly identifies the content and this affects both the display format and the data when it is used within calculations.

To address this issue, use the `CONCAT()` function in your queries. The use of `CONCAT()` forces Excel to treat the value as a string, which Excel will then parse and usually correctly identify the embedded information.

However, even with this option, some data may be incorrectly formatted, even though the source data remains unchanged. Use the `Format Cells` option within Excel to change the format of the displayed information.

#### 18.1.6.2.1.3. Microsoft Visual Basic

To be able to update a table, you must define a primary key for the table.

Visual Basic with ADO can't handle big integers. This means that some queries like `SHOW PROCESSLIST` do not work properly. The fix is to use `OPTION=16384` in the ODBC connect string or to select the `Change BIGINT columns to INT` option in the MyODBC connect screen. You may also want to select the `Return matching rows` option.

#### 18.1.6.2.1.4. Microsoft Visual InterDev

If you have a `BIGINT` in your result, you may get the error `[Microsoft][ODBC Driver Manager] Driver does not support this parameter`. Try selecting the `Change BIGINT columns to INT` option in the MyODBC connect screen.

#### 18.1.6.2.1.5. Visual Objects

You should select the `Don't optimize column widths` option.

#### 18.1.6.2.1.6. Microsoft ADO

When you are coding with the ADO API and MyODBC, you need to pay attention to some default properties that aren't supported by the MySQL server. For example, using the `CursorLocation Property` as `adUseServer` returns a result of `-1` for the `RecordCount Property`. To have the right value, you need to set this property to `adUseClient`, as shown in the VB code here:

```
Dim myconn As New ADOConnection
Dim myrs As New Recordset
Dim mySQL As String
Dim myrows As Long

myconn.Open "DSN=MyODBCsample"
mySQL = "SELECT * from user"
myrs.Source = mySQL
Set myrs.ActiveConnection = myconn
```

```
myrs.CursorLocation = adUseClient
myrs.Open
myrows = myrs.RecordCount

myrs.Close
myconn.Close
```

Another workaround is to use a `SELECT COUNT( *)` statement for a similar query to get the correct row count.

To find the number of rows affected by a specific SQL statement in ADO, use the `RecordsAffected` property in the ADO execute method. For more information on the usage of execute method, refer to <http://msdn.microsoft.com/library/default.asp?url=/library/en-us/ado270/htm/mdmthcnexecute.asp>.

For information, see [ActiveX Data Objects\(ADO\) Frequently Asked Questions](http://support.microsoft.com/default.aspx?scid=kb;EN-US;q183606) [http://support.microsoft.com/default.aspx?scid=kb;EN-US;q183606].

#### 18.1.6.2.1.7. Using MyODBC with Active Server Pages (ASP)

You should select the `Return matching rows` option in the DSN.

For more information about how to access MySQL via ASP using MyODBC, refer to the following articles:

- [Using MyODBC To Access Your MySQL Database Via ASP](http://www.devarticles.com/c/a/ASP/Using-MyODBC-To-Access-Your-MySQL-Database-Via-ASP/) [http://www.devarticles.com/c/a/ASP/Using-MyODBC-To-Access-Your-MySQL-Database-Via-ASP/]
- [ASP and MySQL at DWAM.NT](http://www.dwam.net/mysql/asp_myodbc.asp) [http://www.dwam.net/mysql/asp\_myodbc.asp]

A Frequently Asked Questions list for ASP can be found at <http://support.microsoft.com/default.aspx?scid=/Support/ActiveServer/faq/data/adofaq.asp>.

#### 18.1.6.2.1.8. Using MyODBC with Visual Basic (ADO, DAO and RDO) and ASP

Some articles that may help with Visual Basic and ASP:

- [MySQL BLOB columns and Visual Basic 6](http://dev.mysql.com/tech-resources/articles/vb-blob-handling.html) [http://dev.mysql.com/tech-resources/articles/vb-blob-handling.html] by Mike Hillyer (<mike@openwin.org>).
- [How to map Visual basic data type to MySQL types](http://dev.mysql.com/tech-resources/articles/visual-basic-datatypes.html) [http://dev.mysql.com/tech-resources/articles/visual-basic-datatypes.html] by Mike Hillyer (<mike@openwin.org>).

#### 18.1.6.2.2. Using MyODBC with Borland Applications

With all Borland applications where the Borland Database Engine (BDE) is used, follow these steps to improve compatibility:

- Update to BDE 3.2 or newer.
- Enable the `Don't optimize column widths` option in the DSN.

- Enabled the `Return matching rows` option in the DSN.

#### 18.1.6.2.2.1. Using MyODBC with Borland Builder 4

When you start a query, you can use the `Active` property or the `Open` method. Note that `Active` starts by automatically issuing a `SELECT * FROM ...` query. That may not be a good thing if your tables are large.

#### 18.1.6.2.2.2. Using MyODBC with Delphi

Also, here is some potentially useful Delphi code that sets up both an ODBC entry and a BDE entry for MyODBC. The BDE entry requires a BDE Alias Editor that is free at a Delphi Super Page near you. (Thanks to Bryan Brunton <[bryan@flesherfab.com](mailto:bryan@flesherfab.com)> for this):

```
fReg:= TRegistry.Create;
fReg.OpenKey('\Software\ODBC\ODBC.INI\DocumentsFab', True);
fReg.WriteString('Database', 'Documents');
fReg.WriteString('Description', '');
fReg.WriteString('Driver', 'C:\WINNT\System32\myodbc.dll');
fReg.WriteString('Flag', '1');
fReg.WriteString('Password', '');
fReg.WriteString('Port', '');
fReg.WriteString('Server', 'xmark');
fReg.WriteString('User', 'winuser');
fReg.OpenKey('\Software\ODBC\ODBC.INI\ODBC Data Sources', True);
fReg.WriteString('DocumentsFab', 'MySQL');
fReg.CloseKey;
fReg.Free;

Memol.Lines.Add('DATABASE NAME=');
Memol.Lines.Add('USER NAME=');
Memol.Lines.Add('ODBC DSN=DocumentsFab');
Memol.Lines.Add('OPEN MODE=READ/WRITE');
Memol.Lines.Add('BATCH COUNT=200');
Memol.Lines.Add('LANGDRIVER=');
Memol.Lines.Add('MAX ROWS=-1');
Memol.Lines.Add('SCHEMA CACHE DIR=');
Memol.Lines.Add('SCHEMA CACHE SIZE=8');
Memol.Lines.Add('SCHEMA CACHE TIME=-1');
Memol.Lines.Add('SQLPASSTHRU MODE=SHARED AUTOCOMMIT');
Memol.Lines.Add('SQLQRYMODE=');
Memol.Lines.Add('ENABLE SCHEMA CACHE=FALSE');
Memol.Lines.Add('ENABLE BCD=FALSE');
Memol.Lines.Add('ROWSET SIZE=20');
Memol.Lines.Add('BLOBS TO CACHE=64');
Memol.Lines.Add('BLOB SIZE=32');

AliasEditor.Add('DocumentsFab', 'MySQL', Memol.Lines);
```

#### 18.1.6.2.2.3. Using MyODBC with C++ Builder

Tested with BDE 3.0. The only known problem is that when the table schema changes, query fields are not updated. BDE, however, does not seem to recognize primary keys, only the index named `PRIMARY`, although this has not been a problem.

#### 18.1.6.2.3. Using MyODBC with ColdFusion

The following information is taken from the ColdFusion documentation:

Use the following information to configure ColdFusion Server for Linux to use the `unixODBC` driver with MyODBC for MySQL data sources. Allaire has verified that MyODBC 2.50.26 works with MySQL 3.22.27 and ColdFusion for Linux. (Any newer version should also work.) You can download MyODBC at <http://dev.mysql.com/downloads/connector/odbc/>.

ColdFusion version 4.5.1 allows you to use the ColdFusion Administrator to add the MySQL data source. However, the driver is not included with ColdFusion version 4.5.1. Before the MySQL driver appears in

the ODBC data sources drop-down list, you must build and copy the MyODBC driver to /  
[opt/coldfusion/lib/libmyodbc.so](#).

The Contrib directory contains the program [mydsn-xxx.zip](#) which allows you to build and remove the DSN registry file for the MyODBC driver on ColdFusion applications.

For more information and guides on using ColdFusion and MyODBC, see the following external sites:

- Refer to [MySQL ColdFusion unixODBC MyODBC and Solaris - how to succeed](#) [<http://dbforums.com/showthread.php?threadid=174934>]
- ColdFusion (on Solaris and NT with service pack 5), [How-to: MySQL and ColdFusion](#) [<http://www.njwtech.net/addons/coldfusion/mysql.html>].
- [Troubleshooting Data Sources and Database Connectivity for Unix Platforms](#) [<http://www.macromedia.com/v1/handlers/index.cfm?ID=11328&Method=Full&PageCall=/support/index.cfm>].

#### 18.1.6.2.4. Using MyODBC with OpenOffice

Open Office (<http://www.openoffice.org>) [How-to: MySQL + OpenOffice](#) [[http://dba.openoffice.org/proposals/MySQL\\_OOo.html](http://dba.openoffice.org/proposals/MySQL_OOo.html)]. [How-to: OpenOffice + MyODBC + unixODBC](#) [<http://www.unixodbc.org/doc/OOoMySQL.pdf>].

#### 18.1.6.2.5. Using MyODBC with Sambar Server

Sambar Server (<http://www.sambarserver.info>) [How-to: MyODBC + SambarServer + MySQL](#) [<http://www.sambarserver.info/article.php?sid=66>].

#### 18.1.6.2.6. Using MyODBC with Pervasive Software DataJunction

You have to change it to output `VARCHAR` rather than `ENUM`, as it exports the latter in a manner that causes MySQL problems.

#### 18.1.6.2.7. Using MyODBC with SunSystems Vision

You should select the `Return matching rows` option.

### 18.1.6.3. MyODBC Errors and Resolutions

The following section details some common errors and their suggested fix or alternative solution. If you are still experiencing problems, use the MyODBC mailing list; see [Section 18.1.7.1, “MyODBC Community Support”](#).

Many problems can be resolved by upgrading your MyODBC drivers to the latest available release. On Windows, you should also make sure that you have the latest versions of the Microsoft Data Access Components (MDAC) installed.

#### Questions

- [19.1.6.3.1](#): Are MyODBC 2.50 applications compatible with MyODBC 3.51?
- [19.1.6.3.2](#): I have installed MyODBC on Windows XP x64 Edition or Windows Server 2003 R2 x64. The installation completed successfully, but the MyODBC driver does not appear in [ODBC Data Source Administrator](#).

- [19.1.6.3.3](#): When connecting or using the **Test** button in **ODBC Data Source Administrator** I get error 10061 (Cannot connect to server)
- [19.1.6.3.4](#): The following error is reported when using transactions: **Transactions are not enabled**
- [19.1.6.3.5](#): The following error is reported when I submit a query: **Cursor not found**
- [19.1.6.3.6](#): Access reports records as **#DELETED#** when inserting or updating records in linked tables.
- [19.1.6.3.7](#): How do I handle Write Conflicts or Row Location errors?
- [19.1.6.3.8](#): Exporting data from Access 97 to MySQL reports a **Syntax Error**.
- [19.1.6.3.9](#): Exporting data from Microsoft DTS to MySQL reports a **Syntax Error**.
- [19.1.6.3.10](#): Using ODBC.NET with MyODBC, while fetching empty string (0 length), it starts giving the SQL\_NO\_DATA exception.
- [19.1.6.3.11](#): Using `SELECT COUNT(*) FROM tbl_name` within Visual Basic and ASP returns an error.
- [19.1.6.3.12](#): Using the `AppendChunk()` or `GetChunk()` ADO methods, the **Multiple-step operation generated errors**. Check each status value error is returned.
- [19.1.6.3.13](#): Access Returns **Another user had modified the record that you have modified** while editing records on a Linked Table.

## Questions and Answers

### 19.1.6.3.1: Are MyODBC 2.50 applications compatible with MyODBC 3.51?

Applications based on MyODBC 2.50 should work fine with MyODBC 3.51 and later versions. If you find something is not working with the latest version of MyODBC which previously worked under an earlier version, please file a bug report. See [Section 18.1.7.2](#), “[How to Report MyODBC Problems or Bugs](#)”.

### 19.1.6.3.2: I have installed MyODBC on Windows XP x64 Edition or Windows Server 2003 R2 x64. The installation completed successfully, but the MyODBC driver does not appear in ODBC Data Source Administrator.

This is not a bug, but is related to the way Windows x64 editions operate with the ODBC driver. On Windows x64 editions, the MyODBC driver is installed in the `%SystemRoot%\SysWOW64` folder. However, the default **ODBC Data Source Administrator** that is available through the **Administrative Tools** or **Control Panel** in Windows x64 Editions is located in the `%SystemRoot%\system32` folder, and only searches this folder for ODBC drivers.

On Windows x64 editions, you should use the ODBC administration tool located at `%SystemRoot%\SysWOW64\odbcad32.exe`, this will correctly locate the installed MyODBC drivers and enable you to create a MyODBC DSN.

This issue was originally reported as [Bug#20301](#) [<http://bugs.mysql.com/20301>].

### 19.1.6.3.3: When connecting or using the **Test** button in **ODBC Data Source Administrator** I get error 10061 (Cannot connect to server)

This error can be raised by a number of different issues, including server problems, network problems,

---

and firewall and port blocking problems. For more information, see [Section A.2.2, “Can't connect to \[local\] MySQL server”](#).

#### **19.1.6.3.4: The following error is reported when using transactions: Transactions are not enabled**

This error indicates that you are trying to use transactions with a MySQL table that does not support transactions. Transactions are supported within MySQL when using the [InnoDB](#) and [BDB](#) database engines.

You should check the following before continuing:

- Verify that your MySQL server supports a transactional database engine. Use `SHOW ENGINES` to obtain a list of the available engine types.
- Verify that the tables you are updating use a transaction database engine.
- Ensure that you have not enabled the `disable transactions` option in your DSN.

#### **19.1.6.3.5: The following error is reported when I submit a query: Cursor not found**

This occurs because the application is using the old MyODBC 2.50 version, and it did not set the cursor name explicitly through `SQLSetCursorName`. The fix is to upgrade to MyODBC 3.51 version.

#### **19.1.6.3.6: Access reports records as #DELETED# when inserting or updating records in linked tables.**

If the inserted or updated records are shown as `#DELETED#` in the access, then:

- If you are using Access 2000, you should get and install the newest (version 2.6 or higher) Microsoft MDAC ([Microsoft Data Access Components](#)) from <http://www.microsoft.com/data/>. This fixes a bug in Access that when you export data to MySQL, the table and column names aren't specified. Another way to work around this bug is to upgrade to MyODBC 2.50.33 or higher and MySQL 3.23.x or higher, which together provide a workaround for the problem.

You should also get and apply the Microsoft Jet 4.0 Service Pack 5 (SP5) which can be found at <http://support.microsoft.com/default.aspx?scid=kb;EN-US;q239114>. This fixes some cases where columns are marked as `#DELETED#` in Access.

Note: If you are using MySQL 3.22, you must apply the MDAC patch and use MyODBC 2.50.32 or 2.50.34 and up to work around this problem.

- For all versions of Access, you should enable the MyODBC `Return matching rows` option. For Access 2.0, you should additionally enable the `Simulate ODBC 1.0` option.
- You should have a timestamp in all tables that you want to be able to update..
- You should have a primary key in the table. If not, new or updated rows may show up as `#DELETED#`.
- Use only `DOUBLE` float fields. Access fails when comparing with single-precision floats. The symptom usually is that new or updated rows may show up as `#DELETED#` or that you can't find or update rows.
- If you are using MyODBC to link to a table that has a `BIGINT` column, the results are displayed as `#DELETED`. The work around solution is:



- Have one more dummy column with `TIMESTAMP` as the data type.
- Select the `Change BIGINT columns to INT` option in the connection dialog in ODBC DSN Administrator.
- Delete the table link from Access and re-create it.

Old records still display as `#DELETED#`, but newly added/updated records are displayed properly.

#### 19.1.6.3.7: How do I handle Write Conflicts or Row Location errors?

If you see the following errors, select the `Return Matching Rows` option in the DSN configuration dialog, or specify `OPTION=2`, as the connection parameter:

```
Write Conflict. Another user has changed your data.  
Row cannot be located for updating. Some values may have been changed  
since it was last read.
```

#### 19.1.6.3.8: Exporting data from Access 97 to MySQL reports a `Syntax Error`.

This error is specific to Access 97 and versions of MyODBC earlier than 3.51.02. Update to the latest version of the MyODBC driver to resolve this problem.

#### 19.1.6.3.9: Exporting data from Microsoft DTS to MySQL reports a `Syntax Error`.

This error occurs only with MySQL tables using the `TEXT` or `VARCHAR` data types. You can fix this error by upgrading your MyODBC driver to version 3.51.02 or higher.

#### 19.1.6.3.10: Using ODBC.NET with MyODBC, while fetching empty string (0 length), it starts giving the `SQL_NO_DATA` exception.

You can get the patch that addresses this problem from <http://support.microsoft.com/default.aspx?scid=kb;EN-US;q319243>.

#### 19.1.6.3.11: Using `SELECT COUNT(*) FROM tbl_name` within Visual Basic and ASP returns an error.

This error occurs because the `COUNT(*)` expression is returning a `BIGINT`, and ADO can't make sense of a number this big. Select the `Change BIGINT columns to INT` option (option value 16384).

#### 19.1.6.3.12: Using the `AppendChunk()` or `GetChunk()` ADO methods, the `Multiple-step operation generated errors. Check each status value error is returned`.

The `GetChunk()` and `AppendChunk()` methods from ADO doesn't work as expected when the cursor location is specified as `adUseServer`. On the other hand, you can overcome this error by using `adUseClient`.

A simple example can be found from [http://www.dwam.net/iishelp/ado/docs/adomth02\\_4.htm](http://www.dwam.net/iishelp/ado/docs/adomth02_4.htm)

#### 19.1.6.3.13: Access Returns `Another user had modified the record that you have modified while editing records on a Linked Table`.

In most cases, this can be solved by doing one of the following things:

- Add a primary key for the table if one doesn't exist.

- Add a timestamp column if one doesn't exist.
- Only use double-precision float fields. Some programs may fail when they compare single-precision floats.

If these strategies don't help, you should start by making a log file from the ODBC manager (the log you get when requesting logs from ODBCADMIN) and a MyODBC log to help you figure out why things go wrong. For instructions, see [Section 18.1.3.8, “Getting an ODBC Trace File”](#).

## 18.1.7. MyODBC Support

There are many different places where you can get support for using MyODBC. You should always try the MyODBC Mailing List or MyODBC Forum. See [Section 18.1.7.1, “MyODBC Community Support”](#), for help before reporting a specific bug or issue to MySQL.

### 18.1.7.1. MyODBC Community Support

MySQL AB provides assistance to the user community by means of its mailing lists. For MyODBC-related issues, you can get help from experienced users by using the `<myodbc@lists.mysql.com>` mailing list. Archives are available online at <http://lists.mysql.com/myodbc>.

For information about subscribing to MySQL mailing lists or to browse list archives, visit <http://lists.mysql.com/>. See [Section 1.7.1, “MySQL Mailing Lists”](#).

Community support from experienced users is also available through the [MyODBC Forum](#) [<http://forums.mysql.com/list.php?37>]. You may also find help from other users in the other MySQL Forums, located at <http://forums.mysql.com>. See [Section 1.7.2, “MySQL Community Support at the MySQL Forums”](#).

### 18.1.7.2. How to Report MyODBC Problems or Bugs

If you encounter difficulties or problems with MyODBC, you should start by making a log file from the [ODBC Manager](#) (the log you get when requesting logs from `ODBC ADMIN`) and MyODBC. The procedure for doing this is described in [Section 18.1.3.8, “Getting an ODBC Trace File”](#).

Check the MyODBC trace file to find out what could be wrong. You should be able to determine what statements were issued by searching for the string `>mysql_real_query` in the `myodbc.log` file.

You should also try issuing the statements from the `mysql` client program or from `admndemo`. This helps you determine whether the error is in MyODBC or MySQL.

If you find out something is wrong, please only send the relevant rows (maximum 40 rows) to the `myo-  
dbc` mailing list. See [Section 1.7.1, “MySQL Mailing Lists”](#). Please never send the whole MyODBC or ODBC log file!

You should ideally include the following information with the email:

- Operating system and version
- MyODBC version
- ODBC Driver Manager type and version
- MySQL server version
- ODBC trace from Driver Manager

- MyODBC log file from MyODBC driver
- Simple reproducible sample

Remember that the more information you can supply to us, the more likely it is that we can fix the problem!

Also, before posting the bug, check the MyODBC mailing list archive at <http://lists.mysql.com/myodbc>.

If you are unable to find out what's wrong, the last option is to create an archive in `tar` or Zip format that contains a MyODBC trace file, the ODBC log file, and a `README` file that explains the problem. You can send this to <ftp://ftp.mysql.com/pub/mysql/upload/>. Only MySQL engineers have access to the files you upload, and we are very discreet with the data.

If you can create a program that also demonstrates the problem, please include it in the archive as well.

If the program works with another SQL server, you should include an ODBC log file where you perform exactly the same SQL statements so that we can compare the results between the two systems.

Remember that the more information you can supply to us, the more likely it is that we can fix the problem.

### 18.1.7.3. How to Submit a MyODBC Patch

You can send a patch or suggest a better solution for any existing code or problems by sending a mail message to [<myodbc@lists.mysql.com>](mailto:myodbc@lists.mysql.com).

### 18.1.7.4. MyODBC Change History

The MyODBC Change History (Changelog) is located with the main Changelog for MySQL. See [Section D.6, “MySQL Connector/ODBC \(MyODBC\) Change History”](#).

### 18.1.7.5. Credits

These are the developers that have worked on the MyODBC and MyODBC 3.51 Drivers from MySQL AB.

- Michael (Monty) Widenius
- Venu Anuganti
- Peter Harvey

## 18.2. Connector/NET

Connector/NET enables developers to easily create .NET applications that require secure, high-performance data connectivity with MySQL. It implements the required ADO.NET interfaces and integrates into ADO.NET aware tools. Developers can build applications using their choice of .NET languages. Connector/NET is a fully managed ADO.NET driver written in 100% pure C#.

Connector/NET includes full support for:

- MySQL 5.0 features (such as stored procedures)

- MySQL 4.1 features (server-side prepared statements, Unicode, and shared memory access, and so forth)
- Large-packet support for sending and receiving rows and BLOBs up to 2 gigabytes in size.
- Protocol compression which allows for compressing the data stream between the client and server.
- Support for connecting using TCP/IP sockets, named pipes, or shared memory on Windows.
- Support for connecting using TCP/IP sockets or Unix sockets on Unix.
- Support for the Open Source Mono framework developed by Novell.
- Fully managed, does not utilize the MySQL client library.

This document is intended as a user's guide to Connector/NET and includes a full syntax reference. Syntax information is also included within the [Documentation.chm](#) file included with the Connector/NET distribution.

## 18.2.1. Connector/NET Versions

There is currently one version of the Connector/NET available:

- Connector/NET 1.0 includes support for MySQL 4.0, and MySQL 5.0 features, and full compatibility with the ADO.NET driver interface.

## 18.2.2. How to install Connector/NET

Connector/NET runs on any platform that supports the .NET framework. The .NET framework is primarily supported on recent versions of Microsoft Windows, and is supported on Linux through the Open Source Mono framework (see <http://www.mono-project.com>).

Connector/NET is available for download from <http://dev.mysql.com/downloads/connector/net/1.0.html>.

### 18.2.2.1. Installing Connector/NET on Windows

On Windows, installation is supported either through a binary installation process or by downloading a Zip file with the Connector/NET components.

Before installing, you should ensure that your system is up to date, including installing the latest version of the .NET Framework.


#### 18.2.2.1.1. Installing Connector/NET using the Installer

Using the installer is the most straightforward method of installing Connector/NET on Windows and the installed components include the source code, test code and full reference documentation.


Connector/NET is installed through the use of a Windows Installer ([.msi](#)) installation package, which can be used to install Connector/NET on all Windows operating systems. The MSI package is contained within a ZIP archive named [mysql-connector-net-version.zip](#), where *version* indicates the Connector/NET version.

To install Connector/NET:

1. Double click on the MSI installer file extracted from the Zip you downloaded. Click [Next](#) to start the installation.




2. You must choose the type of installation that you want to perform.




For most situations, the Typical installation will be suitable. Click the [Typical](#) button and proceed to Step 5. A Complete installation installs all the available files. To conduct a Complete installation, click the [Complete](#) button and proceed to step 5. If you want to customize your installation, including choosing the components to install and some installation options, click the [Custom](#) button and proceed to Step 3.

3. If you have chosen a custom installation, you can select the individual components that you want to install, including the core interface component, supporting documentation (a CHM file) samples and examples and the source code. Select the items, and their installation level, and then click [Next](#) to continue the installation.




4. For a custom installation you can also decide whether the Connector/NET component should be registered in the Global Assembly Cache - this will make the Connector/NET component available to all applications, not just those where you explicitly reference the Connector/NET component. You can also enable, or disable, the creation or appropriate items in the Start menu. Click [Next](#) when you have selected the required options.



5. You will be given a final opportunity to confirm the installation. Click [Install](#) to copy and install the files onto your machine.





6. Once the installation has been completed, click **Finish** to exit the installer.

Unless you choose otherwise, Connector/NET is installed in `C:\Program Files\MySQL\MySQL Connector Net X.X.X`, where `X.X.X` is replaced with the version of Connector/NET you are installing. New installations do not overwrite existing versions of Connector/NET.

Depending on your installation type, the installed components will include some or all of the following components:

- `bin` - Connector/NET MySQL libraries for different versions of the .NET environment.
- `docs` - contains a CHM of the Connector/NET documentation.
- `samples` - sample code and applications that use the Connector/NET component.
- `src` - the source code for the Connector/NET component.

You may also use the `/quiet` or `/q` command line option with the `msiexec` tool to install the Connector/NET package automatically (using the default options) with no notification to the user. Using this option you cannot select options and no prompts, messages or dialog boxes will be displayed.

```
C:\> msiexec /package connector-net.msi /quiet
```

To provide a progress bar to the user during automatic installation, but still without presenting the user with a dialog box of the ability to select options, use the `/passive` option.

### 18.2.2.1.2. Installing Connector/NET using the Zip package

If you are having problems running the installer, you can download a .zip file without an installer as an alternative. That file is called `mysql-connector-net-version-noinstall.zip`. Once downloaded, you can extract the files to a location of your choice.

The .zip file contains the following directories:

- `bin` - Connector/NET MySQL libraries for different versions of the .NET environment.
- `doc` - contains a CHM of the Connector/NET documentation.
- `Samples` - sample code and applications that use the Connector/NET component.
- `mysqlclient` - the source code for the Connector/NET component.
- `testsuite` - the test suite used to verify the operation of the Connector/NET component.

### 18.2.2.2. Installing Connector/NET on Unix with Mono

There is no installer available for installing the Connector/NET component on your Unix installation. However, the installation is very simple. Before installing, please ensure that you have a working Mono project installation.

Note that you should only install the Connector/NET component on Unix environments where you want to connect to a MySQL server through the Mono project. If you are deploying or developing on a different environment such as Java or Perl then you should use a more appropriate connectivity component. See the [Chapter 18, Connectors](#), or [Chapter 17, APIs and Libraries](#), for more information.

To install Connector/NET on Unix/Mono:

1. Download the `mysql-connector-net-version-noinstall.zip` and extract the contents.
2. Copy the `MySQL.Data.dll` file to your Mono project installation folder.
3. You must register the Connector/NET component in the Global Assembly Cache using the `gacutil` command:

```
shell> gacutil /i MySQL.Data.dll
```

Once installed, applications that are compiled with the Connector/NET component need no further changes. However, you must ensure that when you compile your applications you include the Connector/NET component using the `-r:MySQLData.dll` command line option.

### 18.2.2.3. Installing Connector/NET using the Source

**Caution:** You should read this section only if you are interested in helping us test our new code. If you just want to get Connector/NET up and running on your system, you should use a standard release distribution.

To be able to access the Connector/NET source tree, you must have Subversion installed. Subversion is freely available from <http://subversion.tigris.org/>.

The most recent development source tree is available from our public Subversion trees at <http://dev.mysql.com/tech-resources/sources.html>.

To checkout out the Connector/NET sources, change to the directory where you want the copy of the Connector/NET tree to be stored, then use the following command:

```
shell> svn co
http://svn.mysql.com/svnpublic/connector-net
```

A Visual Studio project is included in the source which you can use to build Connector/NET.

## 18.2.3. Connector/NET Examples

Connector/NET comprises several classes that are used to connect to the database, execute queries and statements, and manage query results.

The following are the major classes of Connector/NET:

- [MySqlCommand](#): Represents an SQL statement to execute against a MySQL database.
- [MySqlCommandBuilder](#): Automatically generates single-table commands used to reconcile changes made to a DataSet with the associated MySQL database.
- [MySqlConnection](#): Represents an open connection to a MySQL Server database.
- [MySqlDataAdapter](#): Represents a set of data commands and a database connection that are used to fill a dataset and update a MySQL database.
- [MySqlDataReader](#): Provides a means of reading a forward-only stream of rows from a MySQL database.
- [MySqlException](#): The exception that is thrown when MySQL returns an error.
- [MySqlHelper](#): Helper class that makes it easier to work with the provider.
- [MySqlTransaction](#): Represents an SQL transaction to be made in a MySQL database.

This section contains basic information and examples for each of the above classes. For a more detailed reference guide please see [Section 18.2.4, “Connector/NET Reference”](#).

### 18.2.3.1. MySqlCommand

Represents a SQL statement to execute against a MySQL database. This class cannot be inherited.

[MySqlCommand](#) features the following methods for executing commands at a MySQL database:

Item	Description
<a href="#">ExecuteReader</a>	Executes commands that return rows.
<a href="#">ExecuteNonQuery</a>	Executes commands such as SQL INSERT, DELETE, and UPDATE statements.
<a href="#">ExecuteScalar</a>	Retrieves a single value (for example, an aggregate value) from a database.

You can reset the [CommandText](#) property and reuse the [MySqlCommand](#) object. However, you must

close the [MySqlDataReader](#) before you can execute a new or previous command.

If a [MySqlException](#) is generated by the method executing a [MySqlCommand](#), the [MySqlConnection](#) remains open. It is the responsibility of the programmer to close the connection.

**Note.** Prior versions of the provider used the '@' symbol to mark parameters in SQL. This is incompatible with MySQL user variables, so the provider now uses the '?' symbol to locate parameters in SQL. To support older code, you can set 'old syntax=yes' on your connection string. If you do this, please be aware that an exception will not be throw if you fail to define a parameter that you intended to use in your SQL.

### Examples

The following example creates a [MySqlCommand](#) and a [MySqlConnection](#). The [MySqlConnection](#) is opened and set as the [Connection](#) for the [MySqlCommand](#). The example then calls [ExecuteNonQuery](#), and closes the connection. To accomplish this, the [ExecuteNonQuery](#) is passed a connection string and a query string that is a SQL INSERT statement.

Visual Basic example:

```
Public Sub InsertRow(myConnectionString As String)
    " If the connection string is null, use a default.
    If myConnectionString = "" Then
        myConnectionString = "Database=Test;Data Source=localhost;User Id=username;Password=pass"
    End If
    Dim myConnection As New MySqlConnection(myConnectionString)
    Dim myInsertQuery As String = "INSERT INTO Orders (id, customerId, amount) Values(1001, 23, 30.66)"
    Dim myCommand As New MySqlCommand(myInsertQuery)
    myCommand.Connection = myConnection
    myConnection.Open()
    myCommand.ExecuteNonQuery()
    myCommand.Connection.Close()
End Sub
```

C# example:

```
public void InsertRow(string myConnectionString)
{
    // If the connection string is null, use a default.
    if(myConnectionString == "")
    {
        myConnectionString = "Database=Test;Data Source=localhost;User Id=username;Password=pass";
    }
    MySqlConnection myConnection = new MySqlConnection(myConnectionString);
    string myInsertQuery = "INSERT INTO Orders (id, customerId, amount) Values(1001, 23, 30.66)";
    MySqlCommand myCommand = new MySqlCommand(myInsertQuery);
    myCommand.Connection = myConnection;
    myConnection.Open();
    myCommand.ExecuteNonQuery();
    myCommand.Connection.Close();
}
```

### 18.2.3.1.1. Class MySqlCommand Constructor Form 1

#### Overload methods for MySqlCommand

Initializes a new instance of the MySqlCommand class.

#### Examples

The following example creates a MySqlCommand and sets some of its properties.

**Note.** This example shows how to use one of the overloaded versions of the MySqlCommand constructor. For other examples that might be available, see the individual overload topics.

Visual Basic example:

```
Public Sub CreateMySQLCommand()
    Dim myConnection As New MySqlConnection _
        ("Persist Security Info=False;database=test;server=myServer")
    myConnection.Open()
    Dim myTrans As MySQLTransaction = myConnection.BeginTransaction()
    Dim mySelectQuery As String = "SELECT * FROM MyTable"
    Dim myCommand As New MySqlCommand(mySelectQuery, myConnection, myTrans)
    myCommand.CommandTimeout = 20
End Sub
```

C# example:

```
public void CreateMySQLCommand()
{
    MySqlConnection myConnection = new MySqlConnection("Persist Security Info=False;
        database=test;server=myServer");
    myConnection.Open();
    MySQLTransaction myTrans = myConnection.BeginTransaction();
    string mySelectQuery = "SELECT * FROM myTable";
    MySqlCommand myCommand = new MySqlCommand(mySelectQuery, myConnection, myTrans);
    myCommand.CommandTimeout = 20;
}
```

C++ example:

```
public:
void CreateMySQLCommand()
{
    MySqlConnection* myConnection = new MySqlConnection(S"Persist Security Info=False;
        database=test;server=myServer");
    myConnection->Open();
    MySQLTransaction* myTrans = myConnection->BeginTransaction();
    String* mySelectQuery = S"SELECT * FROM myTable";
    MySqlCommand* myCommand = new MySqlCommand(mySelectQuery, myConnection, myTrans);
    myCommand->CommandTimeout = 20;
};
```

Initializes a new instance of the `MySQLCommand` class.

The base constructor initializes all fields to their default values. The following table shows initial property values for an instance of `MySQLCommand`.

Properties	Initial Value
<code>CommandText</code>	empty string ("")
<code>CommandTimeout</code>	0
<code>CommandType</code>	<code>CommandType.Text</code>
<code>Connection</code>	Null

You can change the value for any of these properties through a separate call to the property.

### Examples

The following example creates a `MySQLCommand` and sets some of its properties.

Visual Basic example:

```
Public Sub CreateMySQLCommand()
    Dim myCommand As New MySqlCommand()
    myCommand.CommandType = CommandType.Text
```

---

End Sub

C# example:

```
public void CreateMySQLCommand()
{
    MySqlCommand myCommand = new MySqlCommand();
    myCommand.CommandType = CommandType.Text;
}
```

### 18.2.3.1.2. Class MySqlCommand Constructor Form 2

Initializes a new instance of the [MySqlCommand](#) class with the text of the query.

**Parameters:** The text of the query.

When an instance of [MySqlCommand](#) is created, the following read/write properties are set to initial values.

Properties	Initial Value
<a href="#">CommandText</a>	<a href="#">cmdText</a>
<a href="#">CommandTimeout</a>	0
<a href="#">CommandType</a>	<a href="#">CommandType.Text</a>
<a href="#">Connection</a>	Null

You can change the value for any of these properties through a separate call to the property.

#### Examples

The following example creates a [MySqlCommand](#) and sets some of its properties.

Visual Basic example:

```
Public Sub CreateMySQLCommand()
    Dim sql as String = "SELECT * FROM mytable"
    Dim myCommand As New MySqlCommand(sql)
    myCommand.CommandType = CommandType.Text
End Sub
```

C# example:

```
public void CreateMySQLCommand()
{
    string sql = "SELECT * FROM mytable";
    MySqlCommand myCommand = new MySqlCommand(sql);
    myCommand.CommandType = CommandType.Text;
}
```

### 18.2.3.1.3. Class MySqlCommand Constructor Form 3

Initializes a new instance of the [MySqlCommand](#) class with the text of the query and a [MySQLConnection](#).

**Parameters:** The text of the query.

**Parameters:** A [MySQLConnection](#) that represents the connection to an instance of SQL Server.

When an instance of [MySqlCommand](#) is created, the following read/write properties are set to initial

---

values.

Properties	Initial Value
<a href="#">CommandText</a>	<a href="#">cmdText</a>
<a href="#">CommandTimeout</a>	<b>0</b>
<a href="#">CommandType</a>	<a href="#">CommandType.Text</a>
<a href="#">Connection</a>	<a href="#">connection</a>

You can change the value for any of these properties through a separate call to the property.

### Examples

The following example creates a [MySQLCommand](#) and sets some of its properties.

Visual Basic example:

```
Public Sub CreateMySQLCommand()
    Dim conn as new MySqlConnection("server=myServer")
    Dim sql as String = "SELECT * FROM mytable"
    Dim myCommand As New MySqlCommand(sql, conn)
    myCommand.CommandType = CommandType.Text
End Sub
```

C# example:

```
public void CreateMySQLCommand()
{
    MySqlConnection conn = new MySqlConnection("server=myserver")
    string sql = "SELECT * FROM mytable";
    MySqlCommand myCommand = new MySqlCommand(sql, conn);
    myCommand.CommandType = CommandType.Text;
}
```

#### 18.2.3.1.4. Class MySqlCommand Constructor Form 4

Initializes a new instance of the [MySQLCommand](#) class with the text of the query, a [MySQLConnection](#), and the [MySQLTransaction](#).

**Parameters:** The text of the query.

**Parameters:** A [MySQLConnection](#) that represents the connection to an instance of SQL Server.

**Parameters:** The [MySQLTransaction](#) in which the [MySQLCommand](#) executes.

When an instance of [MySQLCommand](#) is created, the following read/write properties are set to initial values.

Properties	Initial Value
<a href="#">CommandText</a>	<a href="#">cmdText</a>
<a href="#">CommandTimeout</a>	<b>0</b>
<a href="#">CommandType</a>	<a href="#">CommandType.Text</a>
<a href="#">Connection</a>	<a href="#">connection</a>

You can change the value for any of these properties through a separate call to the property.

### Examples

The following example creates a `MySQLCommand` and sets some of its properties.

Visual Basic example:

```
Public Sub CreateMySQLCommand()
    Dim conn as new MySqlConnection("server=myServer")
    conn.Open();
    Dim txn as MySqlTransaction = conn.BeginTransaction()
    Dim sql as String = "SELECT * FROM mytable"
    Dim myCommand As New MySQLCommand(sql, conn, txn)
    myCommand.CommandType = CommandType.Text
End Sub
```

C# example:

```
public void CreateMySQLCommand()
{
    MySqlConnection conn = new MySqlConnection("server=myserver")
    conn.Open();
    MySqlTransaction txn = conn.BeginTransaction();
    string sql = "SELECT * FROM mytable";
    MySQLCommand myCommand = new MySQLCommand(sql, conn, txn);
    myCommand.CommandType = CommandType.Text;
}
```

### 18.2.3.1.5. ExecuteNonQuery

Executes a SQL statement against the connection and returns the number of rows affected.

**Returns:** Number of rows affected

You can use `ExecuteNonQuery` to perform any type of database operation, however any resultsets returned will not be available. Any output parameters used in calling a stored procedure will be populated with data and can be retrieved after execution is complete. For `UPDATE`, `INSERT`, and `DELETE` statements, the return value is the number of rows affected by the command. For all other types of statements, the return value is -1.

#### Examples

The following example creates a `MySQLCommand` and then executes it using `ExecuteNonQuery`. The example is passed a string that is a SQL statement (such as `UPDATE`, `INSERT`, or `DELETE`) and a string to use to connect to the data source.

Visual Basic example:

```
Public Sub CreateMySQLCommand(myExecuteQuery As String, myConnection As MySqlConnection)
    Dim myCommand As New MySQLCommand(myExecuteQuery, myConnection)
    myCommand.Connection.Open()
    myCommand.ExecuteNonQuery()
    myConnection.Close()
End Sub
```

C# example:

```
public void CreateMySQLCommand(string myExecuteQuery, MySqlConnection myConnection)
{
    MySQLCommand myCommand = new MySQLCommand(myExecuteQuery, myConnection);
    myCommand.Connection.Open();
    myCommand.ExecuteNonQuery();
    myConnection.Close();
}
```



### 18.2.3.1.6. ExecuteReader1

Sends the `CommandText` to the `MySqlConnection` `Connection`, and builds a `MySqlDataReader` using one of the `CommandBehavior` values.

**Parameters:** One of the `CommandBehavior` values.

When the `CommandType` property is set to `StoredProcedure`, the `CommandText` property should be set to the name of the stored procedure. The command executes this stored procedure when you call `ExecuteReader`.

The `MySqlDataReader` supports a special mode that enables large binary values to be read efficiently. For more information, see the `SequentialAccess` setting for `CommandBehavior`.

While the `MySqlDataReader` is in use, the associated `MySqlConnection` is busy serving the `MySqlDataReader`. While in this state, no other operations can be performed on the `MySqlConnection` other than closing it. This is the case until the `MySqlDataReader.Close` method of the `MySqlDataReader` is called. If the `MySqlDataReader` is created with `CommandBehavior` set to `CloseConnection`, closing the `MySqlDataReader` closes the connection automatically.

**Note.** When calling `ExecuteReader` with the `SingleRow` behavior, you should be aware that using a `limit` clause in your SQL will cause all rows (up to the limit given) to be retrieved by the client. The `MySqlDataReader.Read` method will still return `false` after the first row but pulling all rows of data into the client will have a performance impact. If the `limit` clause is not necessary, it should be avoided.

**Returns:** A `MySqlDataReader` object.

### 18.2.3.1.7. ExecuteReader

Sends the `CommandText` to the `MySqlConnection` `Connection` and builds a `MySqlDataReader`.

**Returns:** A `MySqlDataReader` object.

When the `CommandType` property is set to `StoredProcedure`, the `CommandText` property should be set to the name of the stored procedure. The command executes this stored procedure when you call `ExecuteReader`.

While the `MySqlDataReader` is in use, the associated `MySqlConnection` is busy serving the `MySqlDataReader`. While in this state, no other operations can be performed on the `MySqlConnection` other than closing it. This is the case until the `MySqlDataReader.Close` method of the `MySqlDataReader` is called.

#### Examples

The following example creates a `MySqlCommand`, then executes it by passing a string that is a SQL `SELECT` statement, and a string to use to connect to the data source.

Visual Basic example:

```
Public Sub CreateMySqlDataReader(mySelectQuery As String, myConnection As MySqlConnection)
    Dim myCommand As New MySqlCommand(mySelectQuery, myConnection)
    myConnection.Open()
    Dim myReader As MySqlDataReader
    myReader = myCommand.ExecuteReader()
    Try
        While myReader.Read()
            Console.WriteLine(myReader.GetString(0))
        End While
    Finally
        myReader.Close()
        myConnection.Close()
    End Try
End Sub
```

```
End Try
End Sub
```

C# example:

```
public void CreateMySqlDataReader(string mySelectQuery, MySqlConnection myConnection)
{
    MySqlCommand myCommand = new MySqlCommand(mySelectQuery, myConnection);
    myConnection.Open();
    MySqlDataReader myReader;
    myReader = myCommand.ExecuteReader();
    try
    {
        while(myReader.Read())
        {
            Console.WriteLine(myReader.GetString(0));
        }
    }
    finally
    {
        myReader.Close();
        myConnection.Close();
    }
}
```

### 18.2.3.1.8. Prepare

Creates a prepared version of the command on an instance of MySQL Server.

Prepared statements are only supported on MySQL version 4.1 and higher. Calling prepare while connected to earlier versions of MySQL will succeed but will execute the statement in the same way as unprepared.

#### Examples

The following example demonstrates the use of the [Prepare](#) method.

Visual Basic example:

```
public sub PrepareExample()
    Dim cmd as New MySqlCommand("INSERT INTO mytable VALUES (?val)", myConnection)
    cmd.Parameters.Add( "?val", 10 )
    cmd.Prepare()
    cmd.ExecuteNonQuery()

    cmd.Parameters(0).Value = 20
    cmd.ExecuteNonQuery()
end sub
```

C# example:

```
private void PrepareExample()
{
    MySqlCommand cmd = new MySqlCommand("INSERT INTO mytable VALUES (?val)", myConnection);
    cmd.Parameters.Add( "?val", 10 );
    cmd.Prepare();
    cmd.ExecuteNonQuery();

    cmd.Parameters[0].Value = 20;
    cmd.ExecuteNonQuery();
}
```

### 18.2.3.1.9. ExecuteScalar

Executes the query, and returns the first column of the first row in the result set returned by the query. Extra columns or rows are ignored.

**Returns:** The first column of the first row in the result set, or a null reference if the result set is empty

Use the [ExecuteScalar](#) method to retrieve a single value (for example, an aggregate value) from a database. This requires less code than using the [ExecuteReader](#) method, and then performing the operations necessary to generate the single value using the data returned by a [MySqlDataReader](#)

A typical [ExecuteScalar](#) query can be formatted as in the following C# example:

C# example:

```
cmd.CommandText = "select count(*) from region";
Int32 count = (int32) cmd.ExecuteScalar();
```

### Examples

The following example creates a [MySqlCommand](#) and then executes it using [ExecuteScalar](#). The example is passed a string that is a SQL statement that returns an aggregate result, and a string to use to connect to the data source.

Visual Basic example:

```
Public Sub CreateMySqlCommand(myScalarQuery As String, myConnection As MySqlConnection)
    Dim myCommand As New MySqlCommand(myScalarQuery, myConnection)
    myCommand.Connection.Open()
    myCommand.ExecuteScalar()
    myConnection.Close()
End Sub
```

C# example:

```
public void CreateMySqlCommand(string myScalarQuery, MySqlConnection myConnection)
{
    MySqlCommand myCommand = new MySqlCommand(myScalarQuery, myConnection);
    myCommand.Connection.Open();
    myCommand.ExecuteScalar();
    myConnection.Close();
}
```

C++ example:

```
public:
    void CreateMySqlCommand(String* myScalarQuery, MySqlConnection* myConnection)
    {
        MySqlCommand* myCommand = new MySqlCommand(myScalarQuery, myConnection);
        myCommand->Connection->Open();
        myCommand->ExecuteScalar();
        myConnection->Close();
    }
```

#### 18.2.3.1.10. CommandText

Gets or sets the SQL statement to execute at the data source.

**Value:** The SQL statement or stored procedure to execute. The default is an empty string.

When the [CommandType](#) property is set to [StoredProcedure](#), the [CommandText](#) property should be set to the name of the stored procedure. The user may be required to use escape character syntax if the stored procedure name contains any special characters. The command executes this stored pro-

cedure when you call one of the Execute methods.

### Examples

The following example creates a [MySQLCommand](#) and sets some of its properties.

Visual Basic example:

```
Public Sub CreateMySQLCommand()  
    Dim myCommand As New MySQLCommand()  
    myCommand.CommandText = "SELECT * FROM Mytable ORDER BY id"  
    myCommand.CommandType = CommandType.Text  
End Sub
```

C# example:

```
public void CreateMySQLCommand()  
{  
    MySQLCommand myCommand = new MySQLCommand();  
    myCommand.CommandText = "SELECT * FROM mytable ORDER BY id";  
    myCommand.CommandType = CommandType.Text;  
}
```

#### 18.2.3.1.11. CommandTimeout

Gets or sets the wait time before terminating the attempt to execute a command and generating an error.

**Value:** The time (in seconds) to wait for the command to execute. The default is 0 seconds.

MySQL currently does not support any method of canceling a pending or executing operation. All commands issued against a MySQL server will execute until completion or exception occurs.

#### 18.2.3.1.12. CommandType

Gets or sets a value indicating how the [CommandText](#) property is to be interpreted.

**Value:** One of the [System.Data.CommandType](#) values. The default is [Text](#).

When you set the [CommandType](#) property to [StoredProcedure](#), you should set the [CommandText](#) property to the name of the stored procedure. The command executes this stored procedure when you call one of the Execute methods.

### Examples

The following example creates a [MySQLCommand](#) and sets some of its properties.

Visual Basic example:

```
Public Sub CreateMySQLCommand()  
    Dim myCommand As New MySQLCommand()  
    myCommand.CommandType = CommandType.Text  
End Sub
```

C# example:

```
public void CreateMySQLCommand()  
{  
    MySQLCommand myCommand = new MySQLCommand();  
    myCommand.CommandType = CommandType.Text;  
}
```

#### 18.2.3.1.13. Connection

---

Gets or sets the [MySQLConnection](#) used by this instance of the [MySQLCommand](#).

**Value:** The connection to a data source. The default value is a null reference ([Nothing](#) in Visual Basic).

If you set [Connection](#) while a transaction is in progress and the [Transaction](#) property is not null, an [InvalidOperationException](#) is generated. If the [Transaction](#) property is not null and the transaction has already been committed or rolled back, [Transaction](#) is set to null.

### Examples

The following example creates a [MySQLCommand](#) and sets some of its properties.

Visual Basic example:

```
Public Sub CreateMySQLCommand()
    Dim mySelectQuery As String = "SELECT * FROM mytable ORDER BY id"
    Dim myConnectionString As String = "Persist Security Info=False;database=test;server=myServer"
    Dim myCommand As New MySQLCommand(mySelectQuery)
    myCommand.Connection = New MySQLConnection(myConnectionString)
    myCommand.CommandType = CommandType.Text
End Sub
```

C# example:

```
public void CreateMySQLCommand()
{
    string mySelectQuery = "SELECT * FROM mytable ORDER BY id";
    string myConnectionString = "Persist Security Info=False;database=test;server=myServer";
    MySQLCommand myCommand = new MySQLCommand(mySelectQuery);
    myCommand.Connection = new MySQLConnection(myConnectionString);
    myCommand.CommandType = CommandType.Text;
}
```

#### 18.2.3.1.14. IsPrepared

Returns true if the statement is prepared.

#### 18.2.3.1.15. Parameters

Get the [MySQLParameterCollection](#)

**Value:** The parameters of the SQL statement or stored procedure. The default is an empty collection.

Connector/Net does not support unnamed parameters. Every parameter added to the collection must have an associated name.

### Examples

The following example creates a [MySQLCommand](#) and displays its parameters. To accomplish this, the method is passed a [MySQLConnection](#), a query string that is a SQL [SELECT](#) statement, and an array of [MySQLParameter](#) objects.

Visual Basic example:

```
Public Sub CreateMySQLCommand(myConnection As MySQLConnection, _
mySelectQuery As String, myParamArray() As MySQLParameter)
    Dim myCommand As New MySQLCommand(mySelectQuery, myConnection)
    myCommand.CommandText = "SELECT id, name FROM mytable WHERE age=?age"
    myCommand.UpdatedRowSource = UpdateRowSource.Both
    myCommand.Parameters.Add(myParamArray)
    Dim j As Integer
    For j = 0 To myCommand.Parameters.Count - 1
```

```

        myCommand.Parameters.Add(myParamArray(j))
    Next j
    Dim myMessage As String = ""
    Dim i As Integer
    For i = 0 To myCommand.Parameters.Count - 1
        myMessage += myCommand.Parameters(i).ToString() & ControlChars.Cr
    Next i
    Console.WriteLine(myMessage)
End Sub

```

C# example:

```

public void CreateMySQLCommand(MySqlConnection myConnection, string mySelectQuery,
    MySqlParameter[] myParamArray)
{
    MySqlCommand myCommand = new MySqlCommand(mySelectQuery, myConnection);
    myCommand.CommandText = "SELECT id, name FROM mytable WHERE age=?age";
    myCommand.Parameters.Add(myParamArray);
    for (int j=0; j<myParamArray.Length; j++)
    {
        myCommand.Parameters.Add(myParamArray[j]) ;
    }
    string myMessage = "";
    for (int i = 0; i < myCommand.Parameters.Count; i++)
    {
        myMessage += myCommand.Parameters[i].ToString() + "\n";
    }
    MessageBox.Show(myMessage);
}

```

### 18.2.3.1.16. Transaction

Gets or sets the [MySQLTransaction](#) within which the [MySQLCommand](#) executes.

**Value:** The [MySQLTransaction](#). The default value is a null reference ([Nothing](#) in Visual Basic).

You cannot set the [Transaction](#) property if it is already set to a specific value, and the command is in the process of executing. If you set the transaction property to a [MySQLTransaction](#) object that is not connected to the same [MySQLConnection](#) as the [MySQLCommand](#) object, an exception will be thrown the next time you attempt to execute a statement.

### 18.2.3.1.17. UpdatedRowSource

Gets or sets how command results are applied to the [DataRow](#) when used by the [System.Data.Common.DbDataAdapter.Update](#) method of the [System.Data.Common.DbDataAdapter](#).

**Value:** One of the [UpdateRowSource](#) values.

The default [System.Data.UpdateRowSource](#) value is [Both](#) unless the command is automatically generated (as in the case of the [MySQLCommandBuilder](#)), in which case the default is [None](#).

## 18.2.3.2. MySqlCommandBuilder

Automatically generates single-table commands used to reconcile changes made to a [DataSet](#) with the associated [MySQL](#) database. This class cannot be inherited.

The [MySQLDataAdapter](#) does not automatically generate the SQL statements required to reconcile changes made to a [System.Data.DataSet](#) with the associated instance of [MySQL](#). However, you can create a [MySQLCommandBuilder](#) object to automatically generate SQL statements for single-table updates if you set the [MySQLDataAdapter.SelectCommand](#) property of the [MySQLDataAdapter](#). Then, any additional SQL statements that you do not set are generated by the [MySQLCommandBuilder](#).

The `MySqlCommandBuilder` registers itself as a listener for `MySqlDataAdapter.OnRowUpdating` events whenever you set the `DataAdapter` property. You can only associate one `MySqlDataAdapter` or `MySqlCommandBuilder` object with each other at one time.

To generate INSERT, UPDATE, or DELETE statements, the `MySqlCommandBuilder` uses the `SelectCommand` property to retrieve a required set of metadata automatically. If you change the `SelectCommand` after the metadata has been retrieved (for example, after the first update), you should call the `RefreshSchema` method to update the metadata.

The `SelectCommand` must also return at least one primary key or unique column. If none are present, an `InvalidOperationException` exception is generated, and the commands are not generated.

The `MySqlCommandBuilder` also uses the `MySqlCommand.Connection`, `MySqlCommand.CommandTimeout`, and `MySqlCommand.Transaction` properties referenced by the `SelectCommand`. The user should call `RefreshSchema` if any of these properties are modified, or if the `SelectCommand` itself is replaced. Otherwise the `MySqlDataAdapter.InsertCommand`, `MySqlDataAdapter.UpdateCommand`, and `MySqlDataAdapter.DeleteCommand` properties retain their previous values.

If you call `Dispose`, the `MySqlCommandBuilder` is disassociated from the `MySqlDataAdapter`, and the generated commands are no longer used.

**Note.** Caution must be used when using `MySqlCommandBuilder` on MySQL 4.0 systems. With MySQL 4.0, database/schema information is not provided to the connector for a query. This means that a query that pulls columns from two identically named tables in two or more different databases will not cause an exception to be thrown but will not work correctly. Even more dangerous is the situation where your select statement references database X but is executed in database Y and both databases have tables with similar layouts. This situation can cause unwanted changes or deletes. This note does not apply to MySQL versions 4.1 and later.

## Examples

The following example uses the `MySqlCommand`, along `MySqlDataAdapter` and `MySqlConnection`, to select rows from a data source. The example is passed an initialized `System.Data.DataSet`, a connection string, a query string that is a SQL `SELECT` statement, and a string that is the name of the database table. The example then creates a `MySqlCommandBuilder`.

Visual Basic example:

```
Public Shared Function SelectRows(myConnection As String, mySelectQuery As String, myTableName As String)
    Dim myConn As New MySqlConnection(myConnection)
    Dim myDataAdapter As New MySqlDataAdapter()
    myDataAdapter.SelectCommand = New MySqlCommand(mySelectQuery, myConn)
    Dim cb As SqlCommandBuilder = New MySqlCommandBuilder(myDataAdapter)
    myConn.Open()
    Dim ds As DataSet = New DataSet
    myDataAdapter.Fill(ds, myTableName)
    ' Code to modify data in DataSet here
    ' Without the MySqlCommandBuilder this line would fail.
    myDataAdapter.Update(ds, myTableName)
    myConn.Close()
End Function 'SelectRows
```

C# example:

```
public static DataSet SelectRows(string myConnection, string mySelectQuery, string myTableName)
{
    MySqlConnection myConn = new MySqlConnection(myConnection);
    MySqlDataAdapter myDataAdapter = new MySqlDataAdapter();
    myDataAdapter.SelectCommand = new MySqlCommand(mySelectQuery, myConn);
```

```
MySQLCommandBuilder cb = new MySQLCommandBuilder(myDataAdapter);
myConn.Open();
DataSet ds = new DataSet();
myDataAdapter.Fill(ds, myTableName);
//code to modify data in DataSet here
//Without the MySQLCommandBuilder this line would fail
myDataAdapter.Update(ds, myTableName);
myConn.Close();
return ds;
}
```

### 18.2.3.2.1. Class MySQLCommandBuilder Constructor

Initializes a new instance of the [MySQLCommandBuilder](#) class.

### 18.2.3.2.2. Class MySQLCommandBuilder Constructor Form 1

Initializes a new instance of the [MySQLCommandBuilder](#) class and sets the last one wins property.

**Parameters:** False to generate change protection code. True otherwise.

The [lastOneWins](#) parameter indicates whether SQL code should be included with the generated DELETE and UPDATE commands that checks the underlying data for changes. If [lastOneWins](#) is true then this code is not included and data records could be overwritten in a multi-user or multi-threaded environments. Setting [lastOneWins](#) to false will include this check which will cause a concurrency exception to be thrown if the underlying data record has changed without our knowledge.

### 18.2.3.2.3. Class MySQLCommandBuilder Constructor Form 2

Initializes a new instance of the [MySQLCommandBuilder](#) class with the associated [MySQLDataAdapter](#) object.

**Parameters:** The [MySQLDataAdapter](#) to use.

The [MySQLCommandBuilder](#) registers itself as a listener for [MySQLDataAdapter.RowUpdating](#) events that are generated by the [MySQLDataAdapter](#) specified in this property.

When you create a new instance [MySQLCommandBuilder](#), any existing [MySQLCommandBuilder](#) associated with this [MySQLDataAdapter](#) is released.

### 18.2.3.2.4. Class MySQLCommandBuilder Constructor Form 3

Initializes a new instance of the [MySQLCommandBuilder](#) class with the associated [MySQLDataAdapter](#) object.

**Parameters:** The [MySQLDataAdapter](#) to use.

**Parameters:** False to generate change protection code. True otherwise.

The [MySQLCommandBuilder](#) registers itself as a listener for [MySQLDataAdapter.RowUpdating](#) events that are generated by the [MySQLDataAdapter](#) specified in this property.

When you create a new instance [MySQLCommandBuilder](#), any existing [MySQLCommandBuilder](#) associated with this [MySQLDataAdapter](#) is released.

The [lastOneWins](#) parameter indicates whether SQL code should be included with the generated DELETE and UPDATE commands that checks the underlying data for changes. If [lastOneWins](#) is true then this code is not included and data records could be overwritten in a multi-user or multi-threaded en-



vironments. Setting `lastOneWins` to false will include this check which will cause a concurrency exception to be thrown if the underlying data record has changed without our knowledge.

### 18.2.3.2.5. DataAdapter

Gets or sets a `MySqlDataAdapter` object for which SQL statements are automatically generated.

**Value:** A `MySqlDataAdapter` object.

The `MySqlCommandBuilder` registers itself as a listener for `MySqlDataAdapter.RowUpdating` events that are generated by the `MySqlDataAdapter` specified in this property.

When you create a new instance `MySqlCommandBuilder`, any existing `MySqlCommandBuilder` associated with this `MySqlDataAdapter` is released.

### 18.2.3.2.6. QuotePrefix

Gets or sets the beginning character or characters to use when specifying MySQL database objects (for example, tables or columns) whose names contain characters such as spaces or reserved tokens.

**Value:** The beginning character or characters to use. The default value is ` `.

Database objects in MySQL can contain special characters such as spaces that would make normal SQL strings impossible to correctly parse. Use of the `QuotePrefix` and the `QuoteSuffix` properties allows the `MySqlCommandBuilder` to build SQL commands that handle this situation.

### 18.2.3.2.7. QuoteSuffix

Gets or sets the beginning character or characters to use when specifying MySQL database objects (for example, tables or columns) whose names contain characters such as spaces or reserved tokens.

**Value:** The beginning character or characters to use. The default value is ` `.

Database objects in MySQL can contain special characters such as spaces that would make normal SQL strings impossible to correctly parse. Use of the `QuotePrefix` and the `QuoteSuffix` properties allows the `MySqlCommandBuilder` to build SQL commands that handle this situation.

### 18.2.3.2.8. DeriveParameters

### 18.2.3.2.9. GetDeleteCommand

Gets the automatically generated `MySqlCommand` object required to perform deletions on the database.

**Returns:** The `MySqlCommand` object generated to handle delete operations.

An application can use the `GetDeleteCommand` method for informational or troubleshooting purposes because it returns the `MySqlCommand` object to be executed.

You can also use `GetDeleteCommand` as the basis of a modified command. For example, you might call `GetDeleteCommand` and modify the `MySqlCommand.CommandTimeout` value, and then explicitly set that on the `MySqlDataAdapter`.

After the SQL statement is first generated, the application must explicitly call `RefreshSchema` if it changes the statement in any way. Otherwise, the `GetDeleteCommand` will be still be using information from the previous statement, which might not be correct. The SQL statements are first generated either when the application calls `System.Data.Common.DataAdapter.Update` or `GetDe-`

[leteCommand](#).

### 18.2.3.2.10. GetInsertCommand

Gets the automatically generated [MySQLCommand](#) object required to perform insertions on the database.

**Returns:** The [MySQLCommand](#) object generated to handle insert operations.

An application can use the [GetInsertCommand](#) method for informational or troubleshooting purposes because it returns the [MySQLCommand](#) object to be executed.

You can also use the [GetInsertCommand](#) as the basis of a modified command. For example, you might call [GetInsertCommand](#) and modify the [MySQLCommand.CommandTimeout](#) value, and then explicitly set that on the [MySQLDataAdapter](#).

After the SQL statement is first generated, the application must explicitly call [RefreshSchema](#) if it changes the statement in any way. Otherwise, the [GetInsertCommand](#) will be still be using information from the previous statement, which might not be correct. The SQL statements are first generated either when the application calls [System.Data.Common.DataAdapter.Update](#) or [GetInsertCommand](#).

### 18.2.3.2.11. GetUpdateCommand

Gets the automatically generated [MySQLCommand](#) object required to perform updates on the database.

**Returns:** The [MySQLCommand](#) object generated to handle update operations.

An application can use the [GetUpdateCommand](#) method for informational or troubleshooting purposes because it returns the [MySQLCommand](#) object to be executed.

You can also use [GetUpdateCommand](#) as the basis of a modified command. For example, you might call [GetUpdateCommand](#) and modify the [MySQLCommand.CommandTimeout](#) value, and then explicitly set that on the [MySQLDataAdapter](#).

After the SQL statement is first generated, the application must explicitly call [RefreshSchema](#) if it changes the statement in any way. Otherwise, the [GetUpdateCommand](#) will be still be using information from the previous statement, which might not be correct. The SQL statements are first generated either when the application calls [System.Data.Common.DataAdapter.Update](#) or [GetUpdateCommand](#).

### 18.2.3.2.12. RefreshSchema

Refreshes the database schema information used to generate INSERT, UPDATE, or DELETE statements.

An application should call [RefreshSchema](#) whenever the [SELECT](#) statement associated with the [MySQLCommandBuilder](#) changes.

An application should call [RefreshSchema](#) whenever the [MySQLDataAdapter.SelectCommand](#) value of the [MySQLDataAdapter](#) changes.

### 18.2.3.3. MySqlConnection

Represents an open connection to a MySQL Server database. This class cannot be inherited.

A [MySqlConnection](#) object represents a session to a MySQL Server data source. When you create an instance of [MySqlConnection](#), all properties are set to their initial values. For a list of these values, see the [MySqlConnection](#) constructor.

If the `MySqlConnection` goes out of scope, it is not closed. Therefore, you must explicitly close the connection by calling `MySqlConnection.Close` or `MySqlConnection.Dispose`.

### Examples

The following example creates a `MySqlCommand` and a `MySqlConnection`. The `MySqlConnection` is opened and set as the `MySqlCommand.Connection` for the `MySqlCommand`. The example then calls `MySqlCommand.ExecuteNonQuery`, and closes the connection. To accomplish this, the `ExecuteNonQuery` is passed a connection string and a query string that is a SQL INSERT statement.

Visual Basic example:

```
Public Sub InsertRow(myConnectionString As String)
    ' If the connection string is null, use a default.
    If myConnectionString = "" Then
        myConnectionString = "Database=Test;Data Source=localhost;User Id=username;Password=pass"
    End If
    Dim myConnection As New MySqlConnection(myConnectionString)
    Dim myInsertQuery As String = "INSERT INTO Orders (id, customerId, amount) Values(1001, 23, 30.66)"
    Dim myCommand As New MySqlCommand(myInsertQuery)
    myCommand.Connection = myConnection
    myConnection.Open()
    myCommand.ExecuteNonQuery()
    myCommand.Connection.Close()
End Sub
```

C# example:

```
public void InsertRow(string myConnectionString)
{
    // If the connection string is null, use a default.
    if(myConnectionString == "")
    {
        myConnectionString = "Database=Test;Data Source=localhost;User Id=username;Password=pass";
    }
    MySqlConnection myConnection = new MySqlConnection(myConnectionString);
    string myInsertQuery = "INSERT INTO Orders (id, customerId, amount) Values(1001, 23, 30.66)";
    MySqlCommand myCommand = new MySqlCommand(myInsertQuery);
    myCommand.Connection = myConnection;
    myConnection.Open();
    myCommand.ExecuteNonQuery();
    myCommand.Connection.Close();
}
```

#### 18.2.3.3.1. Class MySqlConnection Constructor (Default)

Initializes a new instance of the `MySqlConnection` class.

When a new instance of `MySqlConnection` is created, the read/write properties are set to the following initial values unless they are specifically set using their associated keywords in the `ConnectionString` property.

Properties	Initial Value
<code>ConnectionString</code>	empty string ("" )
<code>ConnectionTimeout</code>	15
<code>Database</code>	empty string ("" )
<code>DataSource</code>	empty string ("" )
<code>ServerVersion</code>	empty string ("" )

You can change the value for these properties only by using the [ConnectionString](#) property.

### Examples

#### Overload methods for MySqlConnection

Initializes a new instance of the [MySqlConnection](#) class.

#### 18.2.3.3.2. Class MySqlConnection Constructor Form 1

Initializes a new instance of the [MySqlConnection](#) class when given a string containing the connection string.

When a new instance of [MySqlConnection](#) is created, the read/write properties are set to the following initial values unless they are specifically set using their associated keywords in the [ConnectionString](#) property.

Properties	Initial Value
<a href="#">ConnectionString</a>	empty string ("" )
<a href="#">ConnectionTimeout</a>	15
<a href="#">Database</a>	empty string ("" )
<a href="#">DataSource</a>	empty string ("" )
<a href="#">ServerVersion</a>	empty string ("" )

You can change the value for these properties only by using the [ConnectionString](#) property.

### Examples

**Parameters:** The connection properties used to open the MySQL database.

#### 18.2.3.3.3. Open

Opens a database connection with the property settings specified by the [ConnectionString](#).

**Exception:** Cannot open a connection without specifying a data source or server.

**Exception:** A connection-level error occurred while opening the connection.

The [MySqlConnection](#) draws an open connection from the connection pool if one is available. Otherwise, it establishes a new connection to an instance of MySQL.

### Examples

The following example creates a [MySqlConnection](#), opens it, displays some of its properties, then closes the connection.

Visual Basic example:

```
Public Sub CreateMySqlConnection(myConnString As String)
    Dim myConnection As New MySqlConnection(myConnString)
    myConnection.Open()
    MessageBox.Show("ServerVersion: " + myConnection.ServerVersion _
        + ControlChars.Cr + "State: " + myConnection.State.ToString())
    myConnection.Close()
End Sub
```

C# example:

```
public void CreateMySQLConnection(string myConnString)
{
    MySqlConnection myConnection = new MySqlConnection(myConnString);
    myConnection.Open();
    MessageBox.Show("ServerVersion: " + myConnection.ServerVersion +
        "\nState: " + myConnection.State.ToString());
    myConnection.Close();
}
```

#### 18.2.3.3.4. Database

Gets the name of the current database or the database to be used after a connection is opened.

**Returns:** The name of the current database or the name of the database to be used after a connection is opened. The default value is an empty string.

The [Database](#) property does not update dynamically. If you change the current database using a SQL statement, then this property may reflect the wrong value. If you change the current database using the [ChangeDatabase](#) method, this property is updated to reflect the new database.

#### Examples

The following example creates a [MySQLConnection](#) and displays some of its read-only properties.

Visual Basic example:

```
Public Sub CreateMySQLConnection()
    Dim myConnString As String = _
        "Persist Security Info=False;database=test;server=localhost;user id=joeuser;pwd=pass"
    Dim myConnection As New MySqlConnection( myConnString )
    myConnection.Open()
    MessageBox.Show( "Server Version: " + myConnection.ServerVersion _
        + ControlChars.NewLine + "Database: " + myConnection.Database )
    myConnection.ChangeDatabase( "test2" )
    MessageBox.Show( "ServerVersion: " + myConnection.ServerVersion _
        + ControlChars.NewLine + "Database: " + myConnection.Database )
    myConnection.Close()
End Sub
```

C# example:

```
public void CreateMySQLConnection()
{
    string myConnString =
        "Persist Security Info=False;database=test;server=localhost;user id=joeuser;pwd=pass";
    MySqlConnection myConnection = new MySqlConnection( myConnString );
    myConnection.Open();
    MessageBox.Show( "Server Version: " + myConnection.ServerVersion
        + "\nDatabase: " + myConnection.Database );
    myConnection.ChangeDatabase( "test2" );
    MessageBox.Show( "ServerVersion: " + myConnection.ServerVersion
        + "\nDatabase: " + myConnection.Database );
    myConnection.Close();
}
```

#### 18.2.3.3.5. State

Gets the current state of the connection.

**Returns:** A bitwise combination of the [System.Data.ConnectionState](#) values. The default is [Closed](#).

The allowed state changes are:

- From [Closed](#) to [Open](#), using the [Open](#) method of the connection object.
- From [Open](#) to [Closed](#), using either the [Close](#) method or the [Dispose](#) method of the connection object.

### Examples

The following example creates a [MySQLConnection](#), opens it, displays some of its properties, then closes the connection.

Visual Basic example:

```
Public Sub CreateMySQLConnection(myConnString As String)
    Dim myConnection As New MySQLConnection(myConnString)
    myConnection.Open()
    MessageBox.Show("ServerVersion: " + myConnection.ServerVersion _
        + ControlChars.Cr + "State: " + myConnection.State.ToString())
    myConnection.Close()
End Sub
```

C# example:

```
public void CreateMySQLConnection(string myConnString)
{
    MySqlConnection myConnection = new MySqlConnection(myConnString);
    myConnection.Open();
    MessageBox.Show("ServerVersion: " + myConnection.ServerVersion +
        "\nState: " + myConnection.State.ToString());
    myConnection.Close();
}
```

### 18.2.3.3.6. ServerVersion

Gets a string containing the version of the MySQL server to which the client is connected.

**Returns:** The version of the instance of MySQL.

**Exception:** The connection is closed.

### Examples

The following example creates a [MySQLConnection](#), opens it, displays some of its properties, then closes the connection.

Visual Basic example:

```
Public Sub CreateMySQLConnection(myConnString As String)
    Dim myConnection As New MySQLConnection(myConnString)
    myConnection.Open()
    MessageBox.Show("ServerVersion: " + myConnection.ServerVersion _
        + ControlChars.Cr + "State: " + myConnection.State.ToString())
    myConnection.Close()
End Sub
```

C# example:

```
public void CreateMySQLConnection(string myConnString)
{
```

```

MySQLConnection myConnection = new MySQLConnection(myConnString);
myConnection.Open();
MessageBox.Show("ServerVersion: " + myConnection.ServerVersion +
"\nState: " + myConnection.State.ToString());
myConnection.Close();
}

```

### 18.2.3.3.7. Close

Closes the connection to the database. This is the preferred method of closing any open connection.

The `Close` method rolls back any pending transactions. It then releases the connection to the connection pool, or closes the connection if connection pooling is disabled.

An application can call `Close` more than one time. No exception is generated.

#### Examples

The following example creates a `MySQLConnection`, opens it, displays some of its properties, then closes the connection.

Visual Basic example:

```

Public Sub CreateMySQLConnection(myConnString As String)
    Dim myConnection As New MySQLConnection(myConnString)
    myConnection.Open()
    MessageBox.Show("ServerVersion: " + myConnection.ServerVersion _
+ ControlChars.Cr + "State: " + myConnection.State.ToString())
    myConnection.Close()
End Sub

```

C# example:

```

public void CreateMySQLConnection(string myConnString)
{
    MySQLConnection myConnection = new MySQLConnection(myConnString);
    myConnection.Open();
    MessageBox.Show("ServerVersion: " + myConnection.ServerVersion +
"\nState: " + myConnection.State.ToString());
    myConnection.Close();
}

```

### 18.2.3.3.8. CreateCommand

Creates and returns a `MySQLCommand` object associated with the `MySQLConnection`.

**Returns:** A `MySQLCommand` object.

### 18.2.3.3.9. BeginTransaction

Begins a database transaction.

**Returns:** An object representing the new transaction.

**Exception:** Parallel transactions are not supported.

This command is equivalent to the MySQL `BEGIN TRANSACTION` command.

You must explicitly commit or roll back the transaction using the `MySQLTransaction.Commit` or `MySQLTransaction.Rollback` method.

**Note.** If you do not specify an isolation level, the default isolation level is used. To specify an isolation level with the [BeginTransaction](#) method, use the overload that takes the `iso` parameter.

### Examples

The following example creates a [MySQLConnection](#) and a [MySQLTransaction](#). It also demonstrates how to use the [BeginTransaction](#), a [MySQLTransaction.Commit](#), and [MySQLTransaction.Rollback](#) methods.

Visual Basic example:

```
Public Sub RunTransaction(myConnString As String)
    Dim myConnection As New MySqlConnection(myConnString)
    myConnection.Open()

    Dim myCommand As MySqlCommand = myConnection.CreateCommand()
    Dim myTrans As MySQLTransaction

    ' Start a local transaction
    myTrans = myConnection.BeginTransaction()
    ' Must assign both transaction object and connection
    ' to Command object for a pending local transaction
    myCommand.Connection = myConnection
    myCommand.Transaction = myTrans

    Try
        myCommand.CommandText = "Insert into Test (id, desc) VALUES (100, 'Description')"
        myCommand.ExecuteNonQuery()
        myCommand.CommandText = "Insert into Test (id, desc) VALUES (101, 'Description')"
        myCommand.ExecuteNonQuery()
        myTrans.Commit()
        Console.WriteLine("Both records are written to database.")
    Catch e As Exception
        Try
            myTrans.Rollback()
        Catch ex As MySQLException
            If Not myTrans.Connection Is Nothing Then
                Console.WriteLine("An exception of type " + ex.GetType().ToString() + _
                    " was encountered while attempting to roll back the transaction.")
            End If
        End Try

        Console.WriteLine("An exception of type " + e.GetType().ToString() + _
            "was encountered while inserting the data.")
        Console.WriteLine("Neither record was written to database.")
    Finally
        myConnection.Close()
    End Try
End Sub
```

C# example:

```
public void RunTransaction(string myConnString)
{
    MySqlConnection myConnection = new MySqlConnection(myConnString);
    myConnection.Open();
    MySqlCommand myCommand = myConnection.CreateCommand();
    MySQLTransaction myTrans;
    // Start a local transaction
    myTrans = myConnection.BeginTransaction();
    // Must assign both transaction object and connection
    // to Command object for a pending local transaction
    myCommand.Connection = myConnection;
    myCommand.Transaction = myTrans;
    try
    {
        myCommand.CommandText = "insert into Test (id, desc) VALUES (100, 'Description')";
        myCommand.ExecuteNonQuery();
        myCommand.CommandText = "insert into Test (id, desc) VALUES (101, 'Description')";
        myCommand.ExecuteNonQuery();
        myTrans.Commit();
        Console.WriteLine("Both records are written to database.");
    }
    catch(Exception e)
```



```

{
    try
    {
        myTrans.Rollback();
    }
    catch (SqlException ex)
    {
        if (myTrans.Connection != null)
        {
            Console.WriteLine("An exception of type " + ex.GetType() +
                " was encountered while attempting to roll back the transaction.");
        }
    }

    Console.WriteLine("An exception of type " + e.GetType() +
        " was encountered while inserting the data.");
    Console.WriteLine("Neither record was written to database.");
}
finally
{
    myConnection.Close();
}
}

```

### 18.2.3.3.10. BeginTransaction1

Begins a database transaction with the specified isolation level.

**Parameters:** The isolation level under which the transaction should run.

**Returns:** An object representing the new transaction.

**Exception:** Parallel exceptions are not supported.

This command is equivalent to the MySQL BEGIN TRANSACTION command.

You must explicitly commit or roll back the transaction using the [MySQLTransaction.Commit](#) or [MySQLTransaction.Rollback](#) method.

**Note.** If you do not specify an isolation level, the default isolation level is used. To specify an isolation level with the [BeginTransaction](#) method, use the overload that takes the `iso` parameter.

#### Examples

The following example creates a [MySQLConnection](#) and a [MySQLTransaction](#). It also demonstrates how to use the [BeginTransaction](#), a [MySQLTransaction.Commit](#), and [MySQLTransaction.Rollback](#) methods.

Visual Basic example:

```

Public Sub RunTransaction(myConnString As String)
    Dim myConnection As New MySqlConnection(myConnString)
    myConnection.Open()

    Dim myCommand As MySqlCommand = myConnection.CreateCommand()
    Dim myTrans As MySQLTransaction

    ' Start a local transaction
    myTrans = myConnection.BeginTransaction()
    ' Must assign both transaction object and connection
    ' to Command object for a pending local transaction
    myCommand.Connection = myConnection
    myCommand.Transaction = myTrans

    Try
        myCommand.CommandText = "Insert into Test (id, desc) VALUES (100, 'Description')"
        myCommand.ExecuteNonQuery()
        myCommand.CommandText = "Insert into Test (id, desc) VALUES (101, 'Description')"
        myCommand.ExecuteNonQuery()
    
```

```

    myTrans.Commit()
    Console.WriteLine("Both records are written to database.")
Catch e As Exception
    Try
        myTrans.Rollback()
    Catch ex As MySqlException
        If Not myTrans.Connection Is Nothing Then
            Console.WriteLine("An exception of type " + ex.GetType().ToString() + _
                " was encountered while attempting to roll back the transaction.")
        End If
    End Try

    Console.WriteLine("An exception of type " + e.GetType().ToString() + _
        " was encountered while inserting the data.")
    Console.WriteLine("Neither record was written to database.")
Finally
    myConnection.Close()
End Try
End Sub

```

C# example:

```

public void RunTransaction(string myConnString)
{
    MySqlConnection myConnection = new MySqlConnection(myConnString);
    myConnection.Open();
    MySqlCommand myCommand = myConnection.CreateCommand();
    MySqlTransaction myTrans;
    // Start a local transaction
    myTrans = myConnection.BeginTransaction();
    // Must assign both transaction object and connection
    // to Command object for a pending local transaction
    myCommand.Connection = myConnection;
    myCommand.Transaction = myTrans;
    try
    {
        myCommand.CommandText = "insert into Test (id, desc) VALUES (100, 'Description')";
        myCommand.ExecuteNonQuery();
        myCommand.CommandText = "insert into Test (id, desc) VALUES (101, 'Description')";
        myCommand.ExecuteNonQuery();
        myTrans.Commit();
        Console.WriteLine("Both records are written to database.");
    }
    catch(Exception e)
    {
        try
        {
            myTrans.Rollback();
        }
        catch (SqlException ex)
        {
            if (myTrans.Connection != null)
            {
                Console.WriteLine("An exception of type " + ex.GetType() +
                    " was encountered while attempting to roll back the transaction.");
            }
        }

        Console.WriteLine("An exception of type " + e.GetType() +
            " was encountered while inserting the data.");
        Console.WriteLine("Neither record was written to database.");
    }
    finally
    {
        myConnection.Close();
    }
}

```

### 18.2.3.3.11. ChangeDatabase

Changes the current database for an open MySqlConnection.

**Parameters:** The name of the database to use.

The value supplied in the `database` parameter must be a valid database name. The `database` parameter cannot contain a null value, an empty string, or a string with only blank characters.

When you are using connection pooling against MySQL, and you close the connection, it is returned to the connection pool. The next time the connection is retrieved from the pool, the reset connection request executes before the user performs any operations.

**Exception:** The database name is not valid.

**Exception:** The connection is not open.

**Exception:** Cannot change the database.

### Examples

The following example creates a `MySqlConnection` and displays some of its read-only properties.

Visual Basic example:

```
Public Sub CreateMySQLConnection()
    Dim myConnString As String = _
        "Persist Security Info=False;database=test;server=localhost;user id=joeuser;pwd=pass"
    Dim myConnection As New MySqlConnection( myConnString )
    myConnection.Open()
    MessageBox.Show( "Server Version: " + myConnection.ServerVersion _
        + ControlChars.NewLine + "Database: " + myConnection.Database )
    myConnection.ChangeDatabase( "test2" )
    MessageBox.Show( "ServerVersion: " + myConnection.ServerVersion _
        + ControlChars.NewLine + "Database: " + myConnection.Database )
    myConnection.Close()
End Sub
```

C# example:

```
public void CreateMySQLConnection()
{
    string myConnString =
        "Persist Security Info=False;database=test;server=localhost;user id=joeuser;pwd=pass";
    MySqlConnection myConnection = new MySqlConnection( myConnString );
    myConnection.Open();
    MessageBox.Show( "Server Version: " + myConnection.ServerVersion
        + "\nDatabase: " + myConnection.Database );
    myConnection.ChangeDatabase( "test2" );
    MessageBox.Show( "ServerVersion: " + myConnection.ServerVersion
        + "\nDatabase: " + myConnection.Database );
    myConnection.Close();
}
```

### 18.2.3.3.12. StateChange

Occurs when the state of the connection changes.

The `StateChange` event fires whenever the `State` changes from closed to opened, or from opened to closed. `StateChange` fires immediately after the `MySqlConnection` transitions.

If an event handler throws an exception from within the `StateChange` event, the exception propagates to the caller of the `Open` or `Close` method.

The `StateChange` event is not raised unless you explicitly call `Close` or `Dispose`.

The event handler receives an argument of type `System.Data.StateChangeEventArgs` containing data related to this event. The following `StateChangeEventArgs` properties provide information specific to this event.

Property	Description
<code>System.Data.StateChangeEventArgs.CurrentState</code>	Gets the new state of the connection. The connection object will be in the new state already when the event is fired.
<code>System.Data.StateChangeEventArgs.OriginalState</code>	Gets the original state of the connection.

### 18.2.3.3.13. InfoMessage

Occurs when MySQL returns warnings as a result of executing a command or query.

### 18.2.3.3.14. ConnectionTimeout

Gets the time to wait while trying to establish a connection before terminating the attempt and generating an error.

**Exception:** The value set is less than 0.

A value of 0 indicates no limit, and should be avoided in a `MySQLConnection.ConnectionString` because an attempt to connect will wait indefinitely.

#### Examples

The following example creates a `MySQLConnection` and sets some of its properties in the connection string.

Visual Basic example:

```
Public Sub CreateSqlConnection()
    Dim myConnection As New MySqlConnection()
    myConnection.ConnectionString = "Persist Security Info=False;Username=user;Password=pass;database=test"
    myConnection.Open()
End Sub
```

C# example:

```
public void CreateSqlConnection()
{
    MySqlConnection myConnection = new MySqlConnection();
    myConnection.ConnectionString = "Persist Security Info=False;Username=user;Password=pass;database=test";
    myConnection.Open();
}
```

### 18.2.3.3.15. ConnectionString

Gets or sets the string used to connect to a MySQL Server database.

The `ConnectionString` returned may not be exactly like what was originally set but will be identical in terms of keyword/value pairs. Security information will not be included unless the `Persist Security Info` value is set to true.

You can use the `ConnectionString` property to connect to a database. The following example illustrates a typical connection string.

```
"Persist Security Info=False;database=MyDB;server=MySQLServer;user id=myUser;Password=myPass"
```

The `ConnectionString` property can be set only when the connection is closed. Many of the connection string values have corresponding read-only properties. When the connection string is set, all of these properties are updated, except when an error is detected. In this case, none of the properties are updated. `MySQLConnection` properties return only those settings contained in the `ConnectionString`.

To connect to a local machine, specify "localhost" for the server. If you do not specify a server, localhost is assumed.

Resetting the `ConnectionString` on a closed connection resets all connection string values (and related properties) including the password. For example, if you set a connection string that includes "Database= MyDb", and then reset the connection string to "Data Source=myserver;UserId=myUser;Password=myPass", the `MySQLConnection.Database` property is no longer set to MyDb.

The connection string is parsed immediately after being set. If errors in syntax are found when parsing, a runtime exception, such as `ArgumentException`, is generated. Other errors can be found only when an attempt is made to open the connection.

The basic format of a connection string consists of a series of keyword/value pairs separated by semicolons. The equal sign (=) connects each keyword and its value. To include values that contain a semicolon, single-quote character, or double-quote character, the value must be enclosed in double quotes. If the value contains both a semicolon and a double-quote character, the value can be enclosed in single quotes. The single quote is also useful if the value begins with a double-quote character. Conversely, the double quote can be used if the value begins with a single quote. If the value contains both single-quote and double-quote characters, the quote character used to enclose the value must be doubled each time it occurs within the value.

To include preceding or trailing spaces in the string value, the value must be enclosed in either single quotes or double quotes. Any leading or trailing spaces around integer, Boolean, or enumerated values are ignored, even if enclosed in quotes. However, spaces within a string literal keyword or value are preserved. Using .NET Framework version 1.1, single or double quotes may be used within a connection string without using delimiters (for example, Data Source= my'Server or Data Source= my"Server), unless a quote character is the first or last character in the value.

To include an equal sign (=) in a keyword or value, it must be preceded by another equal sign. For example, in the hypothetical connection string

```
"key==word=value"
```

the keyword is "key=word" and the value is "value".

If a specific keyword in a keyword= value pair occurs multiple times in a connection string, the last occurrence listed is used in the value set.

Keywords are not case sensitive.

The following table lists the valid names for keyword values within the `ConnectionString`.

Name	Default	Description
Connect Timeout -or- Connection Timeout	15	The length of time (in seconds) to wait for a connection to the server before terminating the attempt and generating an error.
Host -or- Server -or- Data Source -or- DataSource -or- Address -or- Addr -or- Network Address	localhost	The name or network address of the instance of MySQL to which to connect. Multiple hosts can be

		specified separated by &. This can be useful where multiple MySQL servers are configured for replication and you are not concerned about the precise server you are connecting to. No attempt is made by the provider to synchronize writes to the database so care should be taken when using this option. In Unix environment with Mono, this can be a fully qualified path to MySQL socket filename. With this configuration, the Unix socket will be used instead of TCP/IP socket. Currently only a single socket name can be given so accessing MySQL in a replicated environment using Unix sockets is not currently supported.
Port	3306	The port MySQL is using to listen for connections. Specify -1 for this value to use a named pipe connection (Windows only). This value is ignored if Unix socket is used.
Protocol	socket	Specifies the type of connection to make to the server. Values can be: socket or tcp for a socket connection pipe for a named pipe connection unix for a Unix socket connection memory to use MySQL shared memory
CharSet -or Character Set		Specifies the character set that should be used to encode all queries sent to the server. Result-sets are still returned in the character set of the data returned.
Logging	false	When true, various pieces of information is output to any configured TraceListeners.
Allow Batch	true	When true, multiple SQL statements can be sent with one command execution. -Note- Starting with MySQL 4.1.1, batch statements should be separated by the server-defined separator character. Commands sent to earlier versions of MySQL should be separated with ';'.
Encrypt	false	When <code>true</code> , SSL encryption is used for all data sent between the client and server if the server has a certificate installed. Recognized values are <code>true</code> , <code>false</code> , <code>yes</code> ,

		and <a href="#">no</a> . <a href="#">Note</a> This parameter currently has no effect.
Initial Catalog -or- Database	mysql	The name of the database to use initially
Password -or- pwd		The password for the MySQL account being used.
Persist Security Info	false	When set to <a href="#">false</a> or <a href="#">no</a> (strongly recommended), security-sensitive information, such as the password, is not returned as part of the connection if the connection is open or has ever been in an open state. Resetting the connection string resets all connection string values including the password. Recognized values are <a href="#">true</a> , <a href="#">false</a> , <a href="#">yes</a> , and <a href="#">no</a> .
User Id -or- Username -or- Uid -or- User name		The MySQL login account being used.
Shared Memory Name	MYSQL	The name of the shared memory object to use for communication if the connection protocol is set to memory.
Allow Zero Datetime	false	True to have <a href="#">MySqlDataReader.GetValue()</a> return a <a href="#">MySqlDateTime</a> for date or datetime columns that have illegal values. False will cause a <a href="#">DateTime</a> object to be returned for legal values and an exception will be thrown for illegal values.
Convert Zero Datetime	false	True to have <a href="#">MySqlDataReader.GetValue()</a> and <a href="#">MySqlDataReader.GetDateTime()</a> return <a href="#">DateTime.MinValue</a> for date or datetime columns that have illegal values.
Old Syntax -or- OldSyntax	false	Allows use of '@' symbol as a parameter marker. See <a href="#">MySQL-Command</a> for more info. This is for compatibility only. All future code should be written to use the new '?' parameter marker.
Pipe Name -or- Pipe	mysql	When set to the name of a named pipe, the <a href="#">MySqlConnection</a> will attempt to connect to MySQL on that named pipe. This settings only applies to the Windows platform.

The following table lists the valid names for connection pooling values within the [ConnectionString](#). For more information about connection pooling, see [Connection Pooling for the MySQL Data](#)

Provider.

Name	Default	Description
Connection Lifetime	0	When a connection is returned to the pool, its creation time is compared with the current time, and the connection is destroyed if that time span (in seconds) exceeds the value specified by <code>Connection Lifetime</code> . This is useful in clustered configurations to force load balancing between a running server and a server just brought online. A value of zero (0) causes pooled connections to have the maximum connection timeout.
Max Pool Size	100	The maximum number of connections allowed in the pool.
Min Pool Size	0	The minimum number of connections allowed in the pool.
Pooling	true	When <code>true</code> , the <code>MySQLConnection</code> object is drawn from the appropriate pool, or if necessary, is created and added to the appropriate pool. Recognized values are <code>true</code> , <code>false</code> , <code>yes</code> , and <code>no</code> .
Reset Pooled Connections -or- ResetConnections -or- Reset- PooledConnections	true	Specifies whether a ping and a reset should be sent to the server before a pooled connection is returned. Not resetting will yeild faster connection opens but also will not clear out session items such as temp tables.
Cache Server Configuration -or- CacheServerConfiguration -or- CacheServerConfig	false	Specifies whether server variables should be updated when a pooled connection is returned. Turning this one will yeild faster opens but will also not catch any server changes made by other connections.

When setting keyword or connection pooling values that require a Boolean value, you can use 'yes' instead of 'true', and 'no' instead of 'false'.

**Note** The MySQL Data Provider uses the native socket protocol to communicate with MySQL. Therefore, it does not support the use of an ODBC data source name (DSN) when connecting to MySQL because it does not add an ODBC layer.

**CAUTION** In this release, the application should use caution when constructing a connection string based on user input (for example when retrieving user ID and password information from a dialog box, and appending it to the connection string). The application should ensure that a user cannot embed extra connection string parameters in these values (for example, entering a password as "validpassword;database=somedb" in an attempt to attach to a different database).



## Examples

The following example creates a [MySQLConnection](#) and sets some of its properties

Visual Basic example:

```
Public Sub CreateConnection()
    Dim myConnection As New MySqlConnection()
    myConnection.ConnectionString = "Persist Security Info=False;database=myDB;server=myHost;Connect T
    myConnection.Open()
End Sub 'CreateConnection
```

C# example:

```
public void CreateConnection()
{
    MySqlConnection myConnection = new MySqlConnection();
    myConnection.ConnectionString = "Persist Security Info=False;database=myDB;server=myHost;Connect T
    myConnection.Open();
}
```

## Examples

The following example creates a [MySQLConnection](#) in Unix environment with Mono installed. MySQL socket filename used in this example is `"/var/lib/mysql/mysql.sock"`. The actual filename depends on your MySQL configuration.

Visual Basic example:

```
Public Sub CreateConnection()
    Dim myConnection As New MySqlConnection()
    myConnection.ConnectionString = "database=myDB;server=/var/lib/mysql/mysql.sock;user id=myUser; pwo
    myConnection.Open()
End Sub 'CreateConnection
```

C# example:

```
public void CreateConnection()
{
    MySqlConnection myConnection = new MySqlConnection();
    myConnection.ConnectionString = "database=myDB;server=/var/lib/mysql/mysql.sock;user id=myUser; pwo
    myConnection.Open();
}
```

### 18.2.3.4. MySqlConnectionAdapter

Represents a set of data commands and a database connection that are used to fill a dataset and update a MySQL database. This class cannot be inherited.

The [MySQLDataAdapter](#), serves as a bridge between a [System.Data.DataSet](#) and MySQL for retrieving and saving data. The [MySQLDataAdapter](#) provides this bridge by mapping [DbDataAdapter.Fill](#), which changes the data in the [DataSet](#) to match the data in the data source, and [DbDataAdapter.Update](#), which changes the data in the data source to match the data in the [DataSet](#), using the appropriate SQL statements against the data source.

When the [MySQLDataAdapter](#) fills a [DataSet](#), it will create the necessary tables and columns for the returned data if they do not already exist. However, primary key information will not be included in the implicitly created schema unless the [System.Data.MissingSchemaAction](#) property is set to [System.Data.MissingSchemaAction.AddWithKey](#). You may also have the

`MySQLDataAdapter` create the schema of the `DataSet`, including primary key information, before filling it with data using `System.Data.Common.DbDataAdapter.FillSchema`.

`MySQLDataAdapter` is used in conjunction with `MySqlConnection` and `MySqlCommand` to increase performance when connecting to a MySQL database.

The `MySQLDataAdapter` also includes the `MySqlCommand.SelectCommand`, `MySqlCommand.InsertCommand`, `MySqlCommand.DeleteCommand`, `MySqlCommand.UpdateCommand`, and `DataAdapter.TableMappings` properties to facilitate the loading and updating of data.

When an instance of `MySQLDataAdapter` is created, the read/write properties are set to initial values. For a list of these values, see the `MySQLDataAdapter` constructor.

**Note.** Please be aware that the `DataColumn` class in .NET 1.0 and 1.1 does not allow columns with type of `UInt16`, `UInt32`, or `UInt64` to be autoincrement columns. If you plan to use autoincrement columns with MySQL, you should consider using signed integer columns.

### Examples

The following example creates a `MySqlCommand` and a `MySqlConnection`. The `MySqlConnection` is opened and set as the `MySqlCommand.Connection` for the `MySqlCommand`. The example then calls `MySqlCommand.ExecuteNonQuery`, and closes the connection. To accomplish this, the `ExecuteNonQuery` is passed a connection string and a query string that is a SQL INSERT statement.

Visual Basic example:

```
Public Function SelectRows(dataSet As DataSet, connection As String, query As String) As DataSet
    Dim conn As New MySqlConnection(connection)
    Dim adapter As New MySQLDataAdapter()
    adapter.SelectCommand = new MySqlCommand(query, conn)
    adapter.Fill(dataset)
    Return dataset
End Function
```

C# example:

```
public DataSet SelectRows(DataSet dataset,string connection,string query)
{
    MySqlConnection conn = new MySqlConnection(connection);
    MySQLDataAdapter adapter = new MySQLDataAdapter();
    adapter.SelectCommand = new MySqlCommand(query, conn);
    adapter.Fill(dataset);
    return dataset;
}
```

#### 18.2.3.4.1. Class `MySqlCommand` Constructor

##### Overload methods for `MySqlCommand`

Initializes a new instance of the `MySqlCommand` class.

When an instance of `MySqlCommand` is created, the following read/write properties are set to the following initial values.

Properties	Initial Value
<code>MissingMappingAction</code>	<code>MissingMappingAction.Passthrough</code>
<code>MissingSchemaAction</code>	<code>MissingSchemaAction.Add</code>

You can change the value of any of these properties through a separate call to the property.

### Examples

The following example creates a [MySQLDataAdapter](#) and sets some of its properties.

Visual Basic example:

```
Public Sub CreateSqlDataAdapter()
    Dim conn As MySqlConnection = New MySqlConnection("Data Source=localhost;" & _
        "database=test")
    Dim da As MySQLDataAdapter = New MySQLDataAdapter
    da.MissingSchemaAction = MissingSchemaAction.AddWithKey

    da.SelectCommand = New MySqlCommand("SELECT id, name FROM mytable", conn)
    da.InsertCommand = New MySqlCommand("INSERT INTO mytable (id, name) " & _
        "VALUES (?id, ?name)", conn)
    da.UpdateCommand = New MySqlCommand("UPDATE mytable SET id=?id, name=?name " & _
        "WHERE id=?oldId", conn)
    da.DeleteCommand = New MySqlCommand("DELETE FROM mytable WHERE id=?id", conn)
    da.InsertCommand.Parameters.Add("?id", MySqlDbType.VarChar, 5, "id")
    da.InsertCommand.Parameters.Add("?name", MySqlDbType.VarChar, 40, "name")

    da.UpdateCommand.Parameters.Add("?id", MySqlDbType.VarChar, 5, "id")
    da.UpdateCommand.Parameters.Add("?name", MySqlDbType.VarChar, 40, "name")
    da.UpdateCommand.Parameters.Add("?oldId", MySqlDbType.VarChar, 5, "id").SourceVersion = DataRowVersion.Original
    da.DeleteCommand.Parameters.Add("?id", MySqlDbType.VarChar, 5, "id").SourceVersion = DataRowVersion.Original
End Sub
```

C# example:

```
public static void CreateSqlDataAdapter()
{
    MySqlConnection conn = new MySqlConnection("Data Source=localhost;database=test");
    MySQLDataAdapter da = new MySQLDataAdapter();
    da.MissingSchemaAction = MissingSchemaAction.AddWithKey;

    da.SelectCommand = new MySqlCommand("SELECT id, name FROM mytable", conn);
    da.InsertCommand = new MySqlCommand("INSERT INTO mytable (id, name) " +
        "VALUES (?id, ?name)", conn);
    da.UpdateCommand = new MySqlCommand("UPDATE mytable SET id=?id, name=?name " +
        "WHERE id=?oldId", conn);
    da.DeleteCommand = new MySqlCommand("DELETE FROM mytable WHERE id=?id", conn);
    da.InsertCommand.Parameters.Add("?id", MySqlDbType.VarChar, 5, "id");
    da.InsertCommand.Parameters.Add("?name", MySqlDbType.VarChar, 40, "name");

    da.UpdateCommand.Parameters.Add("?id", MySqlDbType.VarChar, 5, "id");
    da.UpdateCommand.Parameters.Add("?name", MySqlDbType.VarChar, 40, "name");
    da.UpdateCommand.Parameters.Add("?oldId", MySqlDbType.VarChar, 5, "id").SourceVersion = DataRowVersion.Original
    da.DeleteCommand.Parameters.Add("?id", MySqlDbType.VarChar, 5, "id").SourceVersion = DataRowVersion.Original
}
```

#### 18.2.3.4.2. Class MySQLDataAdapter Constructor Form 1

Initializes a new instance of the [MySQLDataAdapter](#) class with the specified [MySQLCommand](#) as the [SelectCommand](#) property.

**Parameters:** [MySQLCommand](#) that is a SQL [SELECT](#) statement or stored procedure and is set as the [SelectCommand](#) property of the [MySQLDataAdapter](#).

When an instance of [MySQLDataAdapter](#) is created, the following read/write properties are set to the following initial values.

Properties	Initial Value
<a href="#">MissingMappingAction</a>	<a href="#">MissingMappingAction.Passthrough</a>
<a href="#">MissingSchemaAction</a>	<a href="#">MissingSchemaAction.Add</a>

You can change the value of any of these properties through a separate call to the property.

When [SelectCommand](#) (or any of the other command properties) is assigned to a previously created [MySQLCommand](#), the [MySQLCommand](#) is not cloned. The [SelectCommand](#) maintains a reference to the previously created [MySQLCommand](#) object.

### Examples

The following example creates a [MySQLDataAdapter](#) and sets some of its properties.

Visual Basic example:

```
Public Sub CreateSqlDataAdapter()
    Dim conn As MySqlConnection = New MySqlConnection("Data Source=localhost;" & _
        "database=test")
    Dim cmd As new MySqlCommand("SELECT id, name FROM mytable", conn)
    Dim da As MySQLDataAdapter = New MySQLDataAdapter(cmd)
    da.MissingSchemaAction = MissingSchemaAction.AddWithKey

    da.InsertCommand = New MySqlCommand("INSERT INTO mytable (id, name) " & _
        "VALUES (?id, ?name)", conn)
    da.UpdateCommand = New MySqlCommand("UPDATE mytable SET id=?id, name=?name " & _
        "WHERE id=?oldId", conn)
    da.DeleteCommand = New MySqlCommand("DELETE FROM mytable WHERE id=?id", conn)
    da.InsertCommand.Parameters.Add("?id", MySqlDbType.VarChar, 5, "id")
    da.InsertCommand.Parameters.Add("?name", MySqlDbType.VarChar, 40, "name")

    da.UpdateCommand.Parameters.Add("?id", MySqlDbType.VarChar, 5, "id")
    da.UpdateCommand.Parameters.Add("?name", MySqlDbType.VarChar, 40, "name")
    da.UpdateCommand.Parameters.Add("?oldId", MySqlDbType.VarChar, 5, "id").SourceVersion = DataRowVersion.Original
    da.DeleteCommand.Parameters.Add("?id", MySqlDbType.VarChar, 5, "id").SourceVersion = DataRowVersion.Original
End Sub
```

C# example:

```
public static void CreateSqlDataAdapter()
{
    MySqlConnection conn = new MySqlConnection("Data Source=localhost;database=test");
    MySqlCommand cmd = new MySqlCommand("SELECT id, name FROM mytable", conn);
    MySQLDataAdapter da = new MySQLDataAdapter(cmd);
    da.MissingSchemaAction = MissingSchemaAction.AddWithKey;

    da.InsertCommand = new MySqlCommand("INSERT INTO mytable (id, name) " +
        "VALUES (?id, ?name)", conn);
    da.UpdateCommand = new MySqlCommand("UPDATE mytable SET id=?id, name=?name " +
        "WHERE id=?oldId", conn);
    da.DeleteCommand = new MySqlCommand("DELETE FROM mytable WHERE id=?id", conn);
    da.InsertCommand.Parameters.Add("?id", MySqlDbType.VarChar, 5, "id");
    da.InsertCommand.Parameters.Add("?name", MySqlDbType.VarChar, 40, "name");

    da.UpdateCommand.Parameters.Add("?id", MySqlDbType.VarChar, 5, "id");
    da.UpdateCommand.Parameters.Add("?name", MySqlDbType.VarChar, 40, "name");
    da.UpdateCommand.Parameters.Add("?oldId", MySqlDbType.VarChar, 5, "id").SourceVersion = DataRowVersion.Original;
    da.DeleteCommand.Parameters.Add("?id", MySqlDbType.VarChar, 5, "id").SourceVersion = DataRowVersion.Original;
}
```

#### 18.2.3.4.3. Class MySQLDataAdapter Constructor Form 2

Initializes a new instance of the [MySQLDataAdapter](#) class with a [SelectCommand](#) and a [MySQLConnection](#) object.

**Parameters:** A [String](#) that is a SQL [SELECT](#) statement or stored procedure to be used by the [SelectCommand](#) property of the [MySQLDataAdapter](#).

**Parameters:** A [MySQLConnection](#) that represents the connection.

This implementation of the [MySQLDataAdapter](#) opens and closes a [MySQLConnection](#) if it is not

already open. This can be useful in an application that must call the `DbDataAdapter.Fill` method for two or more `MySqlConnection` objects. If the `MySqlConnection` is already open, you must explicitly call `MySqlConnection.Close` or `MySqlConnection.Dispose` to close it.

When an instance of `MySqlConnection` is created, the following read/write properties are set to the following initial values.

Properties	Initial Value
<code>MissingMappingAction</code>	<code>MissingMappingAction.Passthrough</code>
<code>MissingSchemaAction</code>	<code>MissingSchemaAction.Add</code>

You can change the value of any of these properties through a separate call to the property.

### Examples

The following example creates a `MySqlConnection` and sets some of its properties.

Visual Basic example:

```
Public Sub CreateSqlConnection()
    Dim conn As MySqlConnection = New MySqlConnection("Data Source=localhost;" & _
        "database=test")
    Dim da As MySqlDataAdapter = New MySqlDataAdapter("SELECT id, name FROM mytable", conn)
    da.MissingSchemaAction = MissingSchemaAction.AddWithKey

    da.InsertCommand = New MySqlCommand("INSERT INTO mytable (id, name) " & _
        "VALUES (?id, ?name)", conn)
    da.UpdateCommand = New MySqlCommand("UPDATE mytable SET id=?id, name=?name " & _
        "WHERE id=?oldId", conn)
    da.DeleteCommand = New MySqlCommand("DELETE FROM mytable WHERE id=?id", conn)
    da.InsertCommand.Parameters.Add("?id", MySqlDbType.VarChar, 5, "id")
    da.InsertCommand.Parameters.Add("?name", MySqlDbType.VarChar, 40, "name")

    da.UpdateCommand.Parameters.Add("?id", MySqlDbType.VarChar, 5, "id")
    da.UpdateCommand.Parameters.Add("?name", MySqlDbType.VarChar, 40, "name")
    da.UpdateCommand.Parameters.Add("?oldId", MySqlDbType.VarChar, 5, "id").SourceVersion = DataRowVersion.Original
    da.DeleteCommand.Parameters.Add("?id", MySqlDbType.VarChar, 5, "id").SourceVersion = DataRowVersion.Original
End Sub
```

C# example:

```
public static void CreateSqlConnection()
{
    MySqlConnection conn = new MySqlConnection("Data Source=localhost;database=test");
    MySqlDataAdapter da = new MySqlDataAdapter("SELECT id, name FROM mytable", conn);
    da.MissingSchemaAction = MissingSchemaAction.AddWithKey;

    da.InsertCommand = new MySqlCommand("INSERT INTO mytable (id, name) " +
        "VALUES (?id, ?name)", conn);
    da.UpdateCommand = new MySqlCommand("UPDATE mytable SET id=?id, name=?name " +
        "WHERE id=?oldId", conn);
    da.DeleteCommand = new MySqlCommand("DELETE FROM mytable WHERE id=?id", conn);
    da.InsertCommand.Parameters.Add("?id", MySqlDbType.VarChar, 5, "id");
    da.InsertCommand.Parameters.Add("?name", MySqlDbType.VarChar, 40, "name");

    da.UpdateCommand.Parameters.Add("?id", MySqlDbType.VarChar, 5, "id");
    da.UpdateCommand.Parameters.Add("?name", MySqlDbType.VarChar, 40, "name");
    da.UpdateCommand.Parameters.Add("?oldId", MySqlDbType.VarChar, 5, "id").SourceVersion = DataRowVersion.Original;
    da.DeleteCommand.Parameters.Add("?id", MySqlDbType.VarChar, 5, "id").SourceVersion = DataRowVersion.Original;
}
```

#### 18.2.3.4.4. Class MySqlConnection Constructor Form 3

Initializes a new instance of the `MySqlConnection` class with a `SelectCommand` and a connection string.

**Parameters:** A [string](#) that is a SQL [SELECT](#) statement or stored procedure to be used by the [SelectCommand](#) property of the [MySqlDataAdapter](#).

**Parameters:** The connection string

When an instance of [MySqlDataAdapter](#) is created, the following read/write properties are set to the following initial values.

Properties	Initial Value
<a href="#">MissingMappingAction</a>	<a href="#">MissingMappingAction.Passthrough</a>
<a href="#">MissingSchemaAction</a>	<a href="#">MissingSchemaAction.Add</a>

You can change the value of any of these properties through a separate call to the property.

### Examples

The following example creates a [MySqlDataAdapter](#) and sets some of its properties.

Visual Basic example:

```
Public Sub CreateSqlDataAdapter()
    Dim da As MySqlDataAdapter = New MySqlDataAdapter("SELECT id, name FROM mytable", "Data Source=localhost;")
    Dim conn As MySqlConnection = da.SelectCommand.Connection
    da.MissingSchemaAction = MissingSchemaAction.AddWithKey

    da.InsertCommand = New MySqlCommand("INSERT INTO mytable (id, name) " & _
        "VALUES (?id, ?name)", conn)
    da.UpdateCommand = New MySqlCommand("UPDATE mytable SET id=?id, name=?name " & _
        "WHERE id=?oldId", conn)
    da.DeleteCommand = New MySqlCommand("DELETE FROM mytable WHERE id=?id", conn)
    da.InsertCommand.Parameters.Add("?id", MySqlDbType.VarChar, 5, "id")
    da.InsertCommand.Parameters.Add("?name", MySqlDbType.VarChar, 40, "name")

    da.UpdateCommand.Parameters.Add("?id", MySqlDbType.VarChar, 5, "id")
    da.UpdateCommand.Parameters.Add("?name", MySqlDbType.VarChar, 40, "name")
    da.UpdateCommand.Parameters.Add("?oldId", MySqlDbType.VarChar, 5, "id").SourceVersion = DataRowVersion.Current
    da.DeleteCommand.Parameters.Add("?id", MySqlDbType.VarChar, 5, "id").SourceVersion = DataRowVersion.Current
End Sub
```

C# example:

```
public static void CreateSqlDataAdapter()
{
    MySqlDataAdapter da = new MySqlDataAdapter("SELECT id, name FROM mytable", "Data Source=localhost;");
    MySqlConnection conn = da.SelectCommand.Connection;
    da.MissingSchemaAction = MissingSchemaAction.AddWithKey;

    da.InsertCommand = new MySqlCommand("INSERT INTO mytable (id, name) " +
        "VALUES (?id, ?name)", conn);
    da.UpdateCommand = new MySqlCommand("UPDATE mytable SET id=?id, name=?name " +
        "WHERE id=?oldId", conn);
    da.DeleteCommand = new MySqlCommand("DELETE FROM mytable WHERE id=?id", conn);
    da.InsertCommand.Parameters.Add("?id", MySqlDbType.VarChar, 5, "id");
    da.InsertCommand.Parameters.Add("?name", MySqlDbType.VarChar, 40, "name");

    da.UpdateCommand.Parameters.Add("?id", MySqlDbType.VarChar, 5, "id");
    da.UpdateCommand.Parameters.Add("?name", MySqlDbType.VarChar, 40, "name");
    da.UpdateCommand.Parameters.Add("?oldId", MySqlDbType.VarChar, 5, "id").SourceVersion = DataRowVersion.Current;
    da.DeleteCommand.Parameters.Add("?id", MySqlDbType.VarChar, 5, "id").SourceVersion = DataRowVersion.Current;
}
```

#### 18.2.3.4.5. DeleteCommand

Gets or sets a SQL statement or stored procedure used to delete records from the data set.

**Value:** A `MySqlCommand` used during `System.Data.Common.DataAdapter.Update` to delete records in the database that correspond to deleted rows in the `DataSet`.

During `System.Data.Common.DataAdapter.Update`, if this property is not set and primary key information is present in the `DataSet`, the `DeleteCommand` can be generated automatically if you set the `SelectCommand` property and use the `MySqlCommandBuilder`. Then, any additional commands that you do not set are generated by the `MySqlCommandBuilder`. This generation logic requires key column information to be present in the `DataSet`.

When `DeleteCommand` is assigned to a previously created `MySqlCommand`, the `MySqlCommand` is not cloned. The `DeleteCommand` maintains a reference to the previously created `MySqlCommand` object.

### Examples

The following example creates a `MySqlDataAdapter` and sets the `SelectCommand` and `DeleteCommand` properties. It assumes you have already created a `MySqlConnection` object.

Visual Basic example:

```
Public Shared Function CreateCustomerAdapter(conn As MySqlConnection) As MySqlDataAdapter
    Dim da As MySqlDataAdapter = New MySqlDataAdapter()
    Dim cmd As MySqlCommand
    Dim parm As MySqlParameter
    ' Create the SelectCommand.
    cmd = New MySqlCommand("SELECT * FROM mytable WHERE id=?id AND name=?name", conn)
    cmd.Parameters.Add("?id", MySqlDbType.VarChar, 15)
    cmd.Parameters.Add("?name", MySqlDbType.VarChar, 15)
    da.SelectCommand = cmd
    ' Create the DeleteCommand.
    cmd = New MySqlCommand("DELETE FROM mytable WHERE id=?id", conn)
    parm = cmd.Parameters.Add("?id", MySqlDbType.VarChar, 5, "id")
    parm.SourceVersion = DataRowVersion.Original
    da.DeleteCommand = cmd
    Return da
End Function
```

C# example:

```
public static MySqlDataAdapter CreateCustomerAdapter(MySqlConnection conn)
{
    MySqlDataAdapter da = new MySqlDataAdapter();
    MySqlCommand cmd;
    MySqlParameter parm;
    // Create the SelectCommand.
    cmd = new MySqlCommand("SELECT * FROM mytable WHERE id=?id AND name=?name", conn);
    cmd.Parameters.Add("?id", MySqlDbType.VarChar, 15);
    cmd.Parameters.Add("?name", MySqlDbType.VarChar, 15);
    da.SelectCommand = cmd;
    // Create the DeleteCommand.
    cmd = new MySqlCommand("DELETE FROM mytable WHERE id=?id", conn);
    parm = cmd.Parameters.Add("?id", MySqlDbType.VarChar, 5, "id");
    parm.SourceVersion = DataRowVersion.Original;
    da.DeleteCommand = cmd;
    return da;
}
```

### 18.2.3.4.6. InsertCommand

Gets or sets a SQL statement or stored procedure used to insert records into the data set.

**Value:** A `MySqlCommand` used during `System.Data.Common.DataAdapter.Update` to insert records into the database that correspond to new rows in the `DataSet`.

During `System.Data.Common.DataAdapter.Update`, if this property is not set and primary

key information is present in the [DataSet](#), the [InsertCommand](#) can be generated automatically if you set the [SelectCommand](#) property and use the [MySQLCommandBuilder](#). Then, any additional commands that you do not set are generated by the [MySQLCommandBuilder](#). This generation logic requires key column information to be present in the [DataSet](#).

When [InsertCommand](#) is assigned to a previously created [MySQLCommand](#), the [MySQLCommand](#) is not cloned. The [InsertCommand](#) maintains a reference to the previously created [MySQLCommand](#) object.

**Note.** If execution of this command returns rows, these rows may be added to the [DataSet](#) depending on how you set the [MySQLCommand.UpdatedRowSource](#) property of the [MySQLCommand](#) object.

### Examples

The following example creates a [MySQLDataAdapter](#) and sets the [SelectCommand](#) and [InsertCommand](#) properties. It assumes you have already created a [MySQLConnection](#) object.

Visual Basic example:

```
Public Shared Function CreateCustomerAdapter(conn As MySqlConnection) As MySqlDataAdapter
    Dim da As MySqlDataAdapter = New MySqlDataAdapter()
    Dim cmd As MySqlCommand
    Dim parm As MySqlParameter
    ' Create the SelectCommand.
    cmd = New MySqlCommand("SELECT * FROM mytable WHERE id=?id AND name=?name", conn)
    cmd.Parameters.Add("?id", MySqlDbType.VarChar, 15)
    cmd.Parameters.Add("?name", MySqlDbType.VarChar, 15)
    da.SelectCommand = cmd
    ' Create the InsertCommand.
    cmd = New MySqlCommand("INSERT INTO mytable (id,name) VALUES (?id, ?name)", conn)
    cmd.Parameters.Add( "?id", MySqlDbType.VarChar, 15, "id" )
    cmd.Parameters.Add( "?name", MySqlDbType.VarChar, 15, "name" )
    da.InsertCommand = cmd

    Return da
End Function
```

C# example:

```
public static MySqlDataAdapter CreateCustomerAdapter(MySqlConnection conn)
{
    MySqlDataAdapter da = new MySqlDataAdapter();
    MySqlCommand cmd;
    MySqlParameter parm;
    // Create the SelectCommand.
    cmd = new MySqlCommand("SELECT * FROM mytable WHERE id=?id AND name=?name", conn);
    cmd.Parameters.Add("?id", MySqlDbType.VarChar, 15);
    cmd.Parameters.Add("?name", MySqlDbType.VarChar, 15);
    da.SelectCommand = cmd;
    // Create the InsertCommand.
    cmd = new MySqlCommand("INSERT INTO mytable (id,name) VALUES (?id,?name)", conn);
    cmd.Parameters.Add("?id", MySqlDbType.VarChar, 15, "id" );
    cmd.Parameters.Add("?name", MySqlDbType.VarChar, 15, "name" );

    da.InsertCommand = cmd;
    return da;
}
```

#### 18.2.3.4.7. UpdateCommand

Gets or sets a SQL statement or stored procedure used to updated records in the data source.

**Value:** A [MySQLCommand](#) used during [System.Data.Common.DataAdapter.Update](#) to update records in the database with data from the [DataSet](#).

During [System.Data.Common.DataAdapter.Update](#), if this property is not set and primary key information is present in the [DataSet](#), the [UpdateCommand](#) can be generated automatically if



you set the `SelectCommand` property and use the `MySQLCommandBuilder`. Then, any additional commands that you do not set are generated by the `MySQLCommandBuilder`. This generation logic requires key column information to be present in the `DataSet`.

When `UpdateCommand` is assigned to a previously created `MySQLCommand`, the `MySQLCommand` is not cloned. The `UpdateCommand` maintains a reference to the previously created `MySQLCommand` object.

**Note.** If execution of this command returns rows, these rows may be merged with the `DataSet` depending on how you set the `MySQLCommand.UpdatedRowSource` property of the `MySQLCommand` object.

### Examples

The following example creates a `MySQLDataAdapter` and sets the `SelectCommand` and `UpdateCommand` properties. It assumes you have already created a `MySQLConnection` object.

Visual Basic example:

```
Public Shared Function CreateCustomerAdapter(conn As MySqlConnection) As MySQLDataAdapter
    Dim da As MySQLDataAdapter = New MySQLDataAdapter()
    Dim cmd As MySQLCommand
    Dim parm As MySQLParameter
    ' Create the SelectCommand.
    cmd = New MySQLCommand("SELECT * FROM mytable WHERE id=?id AND name=?name", conn)
    cmd.Parameters.Add("?id", MySQLDbType.VarChar, 15)
    cmd.Parameters.Add("?name", MySQLDbType.VarChar, 15)
    da.SelectCommand = cmd
    ' Create the UpdateCommand.
    cmd = New MySQLCommand("UPDATE mytable SET id=?id, name=?name WHERE id=?oldId", conn)
    cmd.Parameters.Add("?id", MySQLDbType.VarChar, 15, "id" )
    cmd.Parameters.Add("?name", MySQLDbType.VarChar, 15, "name" )

    parm = cmd.Parameters.Add("?oldId", MySQLDbType.VarChar, 15, "id")
    parm.SourceVersion = DataRowVersion.Original

    da.UpdateCommand = cmd

    Return da
End Function
```

C# example:

```
public static MySQLDataAdapter CreateCustomerAdapter(MySqlConnection conn)
{
    MySQLDataAdapter da = new MySQLDataAdapter();
    MySQLCommand cmd;
    MySQLParameter parm;
    // Create the SelectCommand.
    cmd = new MySQLCommand("SELECT * FROM mytable WHERE id=?id AND name=?name", conn);
    cmd.Parameters.Add("?id", MySQLDbType.VarChar, 15);
    cmd.Parameters.Add("?name", MySQLDbType.VarChar, 15);
    da.SelectCommand = cmd;
    // Create the UpdateCommand.
    cmd = new MySQLCommand("UPDATE mytable SET id=?id, name=?name WHERE id=?oldId", conn);
    cmd.Parameters.Add("?id", MySQLDbType.VarChar, 15, "id" );
    cmd.Parameters.Add("?name", MySQLDbType.VarChar, 15, "name" );

    parm = cmd.Parameters.Add( "?oldId", MySQLDbType.VarChar, 15, "id" );
    parm.SourceVersion = DataRowVersion.Original;

    da.UpdateCommand = cmd;
    return da;
}
```

#### 18.2.3.4.8. SelectCommand

Gets or sets a SQL statement or stored procedure used to select records in the data source.

**Value:** A `MySQLCommand` used during `System.Data.Common.DbDataAdapter.Fill` to select records from the database for placement in the `DataSet`.

When `SelectCommand` is assigned to a previously created `MySQLCommand`, the `MySQLCommand` is not cloned. The `SelectCommand` maintains a reference to the previously created `MySQLCommand` object.

If the `SelectCommand` does not return any rows, no tables are added to the `DataSet`, and no exception is raised.

### Examples

The following example creates a `MySQLDataAdapter` and sets the `SelectCommand` and `InsertCommand` properties. It assumes you have already created a `MySQLConnection` object.

Visual Basic example:

```
Public Shared Function CreateCustomerAdapter(conn As MySQLConnection) As MySQLDataAdapter
    Dim da As MySQLDataAdapter = New MySQLDataAdapter()
    Dim cmd As MySQLCommand
    Dim parm As MySQLParameter
    ' Create the SelectCommand.
    cmd = New MySQLCommand("SELECT * FROM mytable WHERE id=?id AND name=?name", conn)
    cmd.Parameters.Add("?id", MySQLDbType.VarChar, 15)
    cmd.Parameters.Add("?name", MySQLDbType.VarChar, 15)
    da.SelectCommand = cmd
    ' Create the InsertCommand.
    cmd = New MySQLCommand("INSERT INTO mytable (id,name) VALUES (?id, ?name)", conn)
    cmd.Parameters.Add( "?id", MySQLDbType.VarChar, 15, "id" )
    cmd.Parameters.Add( "?name", MySQLDbType.VarChar, 15, "name" )
    da.InsertCommand = cmd

    Return da
End Function
```

C# example:

```
public static MySQLDataAdapter CreateCustomerAdapter(MySQLConnection conn)
{
    MySQLDataAdapter da = new MySQLDataAdapter();
    MySQLCommand cmd;
    MySQLParameter parm;
    // Create the SelectCommand.
    cmd = new MySQLCommand("SELECT * FROM mytable WHERE id=?id AND name=?name", conn);
    cmd.Parameters.Add("?id", MySQLDbType.VarChar, 15);
    cmd.Parameters.Add("?name", MySQLDbType.VarChar, 15);
    da.SelectCommand = cmd;
    // Create the InsertCommand.
    cmd = new MySQLCommand("INSERT INTO mytable (id,name) VALUES (?id,?name)", conn);
    cmd.Parameters.Add("?id", MySQLDbType.VarChar, 15, "id" );
    cmd.Parameters.Add("?name", MySQLDbType.VarChar, 15, "name" );

    da.InsertCommand = cmd;
    return da;
}
```

## 18.2.3.5. MySQLDataReader

To create a `MySQLDataReader`, you must call the `MySQLCommand.ExecuteReader` method of the `MySQLCommand` object, rather than directly using a constructor.

While the `MySQLDataReader` is in use, the associated `MySQLConnection` is busy serving the `MySQLDataReader`, and no other operations can be performed on the `MySQLConnection` other than closing it. This is the case until the `MySQLDataReader.Close` method of the `MySQLDataReader` is called.

`MySqlDataReader.IsClosed` and `MySqlDataReader.RecordsAffected` are the only properties that you can call after the `MySqlDataReader` is closed. Though the `RecordsAffected` property may be accessed at any time while the `MySqlDataReader` exists, always call `Close` before returning the value of `RecordsAffected` to ensure an accurate return value.

For optimal performance, `MySqlDataReader` avoids creating unnecessary objects or making unnecessary copies of data. As a result, multiple calls to methods such as `MySqlDataReader.GetValue` return a reference to the same object. Use caution if you are modifying the underlying value of the objects returned by methods such as `GetValue`.

### Examples

The following example creates a `MySqlConnection`, a `MySqlCommand`, and a `MySqlDataReader`. The example reads through the data, writing it out to the console. Finally, the example closes the `MySqlDataReader`, then the `MySqlConnection`.

Visual Basic example:

```
Public Sub ReadMyData(myConnString As String)
    Dim mySelectQuery As String = "SELECT OrderID, CustomerID FROM Orders"
    Dim myConnection As New MySqlConnection(myConnString)
    Dim myCommand As New MySqlCommand(mySelectQuery, myConnection)
    myConnection.Open()
    Dim myReader As MySqlDataReader
    myReader = myCommand.ExecuteReader()
    ' Always call Read before accessing data.
    While myReader.Read()
        Console.WriteLine((myReader.GetInt32(0) & ", " & myReader.GetString(1)))
    End While
    ' always call Close when done reading.
    myReader.Close()
    ' Close the connection when done with it.
    myConnection.Close()
End Sub 'ReadMyData
```

C# example:

```
public void ReadMyData(string myConnString) {
    string mySelectQuery = "SELECT OrderID, CustomerID FROM Orders";
    MySqlConnection myConnection = new MySqlConnection(myConnString);
    MySqlCommand myCommand = new MySqlCommand(mySelectQuery, myConnection);
    myConnection.Open();
    MySqlDataReader myReader;
    myReader = myCommand.ExecuteReader();
    // Always call Read before accessing data.
    while (myReader.Read()) {
        Console.WriteLine(myReader.GetInt32(0) + ", " + myReader.GetString(1));
    }
    // always call Close when done reading.
    myReader.Close();
    // Close the connection when done with it.
    myConnection.Close();
}
```

#### 18.2.3.5.1. GetBytes

`GetBytes` returns the number of available bytes in the field. In most cases this is the exact length of the field. However, the number returned may be less than the true length of the field if `GetBytes` has already been used to obtain bytes from the field. This may be the case, for example, if the `MySqlDataReader` is reading a large data structure into a buffer. For more information, see the `SequentialAccess` setting for `MySqlCommand.CommandBehavior`.

If you pass a buffer that is a null reference (`Nothing` in Visual Basic), `GetBytes` returns the length of the field in bytes.

No conversions are performed; therefore the data retrieved must already be a byte array.

#### 18.2.3.5.2. **GetTimeSpan**

Gets the value of the specified column as a [TimeSpan](#) object.

**Parameters:** The zero-based column ordinal.

**Returns:** The value of the specified column.

#### 18.2.3.5.3. **GetDateTime**

Gets the value of the specified column as a [DateTime](#) object.

**Note.** MySQL allows date columns to contain the value '0000-00-00' and datetime columns to contain the value '0000-00-00 00:00:00'. The [DateTime](#) structure cannot contain or represent these values. To read a datetime value from a column that might contain zero values, use [GetMySQLDateTime](#). The behavior of reading a zero datetime column using this method is defined by the [ZeroDateTimeBehavior](#) connection string option. For more information on this option, please refer to [MySQLConnection.ConnectionString](#).

**Parameters:** The zero-based column ordinal.

**Returns:** The value of the specified column.

#### 18.2.3.5.4. **GetMySQLDateTime**

Gets the value of the specified column as a [MySQL.Data.Types.MySQLDateTime](#) object.

**Parameters:** The zero-based column ordinal.

**Returns:** The value of the specified column.

#### 18.2.3.5.5. **GetString**

Gets the value of the specified column as a [String](#) object.

**Parameters:** The zero-based column ordinal.

**Returns:** The value of the specified column.

#### 18.2.3.5.6. **GetDecimal**

Gets the value of the specified column as a [Decimal](#) object.

**Parameters:** The zero-based column ordinal.

**Returns:** The value of the specified column.

#### 18.2.3.5.7. **GetDouble**

Gets the value of the specified column as a double-precision floating point number.

**Parameters:** The zero-based column ordinal.

**Returns:** The value of the specified column.

#### **18.2.3.5.8. GetFloat**

Gets the value of the specified column as a single-precision floating point number.

**Parameters:** The zero-based column ordinal.

**Returns:** The value of the specified column.

#### **18.2.3.5.9. GetGuid**

Gets the value of the specified column as a globally-unique identifier (GUID).

**Parameters:** The zero-based column ordinal.

**Returns:** The value of the specified column.

#### **18.2.3.5.10. GetInt16**

Gets the value of the specified column as a 16-bit signed integer.

**Parameters:** The zero-based column ordinal.

**Returns:** The value of the specified column.

#### **18.2.3.5.11. GetInt32**

Gets the value of the specified column as a 32-bit signed integer.

**Parameters:** The zero-based column ordinal.

**Returns:** The value of the specified column.

#### **18.2.3.5.12. GetInt64**

Gets the value of the specified column as a 64-bit signed integer.

**Parameters:** The zero-based column ordinal.

**Returns:** The value of the specified column.

#### **18.2.3.5.13. GetUInt16**

Gets the value of the specified column as a 16-bit unsigned integer.

**Parameters:** The zero-based column ordinal.

**Returns:** The value of the specified column.

#### **18.2.3.5.14. GetUInt32**

Gets the value of the specified column as a 32-bit unsigned integer.

**Parameters:** The zero-based column ordinal.

**Returns:** The value of the specified column.

#### **18.2.3.5.15. GetUInt64**

Gets the value of the specified column as a 64-bit unsigned integer.

**Parameters:** The zero-based column ordinal.

**Returns:** The value of the specified column.

### 18.2.3.6. MySqlConnectionException

This class is created whenever the MySQL Data Provider encounters an error generated from the server.

Any open connections are not automatically closed when an exception is thrown. If the client application determines that the exception is fatal, it should close any open [MySQLDataReader](#) objects or [MySqlConnection](#) objects.

#### Examples

The following example generates a [MySqlConnectionException](#) due to a missing server, and then displays the exception.

Visual Basic example:

```
Public Sub ShowException()  
    Dim mySelectQuery As String = "SELECT column1 FROM table1"  
    Dim myConnection As New MySqlConnection ("Data Source=localhost;Database=Sample;")  
    Dim myCommand As New MySqlCommand(mySelectQuery, myConnection)  
    Try  
        myCommand.Connection.Open()  
    Catch e As MySqlConnectionException  
        MessageBox.Show( e.Message )  
    End Try  
End Sub
```

C# example:

```
public void ShowException()  
{  
    string mySelectQuery = "SELECT column1 FROM table1";  
    MySqlConnection myConnection =  
        new MySqlConnection("Data Source=localhost;Database=Sample;");  
    MySqlCommand myCommand = new MySqlCommand(mySelectQuery, myConnection);  
    try  
    {  
        myCommand.Connection.Open();  
    }  
    catch (MySqlConnectionException e)  
    {  
        MessageBox.Show( e.Message );  
    }  
}
```

### 18.2.3.7. MySqlParameter

Parameter names are not case sensitive.

#### Examples

The following example creates multiple instances of [MySqlParameter](#) through the [MySqlParameterCollection](#) collection within the [MySqlDataAdapter](#). These parameters are used to select data from the data source and place the data in the [DataSet](#). This example assumes that a [DataSet](#) and a [MySqlDataAdapter](#) have already been created with the appropriate schema, commands, and connection.

Visual Basic example:

```
Public Sub AddSqlParameters()
    ' ...
    ' create myDataSet and myDataAdapter
    ' ...
    myDataAdapter.SelectCommand.Parameters.Add("@CategoryName", MySqlDbType.VarChar, 80).Value = "toast"
    myDataAdapter.SelectCommand.Parameters.Add("@SerialNum", MySqlDbType.Long).Value = 239

    myDataAdapter.Fill(myDataSet)
End Sub 'AddSqlParameters
```

C# example:

```
public void AddSqlParameters()
{
    // ...
    // create myDataSet and myDataAdapter
    // ...
    myDataAdapter.SelectCommand.Parameters.Add("@CategoryName", MySqlDbType.VarChar, 80).Value = "toaster"
    myDataAdapter.SelectCommand.Parameters.Add("@SerialNum", MySqlDbType.Long).Value = 239;
    myDataAdapter.Fill(myDataSet);
}
```

### 18.2.3.8. MySqlConnectionParameterCollection

The number of the parameters in the collection must be equal to the number of parameter placeholders within the command text, or an exception will be generated.

#### Examples

The following example creates multiple instances of [MySqlConnectionParameter](#) through the [MySqlConnectionParameterCollection](#) collection within the [MySqlConnectionDataAdapter](#). These parameters are used to select data within the data source and place the data in the [DataSet](#). This code assumes that a [DataSet](#) and a [MySqlConnectionDataAdapter](#) have already been created with the appropriate schema, commands, and connection.

Visual Basic example:

```
Public Sub AddParameters()
    ' ...
    ' create myDataSet and myDataAdapter
    ' ...
    myDataAdapter.SelectCommand.Parameters.Add("@CategoryName", MySqlDbType.VarChar, 80).Value = "toast"
    myDataAdapter.SelectCommand.Parameters.Add("@SerialNum", MySqlDbType.Long).Value = 239

    myDataAdapter.Fill(myDataSet)
End Sub 'AddSqlParameters
```

C# example:

```
public void AddSqlParameters()
{
    // ...
    // create myDataSet and myDataAdapter
    // ...
    myDataAdapter.SelectCommand.Parameters.Add("@CategoryName", MySqlDbType.VarChar, 80).Value = "toaster"
    myDataAdapter.SelectCommand.Parameters.Add("@SerialNum", MySqlDbType.Long).Value = 239;
    myDataAdapter.Fill(myDataSet);
}
```

### 18.2.3.9. MySqlConnectionTransaction

Represents a SQL transaction to be made in a MySQL database. This class cannot be inherited.

The application creates a `MySqlTransaction` object by calling `MySqlConnection.BeginTransaction` on the `MySqlConnection` object. All subsequent operations associated with the transaction (for example, committing or aborting the transaction), are performed on the `MySqlTransaction` object.

### Examples

The following example creates a `MySqlConnection` and a `MySqlTransaction`. It also demonstrates how to use the `MySqlConnection.BeginTransaction`, `MySqlTransaction.Commit`, and `MySqlTransaction.Rollback` methods.

Visual Basic example:

```
Public Sub RunTransaction(myConnString As String)
    Dim myConnection As New MySqlConnection(myConnString)
    myConnection.Open()

    Dim myCommand As MySqlCommand = myConnection.CreateCommand()
    Dim myTrans As MySqlTransaction

    ' Start a local transaction
    myTrans = myConnection.BeginTransaction()
    ' Must assign both transaction object and connection
    ' to Command object for a pending local transaction
    myCommand.Connection = myConnection
    myCommand.Transaction = myTrans

    Try
        myCommand.CommandText = "Insert into Region (RegionID, RegionDescription) VALUES (100, 'Description 100)"
        myCommand.ExecuteNonQuery()
        myCommand.CommandText = "Insert into Region (RegionID, RegionDescription) VALUES (101, 'Description 101)"
        myCommand.ExecuteNonQuery()
        myTrans.Commit()
        Console.WriteLine("Both records are written to database.")
    Catch e As Exception
        Try
            myTrans.Rollback()
        Catch ex As MySqlException
            If Not myTrans.Connection Is Nothing Then
                Console.WriteLine("An exception of type " & ex.GetType().ToString() & _
                    " was encountered while attempting to roll back the transaction.")
            End If
        End Try

        Console.WriteLine("An exception of type " & e.GetType().ToString() & _
            " was encountered while inserting the data.")
        Console.WriteLine("Neither record was written to database.")
    Finally
        myConnection.Close()
    End Try
End Sub 'RunTransaction
```

C# example:

```
public void RunTransaction(string myConnString)
{
    MySqlConnection myConnection = new MySqlConnection(myConnString);
    myConnection.Open();
    MySqlCommand myCommand = myConnection.CreateCommand();
    MySqlTransaction myTrans;
    // Start a local transaction
    myTrans = myConnection.BeginTransaction();
    // Must assign both transaction object and connection
    // to Command object for a pending local transaction
    myCommand.Connection = myConnection;
    myCommand.Transaction = myTrans;
    try
    {
        myCommand.CommandText = "Insert into Region (RegionID, RegionDescription) VALUES (100, 'Description 100)"
        myCommand.ExecuteNonQuery();
    }
}
```



```

        myCommand.CommandText = "Insert into Region (RegionID, RegionDescription) VALUES (101, 'Description')";
        myCommand.ExecuteNonQuery();
        myTrans.Commit();
        Console.WriteLine("Both records are written to database.");
    }
    catch (Exception e)
    {
        try
        {
            myTrans.Rollback();
        }
        catch (MySqlException ex)
        {
            if (myTrans.Connection != null)
            {
                Console.WriteLine("An exception of type " + ex.GetType() +
                    " was encountered while attempting to roll back the transaction.");
            }
        }

        Console.WriteLine("An exception of type " + e.GetType() +
            " was encountered while inserting the data.");
        Console.WriteLine("Neither record was written to database.");
    }
    finally
    {
        myConnection.Close();
    }
}

```

### 18.2.3.9.1. Rollback

Rolls back a transaction from a pending state.

The Rollback method is equivalent to the MySQL statement ROLLBACK. The transaction can only be rolled back from a pending state (after BeginTransaction has been called, but before Commit is called).

#### Examples

The following example creates [MySqlConnection](#) and a [MySqlTransaction](#). It also demonstrates how to use the [MySqlConnection.BeginTransaction](#), [Commit](#), and [Rollback](#) methods.

Visual Basic example:

```

Public Sub RunSqlTransaction(myConnString As String)
    Dim myConnection As New MySqlConnection(myConnString)
    myConnection.Open()

    Dim myCommand As MySqlCommand = myConnection.CreateCommand()
    Dim myTrans As MySqlTransaction

    ' Start a local transaction
    myTrans = myConnection.BeginTransaction()

    ' Must assign both transaction object and connection
    ' to Command object for a pending local transaction
    myCommand.Connection = myConnection
    myCommand.Transaction = myTrans

    Try
        myCommand.CommandText = "Insert into mytable (id, desc) VALUES (100, 'Description')";
        myCommand.ExecuteNonQuery()
        myCommand.CommandText = "Insert into mytable (id, desc) VALUES (101, 'Description')";
        myCommand.ExecuteNonQuery()
        myTrans.Commit()
        Console.WriteLine("Success.")
    Catch e As Exception
        Try
            myTrans.Rollback()
        Catch ex As MySqlException
            If Not myTrans.Connection Is Nothing Then
                Console.WriteLine("An exception of type " & ex.GetType().ToString() & _

```

```

        " was encountered while attempting to roll back the transaction.")
    End If
End Try

Console.WriteLine("An exception of type " & e.GetType().ToString() & _
    "was encountered while inserting the data.")
Console.WriteLine("Neither record was written to database.")
Finally
    myConnection.Close()
End Try
End Sub

```

C# example:

```

public void RunSqlTransaction(string myConnString)
{
    MySqlConnection myConnection = new MySqlConnection(myConnString);
    myConnection.Open();
    MySqlCommand myCommand = myConnection.CreateCommand();
    MySqlTransaction myTrans;
    // Start a local transaction
    myTrans = myConnection.BeginTransaction();
    // Must assign both transaction object and connection
    // to Command object for a pending local transaction
    myCommand.Connection = myConnection;
    myCommand.Transaction = myTrans;
    try
    {
        myCommand.CommandText = "Insert into mytable (id, desc) VALUES (100, 'Description')";
        myCommand.ExecuteNonQuery();
        myCommand.CommandText = "Insert into mytable (id, desc) VALUES (101, 'Description')";
        myCommand.ExecuteNonQuery();
        myTrans.Commit();
        Console.WriteLine("Both records are written to database.");
    }
    catch (Exception e)
    {
        try
        {
            myTrans.Rollback();
        }
        catch (MySqlException ex)
        {
            if (myTrans.Connection != null)
            {
                Console.WriteLine("An exception of type " + ex.GetType() +
                    " was encountered while attempting to roll back the transaction.");
            }
        }

        Console.WriteLine("An exception of type " + e.GetType() +
            " was encountered while inserting the data.");
        Console.WriteLine("Neither record was written to database.");
    }
    finally
    {
        myConnection.Close();
    }
}

```

### 18.2.3.9.2. Commit

Commits the database transaction.

The [Commit](#) method is equivalent to the MySQL SQL statement COMMIT.

#### Examples

The following example creates [MySqlConnection](#) and a [MySqlTransaction](#). It also demonstrates how to use the [MySqlConnection.BeginTransaction](#), [Commit](#), and [Rollback](#) methods.

## Visual Basic example:

```

Public Sub RunSqlTransaction(myConnString As String)
    Dim myConnection As New MySqlConnection(myConnString)
    myConnection.Open()

    Dim myCommand As MySqlCommand = myConnection.CreateCommand()
    Dim myTrans As MySqlTransaction

    ' Start a local transaction
    myTrans = myConnection.BeginTransaction()

    ' Must assign both transaction object and connection
    ' to Command object for a pending local transaction
    myCommand.Connection = myConnection
    myCommand.Transaction = myTrans

    Try
        myCommand.CommandText = "Insert into mytable (id, desc) VALUES (100, 'Description')"
        myCommand.ExecuteNonQuery()
        myCommand.CommandText = "Insert into mytable (id, desc) VALUES (101, 'Description')"
        myCommand.ExecuteNonQuery()
        myTrans.Commit()
        Console.WriteLine("Success.")
    Catch e As Exception
        Try
            myTrans.Rollback()
        Catch ex As MySqlException
            If Not myTrans.Connection Is Nothing Then
                Console.WriteLine("An exception of type " & ex.GetType().ToString() & _
                    " was encountered while attempting to roll back the transaction.")
            End If
        End Try

        Console.WriteLine("An exception of type " & e.GetType().ToString() & _
            " was encountered while inserting the data.")
        Console.WriteLine("Neither record was written to database.")
    Finally
        myConnection.Close()
    End Try
End Sub

```

## C# example:

```

public void RunSqlTransaction(string myConnString)
{
    MySqlConnection myConnection = new MySqlConnection(myConnString);
    myConnection.Open();
    MySqlCommand myCommand = myConnection.CreateCommand();
    MySqlTransaction myTrans;
    // Start a local transaction
    myTrans = myConnection.BeginTransaction();
    // Must assign both transaction object and connection
    // to Command object for a pending local transaction
    myCommand.Connection = myConnection;
    myCommand.Transaction = myTrans;
    try
    {
        myCommand.CommandText = "Insert into mytable (id, desc) VALUES (100, 'Description')";
        myCommand.ExecuteNonQuery();
        myCommand.CommandText = "Insert into mytable (id, desc) VALUES (101, 'Description')";
        myCommand.ExecuteNonQuery();
        myTrans.Commit();
        Console.WriteLine("Both records are written to database.");
    }
    catch(Exception e)
    {
        try
        {
            myTrans.Rollback();
        }
        catch (MySqlException ex)
        {
            if (myTrans.Connection != null)
            {
                Console.WriteLine("An exception of type " + ex.GetType() +
                    " was encountered while attempting to roll back the transaction.");
            }
        }
    }
}

```

```

    }
    Console.WriteLine("An exception of type " + e.GetType() +
        " was encountered while inserting the data.");
    Console.WriteLine("Neither record was written to database.");
}
finally
{
    myConnection.Close();
}
}

```

## 18.2.4. Connector/NET Reference

This section of the manual contains a complete reference to the Connector/NET ADO.NET component, automatically generated from the embedded documentation.

### 18.2.4.1. MySql.Data.MySqlClient

[Namespace hierarchy](#)

#### Classes

Class	Description
<a href="#">MySqlCommand</a>	
<a href="#">MySqlCommandBuilder</a>	
<a href="#">MySqlConnection</a>	
<a href="#">MySqlDataAdapter</a>	
<a href="#">MySqlDataReader</a>	Provides a means of reading a forward-only stream of rows from a MySQL database. This class cannot be inherited.
<a href="#">MySqlError</a>	Collection of error codes that can be returned by the server
<a href="#">MySqlException</a>	The exception that is thrown when MySQL returns an error. This class cannot be inherited.
<a href="#">MySqlHelper</a>	Helper class that makes it easier to work with the provider.
<a href="#">MySqlInfoMessageEventArgs</a>	Provides data for the InfoMessage event. This class cannot be inherited.
<a href="#">MySqlParameter</a>	Represents a parameter to a <a href="#">MySqlCommand</a> , and optionally, its mapping to DataSet columns. This class cannot be inherited.
<a href="#">MySqlParameterCollection</a>	Represents a collection of parameters relevant to a <a href="#">MySqlCommand</a> as well as their respective mappings to columns in a DataSet. This class cannot be inherited.
<a href="#">MySqlRowUpdatedEventArgs</a>	Provides data for the RowUpdated event. This class cannot be inherited.
<a href="#">MySqlRowUpdatingEventArgs</a>	Provides data for the RowUpdating event. This class cannot be inherited.
<a href="#">MySqlTransaction</a>	

#### Delegates

Delegate	Description
<a href="#">MySqlInfoMessageEventHandler</a>	Represents the method that will handle the <a href="#">InfoMessage</a> event of a <a href="#">MySqlConnection</a> .
<a href="#">MySqlRowUpdatedEventHandler</a>	Represents the method that will handle the RowUpdatedevent of a <a href="#">MySqlDataAdapter</a> .
<a href="#">MySqlRowUpdatingEventHandler</a>	Represents the method that will handle the RowUpdatingevent of a <a href="#">MySqlDataAdapter</a> .

### Enumerations

Enumeration	Description
<a href="#">MySqlDbType</a>	Specifies MySQL specific data type of a field, property, for use in a <a href="#">MySqlParameter</a> .
<a href="#">MySqlErrorCode</a>	

#### 18.2.4.1.1. MySql.Data.MySqlCommandHierarchy

See Also

[MySql.Data.MySqlCommand Namespace](#)

#### 18.2.4.1.2. MySqlCommand Class

For a list of all members of this type, see [MySqlCommand Members](#) .

##### Syntax: Visual Basic

```
NotInheritable Public Class MySqlCommand_
    Inherits Component_
    Implements IDbCommand, ICloneable
```

##### Syntax: C#

```
public sealed class MySqlCommand : Component, IDbCommand, ICloneable
```

##### Thread Safety

Public static (Shared in Visual Basic) members of this type are safe for multithreaded operations. Instance members are not guaranteed to be thread-safe.

##### Requirements

Namespace: [MySql.Data.MySqlCommand](#)

Assembly: MySql.Data (in MySql.Data.dll)

See Also

[MySqlCommand Members](#) , [MySql.Data.MySqlCommand Namespace](#)

#### 18.2.4.1.2.1. MySqlCommand Members

[MySqlCommand overview](#)

**Public Instance Constructors**

<a href="#"> MySqlCommand </a>	Overloaded. Initializes a new instance of the MySqlCommand class.
--------------------------------	---

**Public Instance Properties**

<a href="#"> CommandText </a>	
<a href="#"> CommandTimeout </a>	
<a href="#"> CommandType </a>	
<a href="#"> Connection </a>	
Container(inherited from Component)	Gets the IContainerthat contains the Component.
<a href="#"> IsPrepared </a>	
<a href="#"> Parameters </a>	
Site(inherited from Component)	Gets or sets the ISiteof the Component.
<a href="#"> Transaction </a>	
<a href="#"> UpdatedRowSource </a>	

**Public Instance Methods**

<a href="#"> Cancel </a>	Attempts to cancel the execution of a MySqlCommand. This operation is not supported.
CreateObjRef(inherited from MarshalByRefObject)	Creates an object that contains all the relevant information required to generate a proxy used to communicate with a remote object.
<a href="#"> CreateParameter </a>	Creates a new instance of a <a href="#"> MySqlCommandParameter </a> object.
Dispose(inherited from Component)	Releases all resources used by the Component.
Equals(inherited from Object)	Determines whether the specified Objectis equal to the current Object.
<a href="#"> ExecuteNonQuery </a>	
<a href="#"> ExecuteReader </a>	Overloaded.
<a href="#"> ExecuteScalar </a>	
GetHashCode(inherited from Object)	Serves as a hash function for a particular type. GetHashCodeis suitable for use in hashing algorithms and data structures like a hash table.
GetLifetimeService(inherited from MarshalByRefObject)	Retrieves the current lifetime service object that controls the lifetime policy for this instance.
GetType(inherited from Object)	Gets the Typeof the current instance.
InitializeLifetimeService(inherited from MarshalByRefObject)	Obtains a lifetime service object to control the lifetime policy for this instance.
<a href="#"> Prepare </a>	
ToString(inherited from Component)	Returns a Stringcontaining the name of the Component, if any. This method should not be overridden.

**Public Instance Events**

Disposed(inherited from Component)	Adds an event handler to listen to the Disposed event on the component.
------------------------------------	---

**See Also**

[MySqlCommand Class](#) ,  [MySql.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.1. MySqlCommand Constructor

Initializes a new instance of the  [MySqlCommand](#) class.

**Overload List**

Initializes a new instance of the  [MySqlCommand](#) class.

- [public MySqlCommand\(\);](#)
- [public MySqlCommand\(string\);](#)
- [public MySqlCommand\(string,MySqlConnection\);](#)
- [public MySqlCommand\(string,MySqlConnection,MySqlTransaction\);](#)

**See Also**

[MySqlCommand Class](#) ,  [MySql.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.1.1. MySqlCommand Constructor ()

Initializes a new instance of the  [MySqlCommand](#) class.

**Syntax: Visual Basic**

```
Overloads Public Sub New()
```

**Syntax: C#**

```
public MySqlCommand();
```

**See Also**

[MySqlCommand Class](#) ,  [MySql.Data.MySqlClient Namespace](#) ,  [MySqlCommand Constructor Overload List](#)

## 18.2.4.1.2.1.1.2. MySqlCommand Constructor (String)

**Syntax: Visual Basic**

```
Overloads Public Sub New( _  
    ByVal cmdText As String _  
)
```

**Syntax: C#**

```
public MySqlCommand(
    string cmdText
);
```

**See Also**

[MySqlCommand Class](#) , [MySQL.Data.MySqlClient Namespace](#) , [MySqlCommand Constructor Overload List](#)

## 18.2.4.1.2.1.1.3. MySqlCommand Constructor (String, MySqlConnection)

**Syntax: Visual Basic**

```
Overloads Public Sub New( _
    ByVal cmdText As String, _
    ByVal connection As MySqlConnection _
)
```

**Syntax: C#**

```
public MySqlCommand(
    string cmdText,
    MySqlConnection connection
);
```

**See Also**

[MySqlCommand Class](#) , [MySQL.Data.MySqlClient Namespace](#) , [MySqlCommand Constructor Overload List](#)

## 18.2.4.1.2.1.1.3.1. MySqlConnection Class

For a list of all members of this type, see [MySqlConnection Members](#) .

**Syntax: Visual Basic**

```
NotInheritable Public Class MySqlConnection_
    Inherits Component_
    Implements IDbConnection, ICloneable
```

**Syntax: C#**

```
public sealed class MySqlConnection : Component, IDbConnection, ICloneable
```

**Thread Safety**

Public static (Shared in Visual Basic) members of this type are safe for multithreaded operations. Instance members are not guaranteed to be thread-safe.

**Requirements**

Namespace: [MySQL.Data.MySqlClient](#)

Assembly: MySQL.Data (in MySQL.Data.dll)

**See Also**

[MySqlConnection Members](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.1.3.1.1. MySqlConnection Members



[MySQLConnection overview](#)**Public Instance Constructors**

<a href="#">MySQLConnection</a>	Overloaded. Initializes a new instance of the MySQLConnection class.
---------------------------------	--

**Public Instance Properties**

<a href="#">ConnectionString</a>	
<a href="#">ConnectionTimeout</a>	
Container(inherited from Component)	Gets the IContainerthat contains the Component.
<a href="#">Database</a>	
<a href="#">DataSource</a>	Gets the name of the MySQL server to which to connect.
<a href="#">ServerThread</a>	Returns the id of the server thread this connection is executing on
<a href="#">ServerVersion</a>	
Site(inherited from Component)	Gets or sets the ISiteof the Component.
<a href="#">State</a>	
<a href="#">UseCompression</a>	Indicates if this connection should use compression when communicating with the server.

**Public Instance Methods**

<a href="#">BeginTransaction</a>	Overloaded.
<a href="#">ChangeDatabase</a>	
<a href="#">Close</a>	
<a href="#">CreateCommand</a>	
CreateObjRef(inherited from MarshalByRefObject)	Creates an object that contains all the relevant information required to generate a proxy used to communicate with a remote object.
Dispose(inherited from Component)	Releases all resources used by the Component.
Equals(inherited from Object)	Determines whether the specified Objectis equal to the current Object.
GetHashCode(inherited from Object)	Serves as a hash function for a particular type. GetHashCodeis suitable for use in hashing algorithms and data structures like a hash table.
GetLifetimeService(inherited from MarshalByRefObject)	Retrieves the current lifetime service object that controls the lifetime policy for this instance.
GetType(inherited from Object)	Gets the Typeof the current instance.
InitializeLifetimeService(inherited from MarshalByRefObject)	Obtains a lifetime service object to control the lifetime policy for this instance.
<a href="#">Open</a>	
<a href="#">Ping</a>	Ping
ToString(inherited from Component)	Returns a Stringcontaining the name of the Component, if any. This method should not be overrid-

	den.
--	------

**Public Instance Events**

Disposed(inherited from Component)	Adds an event handler to listen to the Dispose- devent on the component.
<a href="#">InfoMessage</a>	
<a href="#">StateChange</a>	

**See Also**

[MySQLConnection Class](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.1.3.1.1.1. MySQLConnection Constructor

Initializes a new instance of the [MySQLConnection](#) class.

**Overload List**

Initializes a new instance of the [MySQLConnection](#) class.

- [public MySQLConnection\(\);](#)
- [public MySQLConnection\(string\);](#)

**See Also**

[MySQLConnection Class](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.1.3.1.1.1.1. MySQLConnection Constructor ()

Initializes a new instance of the [MySQLConnection](#) class.

**Syntax: Visual Basic**

```
Overloads Public Sub New()
```

**Syntax: C#**

```
public MySQLConnection();
```

**See Also**

[MySQLConnection Class](#) , [MySQL.Data.MySqlClient Namespace](#) , [MySQLConnection Constructor Overload List](#)

## 18.2.4.1.2.1.1.3.1.1.1.2. MySQLConnection Constructor (String)

**Syntax: Visual Basic**

```
Overloads Public Sub New( _  
    ByVal connectionString As String _  
)
```

**Syntax: C#**

```
public MySqlConnection(  
    string connectionString  
);
```

**See Also**

[MySqlConnection Class](#) , [MySql.Data.MySqlClient Namespace](#) , [MySqlConnection Constructor Overload List](#)

## 18.2.4.1.2.1.1.3.1.1.2. ConnectionString Property

**Syntax: Visual Basic**

```
NotOverridable Public Property ConnectionString As String _  
    Implements IDbConnection.ConnectionString
```

**Syntax: C#**

```
public string ConnectionString {get; set;}
```

**Implements**

IDbConnection.ConnectionString

**See Also**

[MySqlConnection Class](#) , [MySql.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.1.3.1.1.3. ConnectionTimeout Property

**Syntax: Visual Basic**

```
NotOverridable Public ReadOnly Property ConnectionTimeout As Integer _  
    Implements IDbConnection.ConnectionTimeout
```

**Syntax: C#**

```
public int ConnectionTimeout {get;}
```

**Implements**

IDbConnection.ConnectionTimeout

**See Also**

[MySqlConnection Class](#) , [MySql.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.1.3.1.1.4. Database Property

**Syntax: Visual Basic**

```
NotOverridable Public ReadOnly Property Database As String _  
    Implements IDbConnection.Database
```

**Syntax: C#**

```
public string Database {get;}
```

**Implements**

IDbConnection.Database

**See Also**

[MySQLConnection Class](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.1.3.1.1.5. DataSource Property

Gets the name of the MySQL server to which to connect.

**Syntax: Visual Basic**

```
Public ReadOnly Property DataSource As String
```

**Syntax: C#**

```
public string DataSource {get;}
```

**See Also**

[MySQLConnection Class](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.1.3.1.1.6. ServerThread Property

Returns the id of the server thread this connection is executing on

**Syntax: Visual Basic**

```
Public ReadOnly Property ServerThread As Integer
```

**Syntax: C#**

```
public int ServerThread {get;}
```

**See Also**

[MySQLConnection Class](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.1.3.1.1.7. ServerVersion Property

**Syntax: Visual Basic**

```
Public ReadOnly Property ServerVersion As String
```

**Syntax: C#**

```
public string ServerVersion {get;}
```

**See Also**

[MySQLConnection Class](#) , [MySQL.Data.MySqlClient Namespace](#)

#### 18.2.4.1.2.1.1.3.1.1.8. State Property

**Syntax: Visual Basic**

```
NotOverridable Public ReadOnly Property State As ConnectionState _  
    Implements IDbConnection.State
```

**Syntax: C#**

```
public System.Data.ConnectionState State {get;}
```

**Implements**

IDbConnection.State

**See Also**

[MySQLConnection Class](#) , [MySQL.Data.MySqlClient Namespace](#)

#### 18.2.4.1.2.1.1.3.1.1.9. UseCompression Property

Indicates if this connection should use compression when communicating with the server.

**Syntax: Visual Basic**

```
Public ReadOnly Property UseCompression As Boolean
```

**Syntax: C#**

```
public bool UseCompression {get;}
```

**See Also**

[MySQLConnection Class](#) , [MySQL.Data.MySqlClient Namespace](#)

#### 18.2.4.1.2.1.1.3.1.1.10. BeginTransaction Method

**Overload List**

- [public MySQLTransaction BeginTransaction\(\);](#)
- [public MySQLTransaction BeginTransaction\(IsolationLevel\);](#)

**See Also**

[MySQLConnection Class](#) , [MySQL.Data.MySqlClient Namespace](#)

#### 18.2.4.1.2.1.1.3.1.1.10.1. MySQLConnection.BeginTransaction Method ()

**Syntax: Visual Basic**

```
Overloads Public Function BeginTransaction() As MySQLTransaction
```

**Syntax: C#**

```
public MySQLTransaction BeginTransaction();
```

**See Also**

[MySQLConnection Class](#) , [MySQL.Data.MySqlClient Namespace](#) , [MySQLConnection.BeginTransaction Overload List](#)

## 18.2.4.1.2.1.1.3.1.1.10.1.1. MySQLTransaction Class

For a list of all members of this type, see [MySQLTransaction Members](#) .

**Syntax: Visual Basic**

```
NotInheritable Public Class MySQLTransaction_
    Implements IDbTransaction, IDisposable
```

**Syntax: C#**

```
public sealed class MySQLTransaction : IDbTransaction, IDisposable
```

**Thread Safety**

Public static (Shared in Visual Basic) members of this type are safe for multithreaded operations. Instance members are not guaranteed to be thread-safe.

**Requirements**

Namespace: [MySQL.Data.MySqlClient](#)

Assembly: MySQL.Data (in MySQL.Data.dll)

**See Also**

[MySQLTransaction Members](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.1.3.1.1.10.1.1.1. MySQLTransaction Members

[MySQLTransaction overview](#)

**Public Instance Properties**

<a href="#">Connection</a>	Gets the <a href="#">MySQLConnection</a> object associated with the transaction, or a null reference (Nothing in Visual Basic) if the transaction is no longer valid.
<a href="#">IsolationLevel</a>	Specifies the IsolationLevel for this transaction.

**Public Instance Methods**

<a href="#">Commit</a>	
Equals(inherited from Object)	Determines whether the specified Object is equal to the current Object.
GetHashCode(inherited from Object)	Serves as a hash function for a particular type. GetHashCode is suitable for use in hashing algorithms and data structures like a hash table.
GetType(inherited from Object)	Gets the Type of the current instance.
<a href="#">Rollback</a>	
ToString(inherited from Object)	Returns a String that represents the current Object.

**See Also**

[MySQLTransaction Class](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.1.3.1.1.10.1.1.1.1. Connection Property

Gets the [MySQLConnection](#) object associated with the transaction, or a null reference (Nothing in Visual Basic) if the transaction is no longer valid.

**Syntax: Visual Basic**

```
Public ReadOnly Property Connection As MySqlConnection
```

**Syntax: C#**

```
public MySqlConnection Connection {get;}
```

**Property Value**

The [MySQLConnection](#) object associated with this transaction.

**Remarks**

A single application may have multiple database connections, each with zero or more transactions. This property enables you to determine the connection object associated with a particular transaction created by [BeginTransaction](#) .

**See Also**

[MySQLTransaction Class](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.1.3.1.1.10.1.1.1.2. IsolationLevel Property

Specifies the IsolationLevel for this transaction.

**Syntax: Visual Basic**

```
NotOverridable Public ReadOnly Property IsolationLevel As IsolationLevel _  
- Implements IDbTransaction.IsolationLevel
```

**Syntax: C#**

```
public System.Data.IsolationLevel IsolationLevel {get;}
```

**Property Value**

The IsolationLevel for this transaction. The default is ReadCommitted.

**Implements**

IDbTransaction.IsolationLevel

**Remarks**

Parallel transactions are not supported. Therefore, the IsolationLevel applies to the entire transaction.

**See Also**

[MySQLTransaction Class](#) , [MySQL.Data.MySqlClient Namespace](#)

#### 18.2.4.1.2.1.1.3.1.1.10.1.1.1.3. MySQLTransaction.Commit Method

##### Syntax: Visual Basic

```
NotOverridable Public Sub Commit() _  
- Implements IDbTransaction.Commit
```

##### Syntax: C#

```
public void Commit();
```

##### Implements

IDbTransaction.Commit

##### See Also

[MySQLTransaction Class](#) , [MySQL.Data.MySqlClient Namespace](#)

#### 18.2.4.1.2.1.1.3.1.1.10.1.1.1.4. MySQLTransaction.Rollback Method

##### Syntax: Visual Basic

```
NotOverridable Public Sub Rollback() _  
- Implements IDbTransaction.Rollback
```

##### Syntax: C#

```
public void Rollback();
```

##### Implements

IDbTransaction.Rollback

##### See Also

[MySQLTransaction Class](#) , [MySQL.Data.MySqlClient Namespace](#)

#### 18.2.4.1.2.1.1.3.1.1.10.2. MySqlConnection.BeginTransaction Method (IsolationLevel)

##### Syntax: Visual Basic

```
Overloads Public Function BeginTransaction( _  
    ByVal iso As IsolationLevel _  
    ) As MySQLTransaction
```

##### Syntax: C#

```
public MySQLTransaction BeginTransaction(  
    IsolationLevel iso  
);
```

##### See Also

[MySQLConnection Class](#) , [MySQL.Data.MySqlClient Namespace](#) , [MySQLConnection.BeginTransaction](#)



[Overload List](#)

## 18.2.4.1.2.1.1.3.1.1.11. MySqlConnection.ChangeDatabase Method

**Syntax: Visual Basic**

```
NotOverridable Public Sub ChangeDatabase( _  
    ByVal databaseName As String _  
) _  
- Implements IDbConnection.ChangeDatabase
```

**Syntax: C#**

```
public void ChangeDatabase(  
string databaseName  
);
```

**Implements**

IDbConnection.ChangeDatabase

**See Also**[MySqlConnection Class](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.1.3.1.1.12. MySqlConnection.Close Method

**Syntax: Visual Basic**

```
NotOverridable Public Sub Close() _  
- Implements IDbConnection.Close
```

**Syntax: C#**

```
public void Close();
```

**Implements**

IDbConnection.Close

**See Also**[MySqlConnection Class](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.1.3.1.1.13. MySqlConnection.CreateCommand Method

**Syntax: Visual Basic**

```
Public Function CreateCommand() As MySqlCommand
```

**Syntax: C#**

```
public MySqlCommand CreateCommand();
```

**See Also**[MySqlConnection Class](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.1.3.1.1.14. MySqlConnection.Open Method

**Syntax: Visual Basic**

```
NotOverridable Public Sub Open() _  
    Implements IDbConnection.Open
```

**Syntax: C#**

```
public void Open();
```

**Implements**

IDbConnection.Open

**See Also**

[MySqlConnection Class](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.1.3.1.1.15. MySqlConnection.Ping Method

Ping

**Syntax: Visual Basic**

```
Public Function Ping() As Boolean
```

**Syntax: C#**

```
public bool Ping();
```

**Return Value****See Also**

[MySqlConnection Class](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.1.3.1.1.16. MySqlConnection.InfoMessage Event

**Syntax: Visual Basic**

```
Public Event InfoMessage As MySqlInfoMessageEventHandler
```

**Syntax: C#**

```
public event MySqlInfoMessageEventHandler InfoMessage;
```

**See Also**

[MySqlConnection Class](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.1.3.1.1.16.1. MySqlInfoMessageEventHandler Delegate

Represents the method that will handle the [InfoMessage](#) event of a [MySqlConnection](#) .

**Syntax: Visual Basic**

```
Public Delegate Sub MySqlInfoMessageEventHandler( _
    ByVal sender As Object, _
    ByVal args As MySqlInfoMessageEventArgs _
)
```

**Syntax: C#**

```
public delegate void MySqlInfoMessageEventHandler(
    object sender,
    MySqlInfoMessageEventArgs args
);
```

**Requirements**

Namespace: [MySql.Data.MySqlClient](#)

Assembly: MySql.Data (in MySql.Data.dll)

**See Also**

[MySql.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.1.3.1.1.16.1.1. MySqlInfoMessageEventArgs Class

Provides data for the InfoMessage event. This class cannot be inherited.

For a list of all members of this type, see [MySqlInfoMessageEventArgs Members](#) .

**Syntax: Visual Basic**

```
Public Class MySqlInfoMessageEventArgs_
    Inherits EventArgs
```

**Syntax: C#**

```
public class MySqlInfoMessageEventArgs : EventArgs
```

**Thread Safety**

Public static (Shared in Visual Basic) members of this type are safe for multithreaded operations. Instance members are not guaranteed to be thread-safe.

**Requirements**

Namespace: [MySql.Data.MySqlClient](#)

Assembly: MySql.Data (in MySql.Data.dll)

**See Also**

[MySqlInfoMessageEventArgs Members](#) , [MySql.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.1.3.1.1.16.1.1.1. MySqlInfoMessageEventArgs Members

[MySqlInfoMessageEventArgs overview](#)

**Public Instance Constructors**

<a href="#">MySqlInfoMessageEventArgs Constructor</a>	Initializes a new instance of the <a href="#">MySqlInfoMessageEventArgs</a> class.
---	--

**Public Instance Fields**

<a href="#">errors</a>	
------------------------	--

**Public Instance Methods**

Equals(inherited from Object)	Determines whether the specified Object is equal to the current Object.
GetHashCode(inherited from Object)	Serves as a hash function for a particular type. GetHashCode is suitable for use in hashing algorithms and data structures like a hash table.
GetType(inherited from Object)	Gets the Type of the current instance.
ToString(inherited from Object)	Returns a String that represents the current Object.

**Protected Instance Methods**

Finalize(inherited from Object)	Allows an Object to attempt to free resources and perform other cleanup operations before the Object is reclaimed by garbage collection.
MemberwiseClone(inherited from Object)	Creates a shallow copy of the current Object.

**See Also**

[MySQLInfoMessageEventArgs Class](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.1.3.1.1.16.1.1.1.1. MySQLInfoMessageEventArgs Constructor

Initializes a new instance of the [MySQLInfoMessageEventArgs](#) class.

**Syntax: Visual Basic**

```
Public Sub New()
```

**Syntax: C#**

```
public MySQLInfoMessageEventArgs();
```

**See Also**

[MySQLInfoMessageEventArgs Class](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.1.3.1.1.16.1.1.1.2. MySQLInfoMessageEventArgs.errors Field

**Syntax: Visual Basic**

```
Public errors As MySQLError()
```

**Syntax: C#**

```
public MySQLError[] errors;
```

**See Also**

[MySQLInfoMessageEventArgs Class](#) , [MySQL.Data.MySqlClient Namespace](#)

#### 18.2.4.1.2.1.1.3.1.1.16.1.1.1.2.1. MySQLError Class

Collection of error codes that can be returned by the server

For a list of all members of this type, see [MySQLError Members](#) .

#### Syntax: Visual Basic

```
Public Class MySQLError
```

#### Syntax: C#

```
public class MySQLError
```

#### Thread Safety

Public static (Shared in Visual Basic) members of this type are safe for multithreaded operations. Instance members are not guaranteed to be thread-safe.

#### Requirements

Namespace: [MySQL.Data.MySqlClient](#)

Assembly: MySQL.Data (in MySQL.Data.dll)

#### See Also

[MySQLError Members](#) , [MySQL.Data.MySqlClient Namespace](#)

#### 18.2.4.1.2.1.1.3.1.1.16.1.1.1.2.1.1. MySQLError Members

[MySQLError overview](#)

#### Public Instance Constructors

<a href="#">MySQLError Constructor</a>	
--	--

#### Public Instance Properties

<a href="#">Code</a>	Error code
<a href="#">Level</a>	Error level
<a href="#">Message</a>	Error message

#### Public Instance Methods

<a href="#">Equals</a> (inherited from Object)	Determines whether the specified Object is equal to the current Object.
<a href="#">GetHashCode</a> (inherited from Object)	Serves as a hash function for a particular type. GetHashCode is suitable for use in hashing algorithms and data structures like a hash table.
<a href="#">GetType</a> (inherited from Object)	Gets the Type of the current instance.
<a href="#">ToString</a> (inherited from Object)	Returns a String that represents the current Object.

**Protected Instance Methods**

Finalize(inherited from Object)	Allows an Object to attempt to free resources and perform other cleanup operations before the Object is reclaimed by garbage collection.
MemberwiseClone(inherited from Object)	Creates a shallow copy of the current Object.

**See Also**

[MySQLError Class](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.1.3.1.1.16.1.1.1.2.1.1.1. MySQLError Constructor

**Syntax: Visual Basic**

```
Public Sub New( _
    ByVal level As String, _
    ByVal code As Integer, _
    ByVal message As String _
)
```

**Syntax: C#**

```
public MySQLError(
    string level,
    int code,
    string message
);
```

**Parameters**

- `level`:
- `code`:
- `message`:

**See Also**

[MySQLError Class](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.1.3.1.1.16.1.1.1.2.1.1.2. Code Property

Error code

**Syntax: Visual Basic**

```
Public ReadOnly Property Code As Integer
```

**Syntax: C#**

```
public int Code {get;}
```

**See Also**

[MySQLError Class](#) , [MySQL.Data.MySqlClient Namespace](#)

#### 18.2.4.1.2.1.1.3.1.1.16.1.1.1.2.1.1.3. Level Property

Error level

**Syntax: Visual Basic**

```
Public ReadOnly Property Level As String
```

**Syntax: C#**

```
public string Level {get;}
```

**See Also**

[MySQLError Class](#) , [MySQL.Data.MySqlClient Namespace](#)

#### 18.2.4.1.2.1.1.3.1.1.16.1.1.1.2.1.1.4. Message Property

Error message

**Syntax: Visual Basic**

```
Public ReadOnly Property Message As String
```

**Syntax: C#**

```
public string Message {get;}
```

**See Also**

[MySQLError Class](#) , [MySQL.Data.MySqlClient Namespace](#)

#### 18.2.4.1.2.1.1.3.1.1.17. MySqlConnection.StateChange Event

**Syntax: Visual Basic**

```
Public Event StateChange As StateChangeEventHandler
```

**Syntax: C#**

```
public event StateChangeEventHandler StateChange;
```

**See Also**

[MySqlConnection Class](#) , [MySQL.Data.MySqlClient Namespace](#)

#### 18.2.4.1.2.1.1.4. MySqlCommand Constructor (String, MySqlConnection, MySqlTransaction)

**Syntax: Visual Basic**

```
Overloads Public Sub New( _  
    ByVal cmdText As String, _  
    ByVal connection As MySqlConnection, _  
    ByVal transaction As MySqlTransaction _  
)
```

**Syntax: C#**

```
public MySqlCommand(  
    string cmdText,  
    MySqlConnection connection,  
    MySqlTransaction transaction  
);
```

**See Also**

[MySqlCommand Class](#), [MySql.Data.MySqlClient Namespace](#), [MySqlCommand Constructor Overload List](#)

## 18.2.4.1.2.1.2. CommandText Property

**Syntax: Visual Basic**

```
NotOverridable Public Property CommandText As String _  
    Implements IDbCommand.CommandText
```

**Syntax: C#**

```
public string CommandText {get; set;}
```

**Implements**

IDbCommand.CommandText

**See Also**

[MySqlCommand Class](#), [MySql.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.3. CommandTimeout Property

**Syntax: Visual Basic**

```
NotOverridable Public Property CommandTimeout As Integer _  
    Implements IDbCommand.CommandTimeout
```

**Syntax: C#**

```
public int CommandTimeout {get; set;}
```

**Implements**

IDbCommand.CommandTimeout

**See Also**

[MySqlCommand Class](#), [MySql.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.4. CommandType Property

**Syntax: Visual Basic**

```
NotOverridable Public Property CommandType As CommandType _  
    Implements IDbCommand.CommandType
```

**Syntax: C#**



```
public System.Data.CommandType CommandType {get; set;}
```

**Implements**

IDbCommand.CommandType

**See Also**

[MySQLCommand Class](#) , [MySql.Data.MySqlClient Namespace](#)

18.2.4.1.2.1.5. Connection Property

**Syntax: Visual Basic**

```
Public Property Connection As MySqlConnection
```

**Syntax: C#**

```
public MySqlConnection Connection {get; set;}
```

**See Also**

[MySQLCommand Class](#) , [MySql.Data.MySqlClient Namespace](#)

18.2.4.1.2.1.6. IsPrepared Property

**Syntax: Visual Basic**

```
Public ReadOnly Property IsPrepared As Boolean
```

**Syntax: C#**

```
public bool IsPrepared {get;}
```

**See Also**

[MySQLCommand Class](#) , [MySql.Data.MySqlClient Namespace](#)

18.2.4.1.2.1.7. Parameters Property

**Syntax: Visual Basic**

```
Public ReadOnly Property Parameters As MySqlParameterCollection
```

**Syntax: C#**

```
public MySqlParameterCollection Parameters {get;}
```

**See Also**

[MySQLCommand Class](#) , [MySql.Data.MySqlClient Namespace](#)

18.2.4.1.2.1.7.1. MySqlParameterCollection Class

Represents a collection of parameters relevant to a [MySQLCommand](#) as well as their respective mappings to columns in a DataSet. This class cannot be inherited.

For a list of all members of this type, see [MySQLParameterCollection Members](#) .

### Syntax: Visual Basic

```
NotInheritable Public Class MySQLParameterCollection_
    Inherits MarshalByRefObject_
    Implements IDataParameterCollection, IList, ICollection, IEnumerable
```

### Syntax: C#

```
public sealed class MySQLParameterCollection : MarshalByRefObject, IDataParameterCollection, IList, IC
```

### Thread Safety

Public static (Shared in Visual Basic) members of this type are safe for multithreaded operations. Instance members are not guaranteed to be thread-safe.

### Requirements

Namespace: [MySQL.Data.MySqlClient](#)

Assembly: MySQL.Data (in MySQL.Data.dll)

### See Also

[MySQLParameterCollection Members](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.7.1.1. MySQLParameterCollection Members

[MySQLParameterCollection overview](#)

### Public Instance Constructors

<a href="#">MySQLParameterCollection Constructor</a>	Initializes a new instance of the <a href="#">MySQLParameterCollection</a> class.
--	---

### Public Instance Properties

<a href="#">Count</a>	Gets the number of MySQLParameter objects in the collection.
<a href="#">Item</a>	Overloaded. Gets the <a href="#">MySQLParameter</a> with a specified attribute. In C#, this property is the indexer for the <a href="#">MySQLParameterCollection</a> class.

### Public Instance Methods

<a href="#">Add</a>	Overloaded. Adds the specified <a href="#">MySQLParameter</a> object to the <a href="#">MySQLParameterCollection</a> .
<a href="#">Clear</a>	Removes all items from the collection.
<a href="#">Contains</a>	Overloaded. Gets a value indicating whether a <a href="#">MySQLParameter</a> exists in the collection.
<a href="#">CopyTo</a>	Copies MySQLParameter objects from the MySQLParameterCollection to the specified array.
<a href="#">CreateObjRef</a> (inherited from MarshalByRefOb-	Creates an object that contains all the relevant in-

ject)	formation required to generate a proxy used to communicate with a remote object.
Equals(inherited from Object)	Determines whether the specified Object is equal to the current Object.
GetHashCode(inherited from Object)	Serves as a hash function for a particular type. GetHashCode is suitable for use in hashing algorithms and data structures like a hash table.
GetLifetimeService(inherited from MarshalByRefObject)	Retrieves the current lifetime service object that controls the lifetime policy for this instance.
GetType(inherited from Object)	Gets the Type of the current instance.
<a href="#">IndexOf</a>	Overloaded. Gets the location of a <a href="#">MySQLParameter</a> in the collection.
InitializeLifetimeService(inherited from MarshalByRefObject)	Obtains a lifetime service object to control the lifetime policy for this instance.
<a href="#">Insert</a>	Inserts a <a href="#">MySQLParameter</a> into the collection at the specified index.
<a href="#">Remove</a>	Removes the specified <a href="#">MySQLParameter</a> from the collection.
<a href="#">RemoveAt</a>	Overloaded. Removes the specified <a href="#">MySQLParameter</a> from the collection.
ToString(inherited from Object)	Returns a String that represents the current Object.

**See Also**

[MySQLParameterCollection Class](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.7.1.1.1. MySQLParameterCollection Constructor

Initializes a new instance of the [MySQLParameterCollection](#) class.

**Syntax: Visual Basic**

```
Public Sub New()
```

**Syntax: C#**

```
public MySQLParameterCollection();
```

**See Also**

[MySQLParameterCollection Class](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.7.1.1.2. Count Property

Gets the number of [MySQLParameter](#) objects in the collection.

**Syntax: Visual Basic**

```
NotOverridable Public ReadOnly Property Count As Integer _  
_ Implements ICollection.Count
```

**Syntax: C#**

```
public int Count {get;}
```

**Implements**

ICollection.Count

**See Also**

[MySQLParameterCollection Class](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.7.1.1.3. Item Property

Gets the [MySQLParameter](#) with a specified attribute. In C#, this property is the indexer for the [MySQLParameterCollection](#) class.

**Overload List**

Gets the [MySQLParameter](#) at the specified index.

- `public MySQLParameter this[int] {get; set;}`

Gets the [MySQLParameter](#) with the specified name.

- `public MySQLParameter this[string] {get; set;}`

**See Also**

[MySQLParameterCollection Class](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.7.1.1.3.1. MySQLParameter Class

Represents a parameter to a [MySQLCommand](#) , and optionally, its mapping to DataSetcolumns. This class cannot be inherited.

For a list of all members of this type, see [MySQLParameter Members](#) .

**Syntax: Visual Basic**

```
NotInheritable Public Class MySQLParameter_
    Inherits MarshalByRefObject_
    Implements IDataParameter, IDbDataParameter, ICloneable
```

**Syntax: C#**

```
public sealed class MySQLParameter : MarshalByRefObject, IDataParameter, IDbDataParameter, ICloneable
```

**Thread Safety**

Public static (Shared in Visual Basic) members of this type are safe for multithreaded operations. Instance members are not guaranteed to be thread-safe.

**Requirements**

Namespace: [MySQL.Data.MySqlClient](#)

Assembly: MySql.Data (in MySql.Data.dll)

**See Also**

[MySqlParameter Members](#) , [MySql.Data.MySqlClient Namespace](#)

18.2.4.1.2.1.7.1.1.3.1.1. MySqlConnection Members

[MySqlParameter overview](#)

**Public Instance Constructors**

<a href="#">MySqlParameter</a>	Overloaded. Initializes a new instance of the MySqlConnection class.
--------------------------------	--

**Public Instance Properties**

<a href="#">DbType</a>	Gets or sets the DbType of the parameter.
<a href="#">Direction</a>	Gets or sets a value indicating whether the parameter is input-only, output-only, bidirectional, or a stored procedure return value parameter. As of MySql version 4.1 and earlier, input-only is the only valid choice.
<a href="#">IsNullable</a>	Gets or sets a value indicating whether the parameter accepts null values.
<a href="#">IsUnsigned</a>	
<a href="#">MySqlDbType</a>	Gets or sets the MySqlDbType of the parameter.
<a href="#">ParameterName</a>	Gets or sets the name of the MySqlConnection.
<a href="#">Precision</a>	Gets or sets the maximum number of digits used to represent the <a href="#">Value</a> property.
<a href="#">Scale</a>	Gets or sets the number of decimal places to which <a href="#">Value</a> is resolved.
<a href="#">Size</a>	Gets or sets the maximum size, in bytes, of the data within the column.
<a href="#">SourceColumn</a>	Gets or sets the name of the source column that is mapped to the DataSet and used for loading or returning the <a href="#">Value</a> .
<a href="#">SourceVersion</a>	Gets or sets the DataRowVersion to use when loading <a href="#">Value</a> .
<a href="#">Value</a>	Gets or sets the value of the parameter.

**Public Instance Methods**

<a href="#">CreateObjRef</a> (inherited from MarshalByRefObject)	Creates an object that contains all the relevant information required to generate a proxy used to communicate with a remote object.
<a href="#">Equals</a> (inherited from Object)	Determines whether the specified Object is equal to the current Object.
<a href="#">GetHashCode</a> (inherited from Object)	Serves as a hash function for a particular type. GetHashCode is suitable for use in hashing algorithms and data structures like a hash table.

GetLifetimeService(inherited from MarshalByRefObject)	Retrieves the current lifetime service object that controls the lifetime policy for this instance.
GetType(inherited from Object)	Gets the Type of the current instance.
InitializeLifetimeService(inherited from MarshalByRefObject)	Obtains a lifetime service object to control the lifetime policy for this instance.
<a href="#">ToString</a>	Overridden. Gets a string containing the <a href="#">ParameterName</a> .

**See Also**

[MySQLParameter Class](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.7.1.1.3.1.1.1. MySQLParameter Constructor

Initializes a new instance of the MySQLParameter class.

**Overload List**

Initializes a new instance of the MySQLParameter class.

- [public MySQLParameter\(\);](#)

Initializes a new instance of the [MySQLParameter](#) class with the parameter name and the data type.

- [public MySQLParameter\(string,MySQLDbType\);](#)

Initializes a new instance of the [MySQLParameter](#) class with the parameter name, the [MySQLDbType](#) , and the size.

- [public MySQLParameter\(string,MySQLDbType,int\);](#)

Initializes a new instance of the [MySQLParameter](#) class with the parameter name, the type of the parameter, the size of the parameter, a [ParameterDirection](#), the precision of the parameter, the scale of the parameter, the source column, a [DataRowVersion](#) to use, and the value of the parameter.

- [public MySQLParameter\(string,MySQLDbType,int,ParameterDirection,bool,byte,byte,string,DataRowVersion,object\);](#)

Initializes a new instance of the [MySQLParameter](#) class with the parameter name, the [MySQLDbType](#) , the size, and the source column name.

- [public MySQLParameter\(string,MySQLDbType,int,string\);](#)

Initializes a new instance of the [MySQLParameter](#) class with the parameter name and a value of the new MySQLParameter.

- [public MySQLParameter\(string,object\);](#)

**See Also**

[MySQLParameter Class](#) , [MySQL.Data.MySqlClient Namespace](#)

#### 18.2.4.1.2.1.7.1.1.3.1.1.1.1. MySQLParameter Constructor ()

Initializes a new instance of the [MySQLParameter](#) class.

**Syntax: Visual Basic**

```
Overloads Public Sub New()
```

**Syntax: C#**

```
public MySQLParameter();
```

**See Also**

[MySQLParameter Class](#) , [MySQL.Data.MySqlClient Namespace](#) , [MySQLParameter Constructor Overload List](#)

#### 18.2.4.1.2.1.7.1.1.3.1.1.1.2. MySQLParameter Constructor (String, MySQLDbType)

Initializes a new instance of the [MySQLParameter](#) class with the parameter name and the data type.

**Syntax: Visual Basic**

```
Overloads Public Sub New( _  
    ByVal parameterName As String, _  
    ByVal dbType As MySQLDbType _  
)
```

**Syntax: C#**

```
public MySQLParameter(  
    stringparameterName,  
    MySQLDbTypedbType  
);
```

**Parameters**

- [parameterName](#): The name of the parameter to map.
- [dbType](#): One of the [MySQLDbType](#) values.

**See Also**

[MySQLParameter Class](#) , [MySQL.Data.MySqlClient Namespace](#) , [MySQLParameter Constructor Overload List](#)

#### 18.2.4.1.2.1.7.1.1.3.1.1.2.1. MySQLDbType Enumeration

Specifies MySQL specific data type of a field, property, for use in a [MySQLParameter](#) .

**Syntax: Visual Basic**

```
Public Enum MySqlDbType
```

**Syntax: C#**

```
public enum MySqlDbType
```

**Members**

Member Name	Description
VarChar	A variable-length string containing 0 to 65535 characters
Timestamp	A timestamp. The range is '1970-01-01 00:00:00' to sometime in the year 2037
LongBlob	A BLOB or TEXT column with a maximum length of 4294967295 or 4G (2 <sup>32</sup> - 1) characters
Time	Time The range is '-838:59:59' to '838:59:59'.
TinyBlob	A BLOB or TEXT column with a maximum length of 255 (2 <sup>8</sup> - 1) characters
Datetime	DateTime The supported range is '1000-01-01 00:00:00' to '9999-12-31 23:59:59'.
Decimal	Decimal A fixed precision and scale numeric value between -1038 -1 and 10 38 -1.
UByte	
Blob	A BLOB or TEXT column with a maximum length of 65535 (2 <sup>16</sup> - 1) characters
Double	Double A normal-size (double-precision) floating-point number. Allowable values are - 1.7976931348623157E+308 to - 2.2250738585072014E-308, 0, and 2.2250738585072014E-308 to 1.7976931348623157E+308.
Newdate	Obsolete Use Datetime or Date type
Byte	Byte The signed range is -128 to 127. The unsigned range is 0 to 255.
Date	Date The supported range is '1000-01-01' to '9999-12-31'.
VarChar	A variable-length string containing 0 to 255 characters
UInt16	
UInt24	
Int16	Int16 A 16-bit signed integer. The signed range is -



	32768 to 32767. The unsigned range is 0 to 65535
NewDecimal	New Decimal
Set	A set. A string object that can have zero or more values, each of which must be chosen from the list of values 'value1', 'value2', ... A SET can have a maximum of 64 members.
String	Obsolete Use VarChar type
Enum	An enumeration. A string object that can have only one value, chosen from the list of values 'value1', 'value2', ..., NULL or the special "" error value. An ENUM can have a maximum of 65535 distinct values
Geometry	
UInt64	
Int64	Int64 A 64-bit signed integer.
UInt32	
Int24	Specifies a 24 (3 byte) signed or unsigned value.
Bit	Bit-field data type
Float	Single A small (single-precision) floating-point number. Allowable values are -3.402823466E+38 to -1.175494351E-38, 0, and 1.175494351E-38 to 3.402823466E+38.
Year	A year in 2- or 4-digit format (default is 4-digit). The allowable values are 1901 to 2155, 0000 in the 4-digit year format, and 1970-2069 if you use the 2-digit format (70-69)
Int32	Int32 A 32-bit signed integer
MediumBlob	A BLOB or TEXT column with a maximum length of 16777215 (2 <sup>24</sup> - 1) characters

**Requirements**

Namespace: [MySql.Data.MySqlClient](#)

Assembly: MySql.Data (in MySql.Data.dll)

**See Also**

[MySql.Data.MySqlClient Namespace](#)

#### 18.2.4.1.2.1.7.1.1.3.1.1.3. MySqlParameter Constructor (String, MySqlDbType, Int32)

Initializes a new instance of the [MySqlParameter](#) class with the parameter name, the [MySqlDbType](#) , and the size.

**Syntax: Visual Basic**

```
Overloads Public Sub New( _
    ByVal parameterName As String, _
    ByVal dbType As MySqlDbType, _
    ByVal size As Integer _
)
```

**Syntax: C#**

```
public MySqlParameter(
    stringparameterName,
    MySqlDbTypedbType,
    intsize
);
```

**Parameters**

- `parameterName`: The name of the parameter to map.
- `dbType`: One of the [MySqlDbType](#) values.
- `size`: The length of the parameter.

**See Also**

[MySqlParameter Class](#) , [MySql.Data.MySqlClient Namespace](#) , [MySqlParameter Constructor Overload List](#)

#### 18.2.4.1.2.1.7.1.1.3.1.1.1.4. MySqlParameter Constructor (String, MySqlDbType, Int32, ParameterDirection, Boolean, Byte, Byte, String, DataRowVersion, Object)

Initializes a new instance of the [MySqlParameter](#) class with the parameter name, the type of the parameter, the size of the parameter, a [ParameterDirection](#), the precision of the parameter, the scale of the parameter, the source column, a [DataRowVersion](#) to use, and the value of the parameter.

**Syntax: Visual Basic**

```
Overloads Public Sub New( _
    ByVal parameterName As String, _
    ByVal dbType As MySqlDbType, _
    ByVal size As Integer, _
    ByVal direction As ParameterDirection, _
    ByVal isNullable As Boolean, _
    ByVal precision As Byte, _
    ByVal scale As Byte, _
    ByVal sourceColumn As String, _
    ByVal sourceVersion As DataRowVersion, _
    ByVal value As Object _
)
```

**Syntax: C#**

```
public MySqlParameter(
    stringparameterName,
    MySqlDbTypedbType,
    intsize,
    ParameterDirectiondirection,
    boolisNullable,
    byteprecision,
    bytescale,
    stringsourceColumn,
    DataRowVersionsourceVersion,
    objectvalue
);
```

```
);
```

## Parameters

- `parameterName`: The name of the parameter to map.
- `dbType`: One of the [MySQLDbType](#) values.
- `size`: The length of the parameter.
- `direction`: One of the [ParameterDirection](#) values.
- `isNullable`: true if the value of the field can be null, otherwise false.
- `precision`: The total number of digits to the left and right of the decimal point to which [Value](#) is resolved.
- `scale`: The total number of decimal places to which [Value](#) is resolved.
- `sourceColumn`: The name of the source column.
- `sourceVersion`: One of the [DataRowVersion](#) values.
- `value`: An Object that is the value of the [MySQLParameter](#).

## Exceptions

Exception Type	Condition
ArgumentException	

## See Also

[MySQLParameter Class](#), [MySQL.Data.MySqlClient Namespace](#), [MySQLParameter Constructor Overload List](#)

### 18.2.4.1.2.1.7.1.1.3.1.1.4.1. Value Property

Gets or sets the value of the parameter.

#### Syntax: Visual Basic

```
NotOverridable Public Property Value As Object _
    Implements IDataParameter.Value
```

#### Syntax: C#

```
public object Value {get; set;}
```

## Implements

[IDataParameter.Value](#)

## See Also

[MySQLParameter Class](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.7.1.1.3.1.1.1.5. MySQLParameter Constructor (String, MySqlDbType, Int32, String)

Initializes a new instance of the [MySQLParameter](#) class with the parameter name, the [MySqlDbType](#) , the size, and the source column name.

**Syntax: Visual Basic**

```
Overloads Public Sub New( _  
    ByVal parameterName As String, _  
    ByVal dbType As MySqlDbType, _  
    ByVal size As Integer, _  
    ByVal sourceColumn As String _  
)
```

**Syntax: C#**

```
public MySQLParameter(  
    stringparameterName,  
    MySqlDbTypedbType,  
    intsize,  
    stringsourceColumn  
);
```

**Parameters**

- [parameterName](#): The name of the parameter to map.
- [dbType](#): One of the [MySqlDbType](#) values.
- [size](#): The length of the parameter.
- [sourceColumn](#): The name of the source column.

**See Also**

[MySQLParameter Class](#) , [MySQL.Data.MySqlClient Namespace](#) , [MySQLParameter Constructor Overload List](#)

## 18.2.4.1.2.1.7.1.1.3.1.1.1.6. MySQLParameter Constructor (String, Object)

Initializes a new instance of the [MySQLParameter](#) class with the parameter name and a value of the new [MySQLParameter](#).

**Syntax: Visual Basic**

```
Overloads Public Sub New( _  
    ByVal parameterName As String, _  
    ByVal value As Object _  
)
```

**Syntax: C#**

```
public MySQLParameter(  
    stringparameterName,  
    objectvalue  
);
```

**Parameters**

- `parameterName`: The name of the parameter to map.
- `value`: An Object that is the value of the [MySQLParameter](#) .

**See Also**

[MySQLParameter Class](#) , [MySQL.Data.MySqlClient Namespace](#) , [MySQLParameter Constructor Overload List](#)

## 18.2.4.1.2.1.7.1.1.3.1.1.2. DbType Property

Gets or sets the DbType of the parameter.

**Syntax: Visual Basic**

```
NotOverridable Public Property DbType As DbType _  
    Implements IDataParameter.DbType
```

**Syntax: C#**

```
public System.Data.DbType DbType {get; set;}
```

**Implements**

IDataParameter.DbType

**See Also**

[MySQLParameter Class](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.7.1.1.3.1.1.3. Direction Property

Gets or sets a value indicating whether the parameter is input-only, output-only, bidirectional, or a stored procedure return value parameter. As of MySQL version 4.1 and earlier, input-only is the only valid choice.

**Syntax: Visual Basic**

```
NotOverridable Public Property Direction As ParameterDirection _  
    Implements IDataParameter.Direction
```

**Syntax: C#**

```
public System.Data.ParameterDirection Direction {get; set;}
```

**Implements**

IDataParameter.Direction

**See Also**

[MySQLParameter Class](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.7.1.1.3.1.1.4. IsNullable Property

Gets or sets a value indicating whether the parameter accepts null values.

**Syntax: Visual Basic**

```
NotOverridable Public Property IsNullable As Boolean _  
    Implements IDataParameter.IsNullable
```

**Syntax: C#**

```
public bool IsNullable {get; set;}
```

**Implements**

IDataParameter.IsNullable

**See Also**

[MySQLParameter Class](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.7.1.1.3.1.1.5. IsUnsigned Property

**Syntax: Visual Basic**

```
Public Property IsUnsigned As Boolean
```

**Syntax: C#**

```
public bool IsUnsigned {get; set;}
```

**See Also**

[MySQLParameter Class](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.7.1.1.3.1.1.6. MySQLDbType Property

Gets or sets the MySQLDbType of the parameter.

**Syntax: Visual Basic**

```
Public Property MySQLDbType As MySQLDbType
```

**Syntax: C#**

```
public MySQLDbType MySQLDbType {get; set;}
```

**See Also**

[MySQLParameter Class](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.7.1.1.3.1.1.7. ParameterName Property

Gets or sets the name of the MySQLParameter.

**Syntax: Visual Basic**

```
NotOverridable Public Property ParameterName As String _  
    Implements IDataParameter.ParameterName
```

**Syntax: C#**

```
public string ParameterName {get; set;}
```

**Implements**

IDataParameter.ParameterName

**See Also**

[MySQLParameter Class](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.7.1.1.3.1.1.8. Precision Property

Gets or sets the maximum number of digits used to represent the [Value](#) property.

**Syntax: Visual Basic**

```
NotOverridable Public Property Precision As Byte _  
- Implements IDbDataParameter.Precision
```

**Syntax: C#**

```
public byte Precision {get; set;}
```

**Implements**

IDbDataParameter.Precision

**See Also**

[MySQLParameter Class](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.7.1.1.3.1.1.9. Scale Property

Gets or sets the number of decimal places to which [Value](#) is resolved.

**Syntax: Visual Basic**

```
NotOverridable Public Property Scale As Byte _  
- Implements IDbDataParameter.Scale
```

**Syntax: C#**

```
public byte Scale {get; set;}
```

**Implements**

IDbDataParameter.Scale

**See Also**

[MySQLParameter Class](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.7.1.1.3.1.1.10. Size Property

Gets or sets the maximum size, in bytes, of the data within the column.

**Syntax: Visual Basic**

```
NotOverridable Public Property Size As Integer _  
    Implements IDbDataParameter.Size
```

**Syntax: C#**

```
public int Size {get; set;}
```

**Implements**

IDbDataParameter.Size

**See Also**

[MySQLParameter Class](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.7.1.1.3.1.1.11. SourceColumn Property

Gets or sets the name of the source column that is mapped to the DataSet and used for loading or returning the [Value](#) .

**Syntax: Visual Basic**

```
NotOverridable Public Property SourceColumn As String _  
    Implements IDataParameter.SourceColumn
```

**Syntax: C#**

```
public string SourceColumn {get; set;}
```

**Implements**

IDataParameter.SourceColumn

**See Also**

[MySQLParameter Class](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.7.1.1.3.1.1.12. SourceVersion Property

Gets or sets the DataRowVersion to use when loading [Value](#) .

**Syntax: Visual Basic**

```
NotOverridable Public Property SourceVersion As DataRowVersion _  
    Implements IDataParameter.SourceVersion
```

**Syntax: C#**

```
public System.Data.DataRowVersion SourceVersion {get; set;}
```

**Implements**

IDataParameter.SourceVersion



**See Also**

[MySQLParameter Class](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.7.1.1.3.1.1.13. MySQLParameter.ToString Method

Overridden. Gets a string containing the [ParameterName](#) .

**Syntax: Visual Basic**

```
Overrides Public Function ToString() As String
```

**Syntax: C#**

```
public override string ToString();
```

**Return Value****See Also**

[MySQLParameter Class](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.7.1.1.3.2. Item Property (Int32)

Gets the [MySQLParameter](#) at the specified index.

**Syntax: Visual Basic**

```
Overloads Public Default Property Item( _  
    ByVal index As Integer _  
) As MySQLParameter
```

**Syntax: C#**

```
public MySQLParameter this[  
    int index  
] {get; set;}
```

**See Also**

[MySQLParameterCollection Class](#) , [MySQL.Data.MySqlClient Namespace](#) , [MySQLParameterCollection.Item Overload List](#)

## 18.2.4.1.2.1.7.1.1.3.3. Item Property (String)

Gets the [MySQLParameter](#) with the specified name.

**Syntax: Visual Basic**

```
Overloads Public Default Property Item( _  
    ByVal name As String _  
) As MySQLParameter
```

**Syntax: C#**

```
public MySQLParameter this[  
    string name  
] {get; set;}
```

**See Also**

[MySQLParameterCollection Class](#) , [MySQL.Data.MySqlClient Namespace](#) , [MySQLParameterCollection.Item Overload List](#)

## 18.2.4.1.2.1.7.1.1.4. Add Method

Adds the specified [MySQLParameter](#) object to the [MySQLParameterCollection](#) .

**Overload List**

Adds the specified [MySQLParameter](#) object to the [MySQLParameterCollection](#) .

- `public MySQLParameter Add(MySQLParameter);`

Adds the specified [MySQLParameter](#) object to the [MySQLParameterCollection](#) .

- `public int Add(object);`

Adds a [MySQLParameter](#) to the [MySQLParameterCollection](#) given the parameter name and the data type.

- `public MySQLParameter Add(string,MySQLDbType);`

Adds a [MySQLParameter](#) to the [MySQLParameterCollection](#) with the parameter name, the data type, and the column length.

- `public MySQLParameter Add(string,MySQLDbType,int);`

Adds a [MySQLParameter](#) to the [MySQLParameterCollection](#) with the parameter name, the data type, the column length, and the source column name.

- `public MySQLParameter Add(string,MySQLDbType,int,string);`

Adds a [MySQLParameter](#) to the [MySQLParameterCollection](#) given the specified parameter name and value.

- `public MySQLParameter Add(string,object);`

**See Also**

[MySQLParameterCollection Class](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.7.1.1.4.1. MySQLParameterCollection.Add Method (MySQLParameter)

Adds the specified [MySQLParameter](#) object to the [MySQLParameterCollection](#) .

**Syntax: Visual Basic**

```
Overloads Public Function Add( _
```

```
ByVal value As MySqlParameter _  
) As MySqlParameter
```

**Syntax: C#**

```
public MySqlParameter Add(  
    MySqlParameter value  
);
```

**Parameters**

- `value`: The [MySqlParameter](#) to add to the collection.

**Return Value**

The newly added [MySqlParameter](#) object.

**See Also**

[MySqlParameterCollection Class](#) , [MySql.Data.MySqlClient Namespace](#) , [MySqlParameterCollection.Add Overload List](#)

#### 18.2.4.1.2.1.7.1.1.4.2. MySqlConnection.Add Method (Object)

Adds the specified [MySqlParameter](#) object to the [MySqlParameterCollection](#) .

**Syntax: Visual Basic**

```
NotOverridable Overloads Public Function Add( _  
    ByVal value As Object _  
    ) As Integer _  
    Implements IList.Add
```

**Syntax: C#**

```
public int Add(  
    object value  
);
```

**Parameters**

- `value`: The [MySqlParameter](#) to add to the collection.

**Return Value**

The index of the new [MySqlParameter](#) object.

**Implements**

[IList.Add](#)

**See Also**

[MySqlParameterCollection Class](#) , [MySql.Data.MySqlClient Namespace](#) , [MySqlParameterCollection.Add Overload List](#)

#### 18.2.4.1.2.1.7.1.1.4.3. MySqlConnection.Add Method (String, MySqlDbType)

Adds a [MySqlParameter](#) to the [MySqlParameterCollection](#) given the parameter name and the data type.

##### Syntax: Visual Basic

```
Overloads Public Function Add( _  
    ByVal parameterName As String, _  
    ByVal dbType As MySqlDbType _  
) As MySqlParameter
```

##### Syntax: C#

```
public MySqlParameter Add(  
    stringparameterName,  
    MySqlDbTypedbType  
);
```

##### Parameters

- [parameterName](#): The name of the parameter.
- [dbType](#): One of the [MySqlDbType](#) values.

##### Return Value

The newly added [MySqlParameter](#) object.

##### See Also

[MySqlParameterCollection Class](#) , [MySql.Data.MySqlClient Namespace](#) , [MySqlParameterCollection.Add Overload List](#)

#### 18.2.4.1.2.1.7.1.1.4.4. MySqlConnection.Add Method (String, MySqlDbType, Int32)

Adds a [MySqlParameter](#) to the [MySqlParameterCollection](#) with the parameter name, the data type, and the column length.

##### Syntax: Visual Basic

```
Overloads Public Function Add( _  
    ByVal parameterName As String, _  
    ByVal dbType As MySqlDbType, _  
    ByVal size As Integer _  
) As MySqlParameter
```

##### Syntax: C#

```
public MySqlParameter Add(  
    stringparameterName,  
    MySqlDbTypedbType,  
    intsize  
);
```

##### Parameters

- [parameterName](#): The name of the parameter.
- [dbType](#): One of the [MySqlDbType](#) values.

- `size`: The length of the column.

### Return Value

The newly added [MySQLParameter](#) object.

### See Also

[MySQLParameterCollection Class](#) , [MySQL.Data.MySqlClient Namespace](#) , [MySQLParameterCollection.Add Overload List](#)

#### 18.2.4.1.2.1.7.1.1.4.5. MySQLParameterCollection.Add Method (String, MySqlDbType, Int32, String)

Adds a [MySQLParameter](#) to the [MySQLParameterCollection](#) with the parameter name, the data type, the column length, and the source column name.

### Syntax: Visual Basic

```
Overloads Public Function Add( _  
    ByVal parameterName As String, _  
    ByVal dbType As MySqlDbType, _  
    ByVal size As Integer, _  
    ByVal sourceColumn As String _  
) As MySQLParameter
```

### Syntax: C#

```
public MySQLParameter Add(  
    stringparameterName,  
    MySqlDbTypedbType,  
    intsize,  
    stringsourceColumn  
);
```

### Parameters

- `parameterName`: The name of the parameter.
- `dbType`: One of the [MySQLDbType](#) values.
- `size`: The length of the column.
- `sourceColumn`: The name of the source column.

### Return Value

The newly added [MySQLParameter](#) object.

### See Also

[MySQLParameterCollection Class](#) , [MySQL.Data.MySqlClient Namespace](#) , [MySQLParameterCollection.Add Overload List](#)

#### 18.2.4.1.2.1.7.1.1.4.6. MySQLParameterCollection.Add Method (String, Object)

Adds a [MySQLParameter](#) to the [MySQLParameterCollection](#) given the specified parameter name and value.

### Syntax: Visual Basic

```
Overloads Public Function Add( _  
    ByVal parameterName As String, _  
    ByVal value As Object _  
) As MySqlParameter
```

**Syntax: C#**

```
public MySqlParameter Add(  
    stringparameterName,  
    objectvalue  
);
```

**Parameters**

- `parameterName`: The name of the parameter.
- `value`: The [Value](#) of the [MySqlParameter](#) to add to the collection.

**Return Value**

The newly added [MySqlParameter](#) object.

**See Also**

[MySqlParameterCollection Class](#) , [MySql.Data.MySqlClient Namespace](#) , [MySqlParameterCollection.Add Overload List](#)

#### 18.2.4.1.2.1.7.1.1.5. MySqlConnection.Clear Method

Removes all items from the collection.

**Syntax: Visual Basic**

```
NotOverridable Public Sub Clear() _  
    Implements IList.Clear
```

**Syntax: C#**

```
public void Clear();
```

**Implements**

[IList.Clear](#)

**See Also**

[MySqlParameterCollection Class](#) , [MySql.Data.MySqlClient Namespace](#)

#### 18.2.4.1.2.1.7.1.1.6. Contains Method

Gets a value indicating whether a [MySqlParameter](#) exists in the collection.

**Overload List**

Gets a value indicating whether a [MySqlParameter](#) exists in the collection.

- `public bool Contains(object);`

Gets a value indicating whether a [MySQLParameter](#) with the specified parameter name exists in the collection.

- `public bool Contains(string);`

#### See Also

[MySQLParameterCollection Class](#) , [MySQL.Data.MySqlClient Namespace](#)

#### 18.2.4.1.2.1.7.1.1.6.1. MySQLParameterCollection.Contains Method (Object)

Gets a value indicating whether a [MySQLParameter](#) exists in the collection.

#### Syntax: Visual Basic

```
NotOverridable Overloads Public Function Contains( _
    ByVal value As Object _
) As Boolean _
- Implements IList.Contains
```

#### Syntax: C#

```
public bool Contains(
    object value
);
```

#### Parameters

- `value`: The value of the [MySQLParameter](#) object to find.

#### Return Value

true if the collection contains the [MySQLParameter](#) object; otherwise, false.

#### Implements

[IList.Contains](#)

#### See Also

[MySQLParameterCollection Class](#) , [MySQL.Data.MySqlClient Namespace](#) , [MySQLParameterCollection.Contains Overload List](#)

#### 18.2.4.1.2.1.7.1.1.6.2. MySQLParameterCollection.Contains Method (String)

Gets a value indicating whether a [MySQLParameter](#) with the specified parameter name exists in the collection.

#### Syntax: Visual Basic

```
NotOverridable Overloads Public Function Contains( _
    ByVal name As String _
) As Boolean _
-
```

---

```
Implements IDataParameterCollection.Contains
```

**Syntax: C#**

```
public bool Contains(  
    stringname  
);
```

**Parameters**

- `name`: The name of the [MySQLParameter](#) object to find.

**Return Value**

true if the collection contains the parameter; otherwise, false.

**Implements**

[IDataParameterCollection.Contains](#)

**See Also**

[MySQLParameterCollection Class](#) , [MySQL.Data.MySqlClient Namespace](#) , [MySQLParameterCollection.Contains Overload List](#)

#### 18.2.4.1.2.1.7.1.1.7. MySQLParameterCollection.CopyTo Method

Copies MySQLParameter objects from the MySQLParameterCollection to the specified array.

**Syntax: Visual Basic**

```
NotOverridable Public Sub CopyTo( _  
    ByVal array As Array, _  
    ByVal index As Integer _  
) _  
- Implements ICollection.CopyTo
```

**Syntax: C#**

```
public void CopyTo(  
    Arrayarray,  
    intindex  
);
```

**Parameters**

- `array`:
- `index`:

**Implements**

[ICollection.CopyTo](#)

**See Also**



[MySQLParameterCollection Class](#) , [MySQL.Data.MySqlClient Namespace](#)

#### 18.2.4.1.2.1.7.1.1.8. IndexOf Method

Gets the location of a [MySQLParameter](#) in the collection.

##### Overload List

Gets the location of a [MySQLParameter](#) in the collection.

- `public int IndexOf(object);`

Gets the location of the [MySQLParameter](#) in the collection with a specific parameter name.

- `public int IndexOf(string);`

##### See Also

[MySQLParameterCollection Class](#) , [MySQL.Data.MySqlClient Namespace](#)

#### 18.2.4.1.2.1.7.1.1.8.1. MySQLParameterCollection.IndexOf Method (Object)

Gets the location of a [MySQLParameter](#) in the collection.

##### Syntax: Visual Basic

```
NotOverridable Overloads Public Function IndexOf( _  
    ByVal value As Object _  
) As Integer _  
- Implements IList.IndexOf
```

##### Syntax: C#

```
public int IndexOf(  
    objectvalue  
);
```

##### Parameters

- `value`: The [MySQLParameter](#) object to locate.

##### Return Value

The zero-based location of the [MySQLParameter](#) in the collection.

##### Implements

`IList.IndexOf`

##### See Also

[MySQLParameterCollection Class](#) , [MySQL.Data.MySqlClient Namespace](#) , [MySQLParameterCollection.IndexOf Overload List](#)

#### 18.2.4.1.2.1.7.1.1.8.2. MySqlConnection.IndexOf Method (String)

Gets the location of the [MySqlParameter](#) in the collection with a specific parameter name.

##### Syntax: Visual Basic

```
NotOverridable Overloads Public Function IndexOf( _  
    ByVal parameterName As String _  
) As Integer _  
- Implements IDataParameterCollection.IndexOf
```

##### Syntax: C#

```
public int IndexOf(  
stringparameterName  
);
```

##### Parameters

- `parameterName`: The name of the [MySqlParameter](#) object to retrieve.

##### Return Value

The zero-based location of the [MySqlParameter](#) in the collection.

##### Implements

[IDataParameterCollection.IndexOf](#)

##### See Also

[MySqlParameterCollection Class](#) , [MySql.Data.MySqlClient Namespace](#) , [MySqlParameterCollection.IndexOf Overload List](#)

#### 18.2.4.1.2.1.7.1.1.9. MySqlConnection.Insert Method

Inserts a [MySqlParameter](#) into the collection at the specified index.

##### Syntax: Visual Basic

```
NotOverridable Public Sub Insert( _  
    ByVal index As Integer, _  
    ByVal value As Object _  
) _  
- Implements IList.Insert
```

##### Syntax: C#

```
public void Insert(  
intindex,  
objectvalue  
);
```

##### Parameters

- `index`:

- `value`:

**Implements**

`ICollection.Insert`

**See Also**

[MySQLParameterCollection Class](#) , [MySQL.Data.MySqlClient Namespace](#)

#### 18.2.4.1.2.1.7.1.1.10. MySQLParameterCollection.Remove Method

Removes the specified `MySQLParameter` from the collection.

**Syntax: Visual Basic**

```
NotOverridable Public Sub Remove( _  
    ByVal value As Object _  
) _  
- Implements ICollection.Remove
```

**Syntax: C#**

```
public void Remove(  
    object value  
);
```

**Parameters**

- `value`:

**Implements**

`ICollection.Remove`

**See Also**

[MySQLParameterCollection Class](#) , [MySQL.Data.MySqlClient Namespace](#)

#### 18.2.4.1.2.1.7.1.1.11. RemoveAt Method

Removes the specified `MySQLParameter` from the collection.

**Overload List**

Removes the specified `MySQLParameter` from the collection using a specific index.

- `public void RemoveAt(int);`

Removes the specified `MySQLParameter` from the collection using the parameter name.

- `public void RemoveAt(string);`

**See Also**

[MySQLParameterCollection Class](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.7.1.1.11.1. MySQLParameterCollection.RemoveAt Method (Int32)

Removes the specified [MySQLParameter](#) from the collection using a specific index.

**Syntax: Visual Basic**

```
NotOverridable Overloads Public Sub RemoveAt( _  
    ByVal index As Integer _  
) _  
- Implements IList.RemoveAt
```

**Syntax: C#**

```
public void RemoveAt(  
    int index  
);
```

**Parameters**

- **index**: The zero-based index of the parameter.

**Implements**

IList.RemoveAt

**See Also**

[MySQLParameterCollection Class](#) , [MySQL.Data.MySqlClient Namespace](#) , [MySQLParameterCollection.RemoveAt Overload List](#)

## 18.2.4.1.2.1.7.1.1.11.2. MySQLParameterCollection.RemoveAt Method (String)

Removes the specified [MySQLParameter](#) from the collection using the parameter name.

**Syntax: Visual Basic**

```
NotOverridable Overloads Public Sub RemoveAt( _  
    ByVal name As String _  
) _  
- Implements IDataParameterCollection.RemoveAt
```

**Syntax: C#**

```
public void RemoveAt(  
    string name  
);
```

**Parameters**

- **name**: The name of the [MySQLParameter](#) object to retrieve.

**Implements**

[IDataParameterCollection.RemoveAt](#)

**See Also**

[MySQLParameterCollection Class](#) , [MySQL.Data.MySqlClient Namespace](#) , [MySQLParameterCollection.RemoveAt Overload List](#)

## 18.2.4.1.2.1.8. Transaction Property

**Syntax: Visual Basic**

```
Public Property Transaction As MySQLTransaction
```

**Syntax: C#**

```
public MySQLTransaction Transaction {get; set;}
```

**See Also**

[MySQLCommand Class](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.9. UpdatedRowSource Property

**Syntax: Visual Basic**

```
NotOverridable Public Property UpdatedRowSource As UpdateRowSource _  
    Implements IDbCommand.UpdatedRowSource
```

**Syntax: C#**

```
public System.Data.UpdateRowSource UpdatedRowSource {get; set;}
```

**Implements**

[IDbCommand.UpdatedRowSource](#)

**See Also**

[MySQLCommand Class](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.10. MySQLCommand.Cancel Method

Attempts to cancel the execution of a [MySQLCommand](#). This operation is not supported.

**Syntax: Visual Basic**

```
NotOverridable Public Sub Cancel() _  
    Implements IDbCommand.Cancel
```

**Syntax: C#**

```
public void Cancel();
```

**Implements**

IDbCommand.Cancel

### Remarks

Cancelling an executing command is currently not supported on any version of MySQL.

### Exceptions

Exception Type	Condition
NotSupportedException	This operation is not supported.

### See Also

[MySQLCommand Class](#) , [MySql.Data.MySqlClient Namespace](#)

#### 18.2.4.1.2.1.11. MySqlCommand.CreateParameter Method

Creates a new instance of a [MySqlParameter](#) object.

### Syntax: Visual Basic

```
Public Function CreateParameter() As MySqlParameter
```

### Syntax: C#

```
public MySqlParameter CreateParameter();
```

### Return Value

A [MySqlParameter](#) object.

### Remarks

This method is a strongly-typed version of CreateParameter.

### See Also

[MySQLCommand Class](#) , [MySql.Data.MySqlClient Namespace](#)

#### 18.2.4.1.2.1.12. MySqlCommand.ExecuteNonQuery Method

### Syntax: Visual Basic

```
NotOverridable Public Function ExecuteNonQuery() As Integer _
    Implements IDbCommand.ExecuteNonQuery
```

### Syntax: C#

```
public int ExecuteNonQuery();
```

### Implements

IDbCommand.ExecuteNonQuery

### See Also

[MySqlCommand Class](#) ,  [MySql.Data.MySqlClient Namespace](#)

#### 18.2.4.1.2.1.13. ExecuteReader Method

##### Overload List

- [public MySqlDataReader ExecuteReader\(\);](#)
- [public MySqlDataReader ExecuteReader\(CommandBehavior\);](#)

##### See Also

[MySqlCommand Class](#) ,  [MySql.Data.MySqlClient Namespace](#)

#### 18.2.4.1.2.1.13.1. MySqlCommand.ExecuteReader Method ()

##### Syntax: Visual Basic

```
Overloads Public Function ExecuteReader() As MySqlDataReader
```

##### Syntax: C#

```
public MySqlDataReader ExecuteReader();
```

##### See Also

[MySqlCommand Class](#) ,  [MySql.Data.MySqlClient Namespace](#) ,  [MySqlCommand.ExecuteReader Overload List](#)

#### 18.2.4.1.2.1.13.1.1. MySqlDataReader Class

Provides a means of reading a forward-only stream of rows from a MySQL database. This class cannot be inherited.

For a list of all members of this type, see  [MySqlDataReader Members](#) .

##### Syntax: Visual Basic

```
NotInheritable Public Class MySqlDataReader_  
    Inherits MarshalByRefObject_  
    Implements IEnumerable, IDataReader, IDisposable, IDataRecord
```

##### Syntax: C#

```
public sealed class MySqlDataReader : MarshalByRefObject, IEnumerable, IDataReader, IDisposable, IDataRecord
```

##### Thread Safety

Public static (Shared in Visual Basic) members of this type are safe for multithreaded operations. Instance members are not guaranteed to be thread-safe.

##### Requirements

Namespace:  [MySql.Data.MySqlClient](#)

Assembly: MySql.Data (in MySql.Data.dll)

**See Also**

[MySqlDataReader Members](#) , [MySql.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.13.1.1.1. MySqlDataReader Members

[MySqlDataReader overview](#)

**Public Instance Properties**

<a href="#">Depth</a>	Gets a value indicating the depth of nesting for the current row. This method is not supported currently and always returns 0.
<a href="#">FieldCount</a>	Gets the number of columns in the current row.
<a href="#">HasRows</a>	Gets a value indicating whether the MySqlDataReader contains one or more rows.
<a href="#">IsClosed</a>	Gets a value indicating whether the data reader is closed.
<a href="#">Item</a>	Overloaded. Overloaded. Gets the value of a column in its native format. In C#, this property is the indexer for the MySqlDataReader class.
<a href="#">RecordsAffected</a>	Gets the number of rows changed, inserted, or deleted by execution of the SQL statement.

**Public Instance Methods**

<a href="#">Close</a>	Closes the MySqlDataReader object.
<a href="#">CreateObjRef</a> (inherited from MarshalByRefObject)	Creates an object that contains all the relevant information required to generate a proxy used to communicate with a remote object.
<a href="#">Equals</a> (inherited from Object)	Determines whether the specified Object is equal to the current Object.
<a href="#">GetBoolean</a>	Gets the value of the specified column as a Boolean.
<a href="#">GetByte</a>	Gets the value of the specified column as a byte.
<a href="#">GetBytes</a>	Reads a stream of bytes from the specified column offset into the buffer an array starting at the given buffer offset.
<a href="#">GetChar</a>	Gets the value of the specified column as a single character.
<a href="#">GetChars</a>	Reads a stream of characters from the specified column offset into the buffer as an array starting at the given buffer offset.
<a href="#">GetDataTypeName</a>	Gets the name of the source data type.
<a href="#">GetDateTime</a>	
<a href="#">GetDecimal</a>	
<a href="#">GetDouble</a>	
<a href="#">GetFieldType</a>	Gets the Type that is the data type of the object.
<a href="#">GetFloat</a>	



<a href="#">GetGuid</a>	
GetHashCode(inherited from Object)	Serves as a hash function for a particular type. GetHashCode is suitable for use in hashing algorithms and data structures like a hash table.
<a href="#">GetInt16</a>	
<a href="#">GetInt32</a>	
<a href="#">GetInt64</a>	
GetLifetimeService(inherited from MarshalByRefObject)	Retrieves the current lifetime service object that controls the lifetime policy for this instance.
<a href="#">GetMySqlDateTime</a>	
<a href="#">GetName</a>	Gets the name of the specified column.
<a href="#">GetOrdinal</a>	Gets the column ordinal, given the name of the column.
<a href="#">GetSchemaTable</a>	Returns a DataTable that describes the column metadata of the MySqlDataReader.
<a href="#">GetString</a>	
<a href="#">GetTimeSpan</a>	
GetType(inherited from Object)	Gets the Type of the current instance.
<a href="#">GetUInt16</a>	
<a href="#">GetUInt32</a>	
<a href="#">GetUInt64</a>	
<a href="#">GetValue</a>	Gets the value of the specified column in its native format.
<a href="#">GetValues</a>	Gets all attribute columns in the collection for the current row.
InitializeLifetimeService(inherited from MarshalByRefObject)	Obtains a lifetime service object to control the lifetime policy for this instance.
<a href="#">IsDBNull</a>	Gets a value indicating whether the column contains non-existent or missing values.
<a href="#">NextResult</a>	Advances the data reader to the next result, when reading the results of batch SQL statements.
<a href="#">Read</a>	Advances the MySqlDataReader to the next record.
ToString(inherited from Object)	Returns a String that represents the current Object.

**See Also**

[MySqlDataReader Class](#) , [MySql.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.13.1.1.1.1.1. Depth Property

Gets a value indicating the depth of nesting for the current row. This method is not supported currently and always returns 0.

**Syntax: Visual Basic**

```
NotOverridable Public ReadOnly Property Depth As Integer _
    Implements IDataReader.Depth
```

**Syntax: C#**

```
public int Depth {get;}
```

**Implements**

IDataReader.Depth

**See Also**

[MySQLDataReader Class](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.13.1.1.1.2. FieldCount Property

Gets the number of columns in the current row.

**Syntax: Visual Basic**

```
NotOverridable Public ReadOnly Property FieldCount As Integer _  
- Implements IDataRecord.FieldCount
```

**Syntax: C#**

```
public int FieldCount {get;}
```

**Implements**

IDataRecord.FieldCount

**See Also**

[MySQLDataReader Class](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.13.1.1.1.3. HasRows Property

Gets a value indicating whether the MySQLDataReader contains one or more rows.

**Syntax: Visual Basic**

```
Public ReadOnly Property HasRows As Boolean
```

**Syntax: C#**

```
public bool HasRows {get;}
```

**See Also**

[MySQLDataReader Class](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.13.1.1.1.4. IsClosed Property

Gets a value indicating whether the data reader is closed.

**Syntax: Visual Basic**

```
NotOverridable Public ReadOnly Property IsClosed As Boolean _  
- Implements IDataReader.IsClosed
```

**Syntax: C#**

```
public bool IsClosed {get;}
```

**Implements**

IDataReader.IsClosed

**See Also**

[MySQLDataReader Class](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.13.1.1.1.5. Item Property

Overloaded. Gets the value of a column in its native format. In C#, this property is the indexer for the MySQLDataReader class.

**Overload List**

Overloaded. Gets the value of a column in its native format. In C#, this property is the indexer for the MySQLDataReader class.

- [public object this\[int\] {get;}](#)

Gets the value of a column in its native format. In C#, this property is the indexer for the MySQLDataReader class.

- [public object this\[string\] {get;}](#)

**See Also**

[MySQLDataReader Class](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.13.1.1.1.5.1. Item Property (Int32)

Overloaded. Gets the value of a column in its native format. In C#, this property is the indexer for the MySQLDataReader class.

**Syntax: Visual Basic**

```
NotOverridable Overloads Public Default ReadOnly Property Item( _
    ByVal i As Integer _
) _
- Implements IDataRecord.Item As Object _
- Implements IDataRecord.Item
```

**Syntax: C#**

```
public object this[
    inti
] {get;}
```

**Implements**

IDataRecord.Item

**See Also**

[MySQLDataReader Class](#) , [MySQL.Data.MySqlClient Namespace](#) , [MySQLDataReader.Item Overload List](#)

## 18.2.4.1.2.1.13.1.1.1.5.2. Item Property (String)

Gets the value of a column in its native format. In C#, this property is the indexer for the MySQLDataReader class.

**Syntax: Visual Basic**

```
NotOverridable Overloads Public Default ReadOnly Property Item( _  
    ByVal name As String _  
) _  
- Implements IDataRecord.Item As Object _  
- Implements IDataRecord.Item
```

**Syntax: C#**

```
public object this[  
    stringname  
] {get;}
```

**Implements**

IDataRecord.Item

**See Also**

[MySQLDataReader Class](#) , [MySQL.Data.MySqlClient Namespace](#) , [MySQLDataReader.Item Overload List](#)

## 18.2.4.1.2.1.13.1.1.1.6. RecordsAffected Property

Gets the number of rows changed, inserted, or deleted by execution of the SQL statement.

**Syntax: Visual Basic**

```
NotOverridable Public ReadOnly Property RecordsAffected As Integer _  
- Implements IDataReader.RecordsAffected
```

**Syntax: C#**

```
public int RecordsAffected {get;}
```

**Implements**

IDataReader.RecordsAffected

**See Also**

[MySQLDataReader Class](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.13.1.1.1.7. MySQLDataReader.Close Method

Closes the MySQLDataReader object.

**Syntax: Visual Basic**

```
NotOverridable Public Sub Close() _  
    Implements IDataReader.Close
```

**Syntax: C#**

```
public void Close();
```

**Implements**

IDataReader.Close

**See Also**

[MySQLDataReader Class](#) , [MySQL.Data.MySqlClient Namespace](#)

#### 18.2.4.1.2.1.13.1.1.1.8. MySQLDataReader.GetBoolean Method

Gets the value of the specified column as a Boolean.

**Syntax: Visual Basic**

```
NotOverridable Public Function GetBoolean( _  
    ByVal i As Integer _  
) As Boolean _  
    Implements IDataRecord.GetBoolean
```

**Syntax: C#**

```
public bool GetBoolean(  
    inti  
);
```

**Parameters**

- `i`:

**Return Value****Implements**

IDataRecord.GetBoolean

**See Also**

[MySQLDataReader Class](#) , [MySQL.Data.MySqlClient Namespace](#)

#### 18.2.4.1.2.1.13.1.1.1.9. MySQLDataReader.GetByte Method

Gets the value of the specified column as a byte.

**Syntax: Visual Basic**

```
NotOverridable Public Function GetByte( _  
    ByVal i As Integer _  
) As Byte _  
    Implements IDataRecord.GetByte
```

**Syntax: C#**

```
public byte GetByte(  
    inti  
);
```

**Parameters**

- `i`:

**Return Value****Implements**

IDataRecord.GetByte

**See Also**

[MySQLDataReader Class](#) , [MySQL.Data.MySqlClient Namespace](#)

**18.2.4.1.2.1.13.1.1.1.10. MySQLDataReader.GetBytes Method**

Reads a stream of bytes from the specified column offset into the buffer an array starting at the given buffer offset.

**Syntax: Visual Basic**

```
NotOverridable Public Function GetBytes( _  
    ByVal i As Integer, _  
    ByVal dataIndex As Long, _  
    ByVal buffer As Byte(), _  
    ByVal bufferIndex As Integer, _  
    ByVal length As Integer _  
    ) As Long _  
- Implements IDataRecord.GetBytes
```

**Syntax: C#**

```
public long GetBytes(  
    inti,  
    longdataIndex,  
    byte[]buffer,  
    intbufferIndex,  
    intlength  
);
```

**Parameters**

- `i`: The zero-based column ordinal.
- `dataIndex`: The index within the field from which to begin the read operation.
- `buffer`: The buffer into which to read the stream of bytes.
- `bufferIndex`: The index for buffer to begin the read operation.
- `length`: The maximum length to copy into the buffer.

**Return Value**

The actual number of bytes read.

**Implements**

IDataRecord.GetBytes

**See Also**

[MySQLDataReader Class](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.13.1.1.11. MySQLDataReader.GetChar Method

Gets the value of the specified column as a single character.

**Syntax: Visual Basic**

```
NotOverridable Public Function GetChar( _
    ByVal i As Integer _
) As Char _
    Implements IDataRecord.GetChar
```

**Syntax: C#**

```
public char GetChar(
    inti
);
```

**Parameters**

- *i*:

**Return Value****Implements**

IDataRecord.GetChar

**See Also**

[MySQLDataReader Class](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.13.1.1.12. MySQLDataReader.GetChars Method

Reads a stream of characters from the specified column offset into the buffer as an array starting at the given buffer offset.

**Syntax: Visual Basic**

```
NotOverridable Public Function GetChars( _
    ByVal i As Integer, _
    ByVal fieldOffset As Long, _
    ByVal buffer As Char(), _
    ByVal bufferoffset As Integer, _
    ByVal length As Integer _
) As Long _
    Implements IDataRecord.GetChars
```

**Syntax: C#**

```
public long GetChars(  
    inti,  
    longfieldOffset,  
    char[]buffer,  
    intbufferoffset,  
    intlength  
);
```

**Parameters**

- `i`:
- `fieldOffset`:
- `buffer`:
- `bufferoffset`:
- `length`:

**Return Value****Implements**

IDataRecord.GetChars

**See Also**

[MySQLDataReader Class](#) , [MySql.Data.MySqlClient Namespace](#)

**18.2.4.1.2.1.13.1.1.13. MySQLDataReader.GetDataTypeName Method**

Gets the name of the source data type.

**Syntax: Visual Basic**

```
NotOverridable Public Function GetDataTypeName( _  
    ByVal i As Integer _  
) As String _  
— Implements IDataRecord.GetDataTypeName
```

**Syntax: C#**

```
public string GetDataTypeName(  
    inti  
);
```

**Parameters**

- `i`:

**Return Value****Implements**



IDataRecord.GetDataTypeName

**See Also**

[MySQLDataReader Class](#) , [MySQL.Data.MySqlClient Namespace](#)

#### 18.2.4.1.2.1.13.1.1.1.14. MySQLDataReader.GetDateTime Method

**Syntax: Visual Basic**

```
NotOverridable Public Function GetDateTime( _  
    ByVal index As Integer _  
    ) As Date _  
    Implements IDataRecord.GetDateTime
```

**Syntax: C#**

```
public DateTime GetDateTime(  
    int index  
);
```

**Implements**

IDataRecord.GetDateTime

**See Also**

[MySQLDataReader Class](#) , [MySQL.Data.MySqlClient Namespace](#)

#### 18.2.4.1.2.1.13.1.1.1.15. MySQLDataReader.GetDecimal Method

**Syntax: Visual Basic**

```
NotOverridable Public Function GetDecimal( _  
    ByVal index As Integer _  
    ) As Decimal _  
    Implements IDataRecord.GetDecimal
```

**Syntax: C#**

```
public decimal GetDecimal(  
    int index  
);
```

**Implements**

IDataRecord.GetDecimal

**See Also**

[MySQLDataReader Class](#) , [MySQL.Data.MySqlClient Namespace](#)

#### 18.2.4.1.2.1.13.1.1.1.16. MySQLDataReader.GetDouble Method

**Syntax: Visual Basic**

```
NotOverridable Public Function GetDouble( _  
    ByVal index As Integer _  
    ) As Double _  
    Implements IDataRecord.GetDouble
```

**Syntax: C#**

```
public double GetDouble(  
int index  
);
```

**Implements**

IDataRecord.GetDouble

**See Also**

[MySQLDataReader Class](#) , [MySQL.Data.MySqlClient Namespace](#)

**18.2.4.1.2.1.13.1.1.17. MySQLDataReader.GetFieldType Method**

Gets the Type that is the data type of the object.

**Syntax: Visual Basic**

```
NotOverridable Public Function GetFieldType( _  
    ByVal i As Integer _  
    ) As Type _  
    Implements IDataRecord.GetFieldType
```

**Syntax: C#**

```
public Type GetFieldType(  
int i  
);
```

**Parameters**

- **i**:

**Return Value****Implements**

IDataRecord.GetFieldType

**See Also**

[MySQLDataReader Class](#) , [MySQL.Data.MySqlClient Namespace](#)

**18.2.4.1.2.1.13.1.1.18. MySQLDataReader.GetFloat Method****Syntax: Visual Basic**

```
NotOverridable Public Function GetFloat( _  
    ByVal index As Integer _  
    ) As Single _  
    Implements IDataRecord.GetFloat
```

**Syntax: C#**

```
public float GetFloat(  
int index
```

```
);
```

**Implements**

IDataRecord.GetFloat

**See Also**[MySQLDataReader Class](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.13.1.1.1.19. MySQLDataReader.GetGuid Method

**Syntax: Visual Basic**

```
NotOverridable Public Function GetGuid( _  
    ByVal index As Integer _  
) As Guid _  
Implements IDataRecord.GetGuid
```

**Syntax: C#**

```
public Guid GetGuid(  
    int index  
);
```

**Implements**

IDataRecord.GetGuid

**See Also**[MySQLDataReader Class](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.13.1.1.1.20. MySQLDataReader.GetInt16 Method

**Syntax: Visual Basic**

```
NotOverridable Public Function GetInt16( _  
    ByVal index As Integer _  
) As Short _  
Implements IDataRecord.GetInt16
```

**Syntax: C#**

```
public short GetInt16(  
    int index  
);
```

**Implements**

IDataRecord.GetInt16

**See Also**[MySQLDataReader Class](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.13.1.1.1.21. MySQLDataReader.GetInt32 Method

**Syntax: Visual Basic**

```
NotOverridable Public Function GetInt32( _  
    ByVal index As Integer _  
) As Integer _  
- Implements IDataRecord.GetInt32
```

**Syntax: C#**

```
public int GetInt32(  
    int index  
);
```

**Implements**

IDataRecord.GetInt32

**See Also**

[MySQLDataReader Class](#) , [MySQL.Data.MySqlClient Namespace](#)

#### 18.2.4.1.2.1.13.1.1.1.22. MySQLDataReader.GetInt64 Method

**Syntax: Visual Basic**

```
NotOverridable Public Function GetInt64( _  
    ByVal index As Integer _  
) As Long _  
- Implements IDataRecord.GetInt64
```

**Syntax: C#**

```
public long GetInt64(  
    int index  
);
```

**Implements**

IDataRecord.GetInt64

**See Also**

[MySQLDataReader Class](#) , [MySQL.Data.MySqlClient Namespace](#)

#### 18.2.4.1.2.1.13.1.1.1.23. MySQLDataReader.GetMySqlDateTime Method

**Syntax: Visual Basic**

```
Public Function GetMySqlDateTime( _  
    ByVal index As Integer _  
) As MySqlDateTime
```

**Syntax: C#**

```
public MySqlDateTime GetMySqlDateTime(  
    int index  
);
```

**See Also**

[MySQLDataReader Class](#) , [MySQL.Data.MySqlClient Namespace](#)

#### 18.2.4.1.2.1.13.1.1.24. MySqlDataReader.GetName Method

Gets the name of the specified column.

##### Syntax: Visual Basic

```
NotOverridable Public Function GetName( _  
    ByVal i As Integer _  
) As String _  
- Implements IDataRecord.GetName
```

##### Syntax: C#

```
public string GetName(  
    inti  
);
```

##### Parameters

- `i`:

##### Return Value

##### Implements

IDataRecord.GetName

##### See Also

[MySqlDataReader Class](#) , [MySql.Data.MySqlClient Namespace](#)

#### 18.2.4.1.2.1.13.1.1.25. MySqlDataReader.GetOrdinal Method

Gets the column ordinal, given the name of the column.

##### Syntax: Visual Basic

```
NotOverridable Public Function GetOrdinal( _  
    ByVal name As String _  
) As Integer _  
- Implements IDataRecord.GetOrdinal
```

##### Syntax: C#

```
public int GetOrdinal(  
    stringname  
);
```

##### Parameters

- `name`:

##### Return Value

##### Implements

IDataRecord.GetOrdinal

**See Also**

[MySQLDataReader Class](#) , [MySQL.Data.MySqlClient Namespace](#)

18.2.4.1.2.1.13.1.1.1.26. MySQLDataReader.GetSchemaTable Method

Returns a DataTable that describes the column metadata of the MySQLDataReader.

**Syntax: Visual Basic**

```
NotOverridable Public Function GetSchemaTable() As DataTable _  
    Implements IDataReader.GetSchemaTable
```

**Syntax: C#**

```
public DataTable GetSchemaTable();
```

**Return Value**

**Implements**

IDataReader.GetSchemaTable

**See Also**

[MySQLDataReader Class](#) , [MySQL.Data.MySqlClient Namespace](#)

18.2.4.1.2.1.13.1.1.1.27. MySQLDataReader.GetString Method

**Syntax: Visual Basic**

```
NotOverridable Public Function GetString( _  
    ByVal index As Integer _  
) As String _  
    Implements IDataRecord.GetString
```

**Syntax: C#**

```
public string GetString(  
    int index  
);
```

**Implements**

IDataRecord.GetString

**See Also**

[MySQLDataReader Class](#) , [MySQL.Data.MySqlClient Namespace](#)

18.2.4.1.2.1.13.1.1.1.28. MySQLDataReader.GetTimeSpan Method

**Syntax: Visual Basic**

```
Public Function GetTimeSpan( _  
    ByVal index As Integer _  
) As TimeSpan
```

**Syntax: C#**

```
public TimeSpan GetTimeSpan(  
    int index  
);
```

**See Also**

[MySQLDataReader Class](#) , [MySql.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.13.1.1.1.29. MySQLDataReader.GetUInt16 Method

**Syntax: Visual Basic**

```
Public Function GetUInt16( _  
    ByVal index As Integer _  
) As UInt16
```

**Syntax: C#**

```
public ushort GetUInt16(  
    int index  
);
```

**See Also**

[MySQLDataReader Class](#) , [MySql.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.13.1.1.1.30. MySQLDataReader.GetUInt32 Method

**Syntax: Visual Basic**

```
Public Function GetUInt32( _  
    ByVal index As Integer _  
) As UInt32
```

**Syntax: C#**

```
public uint GetUInt32(  
    int index  
);
```

**See Also**

[MySQLDataReader Class](#) , [MySql.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.13.1.1.1.31. MySQLDataReader.GetUInt64 Method

**Syntax: Visual Basic**

```
Public Function GetUInt64( _  
    ByVal index As Integer _  
) As UInt64
```

**Syntax: C#**

```
public ulong GetUInt64(  
    int index  
);
```

**See Also**

[MySQLDataReader Class](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.13.1.1.1.32. MySQLDataReader.GetValue Method

Gets the value of the specified column in its native format.

**Syntax: Visual Basic**

```
NotOverridable Public Function GetValue( _  
    ByVal i As Integer _  
) As Object _  
- Implements IDataRecord.GetValue
```

**Syntax: C#**

```
public object GetValue(  
    inti  
);
```

**Parameters**

- `i`:

**Return Value****Implements**

IDataRecord.GetValue

**See Also**

[MySQLDataReader Class](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.13.1.1.1.33. MySQLDataReader.GetValues Method

Gets all attribute columns in the collection for the current row.

**Syntax: Visual Basic**

```
NotOverridable Public Function GetValues( _  
    ByVal values As Object() _  
) As Integer _  
- Implements IDataRecord.GetValues
```

**Syntax: C#**

```
public int GetValues(  
    object[] values  
);
```

**Parameters**

- `values`:



**Return Value****Implements**

IDataRecord.GetValues

**See Also**[MySQLDataReader Class](#) , [MySql.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.13.1.1.1.34. MySQLDataReader.IsDBNull Method

Gets a value indicating whether the column contains non-existent or missing values.

**Syntax: Visual Basic**

```
NotOverridable Public Function IsDBNull( _  
    ByVal i As Integer _  
) As Boolean _  
- Implements IDataRecord.IsDBNull
```

**Syntax: C#**

```
public bool IsDBNull(  
    inti  
);
```

**Parameters**

- *i*:

**Return Value****Implements**

IDataRecord.IsDBNull

**See Also**[MySQLDataReader Class](#) , [MySql.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.13.1.1.1.35. MySQLDataReader.NextResult Method

Advances the data reader to the next result, when reading the results of batch SQL statements.

**Syntax: Visual Basic**

```
NotOverridable Public Function NextResult() As Boolean _  
- Implements IDataReader.NextResult
```

**Syntax: C#**

```
public bool NextResult();
```

**Return Value**

**Implements**

IDataReader.NextResult

**See Also**[MySQLDataReader Class](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.13.1.1.1.36. MySQLDataReader.Read Method

Advances the MySQLDataReader to the next record.

**Syntax: Visual Basic**

```
NotOverridable Public Function Read() As Boolean _  
    Implements IDataReader.Read
```

**Syntax: C#**

```
public bool Read();
```

**Return Value****Implements**

IDataReader.Read

**See Also**[MySQLDataReader Class](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.2.1.13.2. MySqlCommand.ExecuteReader Method (CommandBehavior)

**Syntax: Visual Basic**

```
Overloads Public Function ExecuteReader( _  
    ByVal behavior As CommandBehavior _  
) As MySQLDataReader
```

**Syntax: C#**

```
public MySQLDataReader ExecuteReader(  
    CommandBehaviorbehavior  
);
```

**See Also**[MySqlCommand Class](#) , [MySQL.Data.MySqlClient Namespace](#) , [MySqlCommand.ExecuteReader Overload List](#)

## 18.2.4.1.2.1.14. MySqlCommand.ExecuteScalar Method

**Syntax: Visual Basic**

```
NotOverridable Public Function ExecuteScalar() As Object _  
    Implements IDbCommand.ExecuteScalar
```

**Syntax: C#**

```
public object ExecuteScalar();
```

**Implements**

IDbCommand.ExecuteScalar

**See Also**

[MySQLCommand Class](#) , [MySQL.Data.MySqlClient Namespace](#)

#### 18.2.4.1.2.1.15. MySQLCommand.Prepare Method

**Syntax: Visual Basic**

```
NotOverridable Public Sub Prepare() _  
    Implements IDbCommand.Prepare
```

**Syntax: C#**

```
public void Prepare();
```

**Implements**

IDbCommand.Prepare

**See Also**

[MySQLCommand Class](#) , [MySQL.Data.MySqlClient Namespace](#)

#### 18.2.4.1.3. MySQLCommandBuilder Class

For a list of all members of this type, see [MySQLCommandBuilder Members](#) .

**Syntax: Visual Basic**

```
NotInheritable Public Class MySQLCommandBuilder_  
    Inherits Component
```

**Syntax: C#**

```
public sealed class MySQLCommandBuilder : Component
```

**Thread Safety**

Public static (Shared in Visual Basic) members of this type are safe for multithreaded operations. Instance members are not guaranteed to be thread-safe.

**Requirements**

Namespace: [MySQL.Data.MySqlClient](#)

Assembly: MySQL.Data (in MySQL.Data.dll)

**See Also**

[MySQLCommandBuilder Members](#) , [MySQL.Data.MySqlClient Namespace](#)

**18.2.4.1.3.1. MySqlCommandBuilder Members**[MySqlCommandBuilder overview](#)**Public Static (Shared) Methods**

<a href="#">DeriveParameters</a>	Overloaded. Retrieves parameter information from the stored procedure specified in the MySqlCommand and populates the Parameters collection of the specified MySqlCommand object. This method is not currently supported since stored procedures are not available in MySQL.
----------------------------------	--

**Public Instance Constructors**

<a href="#">MySqlCommandBuilder</a>	Overloaded. Initializes a new instance of the MySqlCommandBuilder class.
-------------------------------------	--

**Public Instance Properties**

Container(inherited from Component)	Gets the IContainerthat contains the Component.
<a href="#">DataAdapter</a>	
<a href="#">QuotePrefix</a>	
<a href="#">QuoteSuffix</a>	
Site(inherited from Component)	Gets or sets the ISiteof the Component.

**Public Instance Methods**

CreateObjRef(inherited from MarshalByRefObject)	Creates an object that contains all the relevant information required to generate a proxy used to communicate with a remote object.
Dispose(inherited from Component)	Releases all resources used by the Component.
Equals(inherited from Object)	Determines whether the specified Objectis equal to the current Object.
<a href="#">GetDeleteCommand</a>	
GetHashCode(inherited from Object)	Serves as a hash function for a particular type. GetHashCodeis suitable for use in hashing algorithms and data structures like a hash table.
<a href="#">GetInsertCommand</a>	
GetLifetimeService(inherited from MarshalByRefObject)	Retrieves the current lifetime service object that controls the lifetime policy for this instance.
GetType(inherited from Object)	Gets the Typeof the current instance.
<a href="#">GetUpdateCommand</a>	
InitializeLifetimeService(inherited from MarshalByRefObject)	Obtains a lifetime service object to control the lifetime policy for this instance.
<a href="#">RefreshSchema</a>	
ToString(inherited from Component)	Returns a Stringcontaining the name of the Component, if any. This method should not be overrid-

	den.
--	------

### Public Instance Events

Disposed(inherited from Component)	Adds an event handler to listen to the Dispose- devent on the component.
------------------------------------	---

### See Also

[MySQLCommandBuilder Class](#) , [MySQL.Data.MySqlClient Namespace](#)

#### 18.2.4.1.3.1.1. DeriveParameters Method

Retrieves parameter information from the stored procedure specified in the MySqlCommand and populates the Parameters collection of the specified MySqlCommand object. This method is not currently supported since stored procedures are not available in MySQL.

### Overload List

Retrieves parameter information from the stored procedure specified in the MySqlCommand and populates the Parameters collection of the specified MySqlCommand object. This method is not currently supported since stored procedures are not available in MySQL.

- [public static void DeriveParameters\(MySqlCommand\);](#)
- [public static void DeriveParameters\(MySqlCommand,bool\);](#)

### See Also

[MySQLCommandBuilder Class](#) , [MySQL.Data.MySqlClient Namespace](#)

#### 18.2.4.1.3.1.1.1. MySqlCommandBuilder.DeriveParameters Method (MySqlCommand)

Retrieves parameter information from the stored procedure specified in the MySqlCommand and populates the Parameters collection of the specified MySqlCommand object. This method is not currently supported since stored procedures are not available in MySQL.

### Syntax: Visual Basic

```
Overloads Public Shared Sub DeriveParameters( _
    ByVal command As MySqlCommand _
)
```

### Syntax: C#

```
public static void DeriveParameters(
    MySqlCommandcommand
);
```

### Parameters

- [command](#): The MySqlCommand referencing the stored procedure from which the parameter information is to be derived. The derived parameters are added to the Parameters collection of the MySqlCommand.

**Exceptions**

Exception Type	Condition
InvalidOperationException	The command text is not a valid stored procedure name.

**See Also**

[MySQLCommandBuilder Class](#) , [MySQL.Data.MySqlClient Namespace](#) , [MySQLCommandBuilder.DeriveParameters Overload List](#)

## 18.2.4.1.3.1.1.2. MySQLCommandBuilder.DeriveParameters Method (MySQLCommand, Boolean)

**Syntax: Visual Basic**

```
Overloads Public Shared Sub DeriveParameters( _
    ByVal command As MySQLCommand, _
    ByVal useProc As Boolean _
)
```

**Syntax: C#**

```
public static void DeriveParameters(
    MySQLCommand command,
    bool useProc
);
```

**See Also**

[MySQLCommandBuilder Class](#) , [MySQL.Data.MySqlClient Namespace](#) , [MySQLCommandBuilder.DeriveParameters Overload List](#)

## 18.2.4.1.3.1.2. MySQLCommandBuilder Constructor

Initializes a new instance of the [MySQLCommandBuilder](#) class.

**Overload List**

Initializes a new instance of the [MySQLCommandBuilder](#) class.

- [public MySQLCommandBuilder\(\);](#)
- [public MySQLCommandBuilder\(MySQLDataAdapter\);](#)
- [public MySQLCommandBuilder\(MySQLDataAdapter,bool\);](#)
- [public MySQLCommandBuilder\(bool\);](#)

**See Also**

[MySQLCommandBuilder Class](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.3.1.2.1. MySQLCommandBuilder Constructor ()

Initializes a new instance of the [MySQLCommandBuilder](#) class.

**Syntax: Visual Basic**

```
Overloads Public Sub New()
```

**Syntax: C#**

```
public MySqlCommandBuilder();
```

**See Also**

[MySqlCommandBuilder Class](#) , [MySql.Data.MySqlClient Namespace](#) , [MySqlCommandBuilder Constructor Overload List](#)

#### 18.2.4.1.3.1.2.2. MySqlCommandBuilder Constructor (MySqlDataAdapter)

**Syntax: Visual Basic**

```
Overloads Public Sub New( _  
    ByVal adapter As MySqlDataAdapter _  
)
```

**Syntax: C#**

```
public MySqlCommandBuilder(  
    MySqlDataAdapteradapter  
);
```

**See Also**

[MySqlCommandBuilder Class](#) , [MySql.Data.MySqlClient Namespace](#) , [MySqlCommandBuilder Constructor Overload List](#)

#### 18.2.4.1.3.1.2.2.1. MySqlDataAdapter Class

For a list of all members of this type, see [MySqlDataAdapter Members](#) .

**Syntax: Visual Basic**

```
NotInheritable Public Class MySqlDataAdapter_  
    Inherits DbDataAdapter
```

**Syntax: C#**

```
public sealed class MySqlDataAdapter : DbDataAdapter
```

**Thread Safety**

Public static (Shared in Visual Basic) members of this type are safe for multithreaded operations. Instance members are not guaranteed to be thread-safe.

**Requirements**

Namespace: [MySql.Data.MySqlClient](#)

Assembly: MySql.Data (in MySql.Data.dll)

**See Also**

[MySqlDataAdapter Members](#) , [MySql.Data.MySqlClient Namespace](#)

## 18.2.4.1.3.1.2.2.1.1. MySqlDataAdapter Members

[MySqlDataAdapter overview](#)**Public Instance Constructors**

<a href="#">MySqlDataAdapter</a>	Overloaded. Initializes a new instance of the <code>MySqlDataAdapter</code> class.
----------------------------------	--

**Public Instance Properties**

<code>AcceptChangesDuringFill</code> (inherited from <code>DataAdapter</code> )	Gets or sets a value indicating whether <code>AcceptChanges</code> is called on a <code>DataRow</code> after it is added to the <code>DataTable</code> during any of the Fill operations.
<code>AcceptChangesDuringUpdate</code> (inherited from <code>DataAdapter</code> )	Gets or sets whether <code>AcceptChanges</code> is called during a <code>Update</code> .
<code>Container</code> (inherited from <code>Component</code> )	Gets the <code>IContainer</code> that contains the <code>Component</code> .
<code>ContinueUpdateOnError</code> (inherited from <code>DataAdapter</code> )	Gets or sets a value that specifies whether to generate an exception when an error is encountered during a row update.
<a href="#">DeleteCommand</a>	Overloaded.
<code>FillLoadOption</code> (inherited from <code>DataAdapter</code> )	Gets or sets the <code>LoadOption</code> that determines how the adapter fills the <code>DataTable</code> from the <code>DbDataReader</code> .
<a href="#">InsertCommand</a>	Overloaded.
<code>MissingMappingAction</code> (inherited from <code>DataAdapter</code> )	Determines the action to take when incoming data does not have a matching table or column.
<code>MissingSchemaAction</code> (inherited from <code>DataAdapter</code> )	Determines the action to take when existing <code>DataSets</code> schema does not match incoming data.
<code>ReturnProviderSpecificTypes</code> (inherited from <code>DataAdapter</code> )	Gets or sets whether the <code>Fill</code> method should return provider-specific values or common CLS-compliant values.
<a href="#">SelectCommand</a>	Overloaded.
<code>Site</code> (inherited from <code>Component</code> )	Gets or sets the <code>ISite</code> of the <code>Component</code> .
<code>TableMappings</code> (inherited from <code>DataAdapter</code> )	Gets a collection that provides the master mapping between a source table and a <code>DataTable</code> .
<code>UpdateBatchSize</code> (inherited from <code>DbDataAdapter</code> )	Gets or sets a value that enables or disables batch processing support, and specifies the number of commands that can be executed in a batch.
<a href="#">UpdateCommand</a>	Overloaded.

**Public Instance Methods**

<code>CreateObjRef</code> (inherited from <code>MarshalByRefObject</code> )	Creates an object that contains all the relevant information required to generate a proxy used to communicate with a remote object.
<code>Dispose</code> (inherited from <code>Component</code> )	Releases all resources used by the <code>Component</code> .
<code>Equals</code> (inherited from <code>Object</code> )	Determines whether the specified <code>Object</code> is equal to the current <code>Object</code> .



Fill(inherited from DbDataAdapter)	Overloaded. Adds or refreshes rows in the DataSet to match those in the data source using the DataSetname, and creates a DataTable named "Table."
FillSchema(inherited from DbDataAdapter)	Overloaded. Configures the schema of the specified DataTable based on the specified SchemaType.
GetFillParameters(inherited from DbDataAdapter)	Gets the parameters set by the user when executing an SQL <code>SELECT</code> statement.
GetHashCode(inherited from Object)	Serves as a hash function for a particular type. GetHashCode is suitable for use in hashing algorithms and data structures like a hash table.
GetLifetimeService(inherited from MarshalByRefObject)	Retrieves the current lifetime service object that controls the lifetime policy for this instance.
GetType(inherited from Object)	Gets the Type of the current instance.
InitializeLifetimeService(inherited from MarshalByRefObject)	Obtains a lifetime service object to control the lifetime policy for this instance.
ResetFillLoadOption(inherited from DataAdapter)	Resets FillLoadOption to its default state and causes Fill to honor AcceptChangesDuringFill.
ShouldSerializeAcceptChangesDuringFill(inherited from DataAdapter)	Determines whether the AcceptChangesDuringFill property should be persisted.
ShouldSerializeFillLoadOption(inherited from DataAdapter)	Determines whether the FillLoadOption property should be persisted.
ToString(inherited from Component)	Returns a String containing the name of the Component, if any. This method should not be overridden.
Update(inherited from DbDataAdapter)	Overloaded. Calls the respective INSERT, UPDATE, or DELETE statements for each inserted, updated, or deleted row in the specified DataSet.

### Public Instance Events

Disposed(inherited from Component)	Adds an event handler to listen to the Dispose event on the component.
FillError(inherited from DataAdapter)	Returned when an error occurs during a fill operation.
<a href="#">RowUpdated</a>	Occurs during Update after a command is executed against the data source. The attempt to update is made, so the event fires.
<a href="#">RowUpdating</a>	Occurs during Update before a command is executed against the data source. The attempt to update is made, so the event fires.

### Protected Internal Instance Properties

FillCommandBehavior(inherited from DbDataAdapter)	Gets or sets the behavior of the command used to fill the data adapter.
---	---

**See Also**

[MySQLDataAdapter Class](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.3.1.2.2.1.1.1. MySQLDataAdapter Constructor

Initializes a new instance of the [MySQLDataAdapter](#) class.

**Overload List**

Initializes a new instance of the [MySQLDataAdapter](#) class.

- [public MySQLDataAdapter\(\);](#)
- [public MySQLDataAdapter\(MySqlCommand\);](#)
- [public MySQLDataAdapter\(string,MySQLConnection\);](#)
- [public MySQLDataAdapter\(string,string\);](#)

**See Also**

[MySQLDataAdapter Class](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.3.1.2.2.1.1.1.1. MySQLDataAdapter Constructor ()

Initializes a new instance of the [MySQLDataAdapter](#) class.

**Syntax: Visual Basic**

```
Overloads Public Sub New()
```

**Syntax: C#**

```
public MySQLDataAdapter();
```

**See Also**

[MySQLDataAdapter Class](#) , [MySQL.Data.MySqlClient Namespace](#) , [MySQLDataAdapter Constructor Overload List](#)

## 18.2.4.1.3.1.2.2.1.1.1.2. MySQLDataAdapter Constructor (MySqlCommand)

**Syntax: Visual Basic**

```
Overloads Public Sub New( _  
    ByVal selectCommand As MySqlCommand _  
)
```

**Syntax: C#**

```
public MySQLDataAdapter(  
    MySqlCommandselectCommand  
);
```

**See Also**

[MySQLDataAdapter Class](#) , [MySQL.Data.MySqlClient Namespace](#) , [MySQLDataAdapter Constructor Overload List](#)

#### 18.2.4.1.3.1.2.2.1.1.3. MySQLDataAdapter Constructor (String, MySqlConnection)

##### Syntax: Visual Basic

```
Overloads Public Sub New( _  
    ByVal selectCommandText As String, _  
    ByVal connection As MySqlConnection _  
)
```

##### Syntax: C#

```
public MySQLDataAdapter(  
    stringselectCommandText,  
    MySqlConnectionconnection  
) ;
```

##### See Also

[MySQLDataAdapter Class](#) , [MySQL.Data.MySqlClient Namespace](#) , [MySQLDataAdapter Constructor Overload List](#)

#### 18.2.4.1.3.1.2.2.1.1.4. MySQLDataAdapter Constructor (String, String)

##### Syntax: Visual Basic

```
Overloads Public Sub New( _  
    ByVal selectCommandText As String, _  
    ByVal selectConnString As String _  
)
```

##### Syntax: C#

```
public MySQLDataAdapter(  
    stringselectCommandText,  
    stringselectConnString  
) ;
```

##### See Also

[MySQLDataAdapter Class](#) , [MySQL.Data.MySqlClient Namespace](#) , [MySQLDataAdapter Constructor Overload List](#)

#### 18.2.4.1.3.1.2.2.1.1.2. DeleteCommand Property

##### Syntax: Visual Basic

```
Overloads Public Property DeleteCommand As MySqlCommand
```

##### Syntax: C#

```
new public MySqlCommand DeleteCommand {get; set;}
```

##### See Also

[MySQLDataAdapter Class](#) , [MySQL.Data.MySqlClient Namespace](#)

#### 18.2.4.1.3.1.2.2.1.1.3. InsertCommand Property

**Syntax: Visual Basic**

```
Overloads Public Property InsertCommand As MySqlCommand
```

**Syntax: C#**

```
new public MySqlCommand InsertCommand {get; set;}
```

**See Also**

[MySqlDataAdapter Class](#) , [MySql.Data.MySqlClient Namespace](#)

## 18.2.4.1.3.1.2.2.1.1.4. SelectCommand Property

**Syntax: Visual Basic**

```
Overloads Public Property SelectCommand As MySqlCommand
```

**Syntax: C#**

```
new public MySqlCommand SelectCommand {get; set;}
```

**See Also**

[MySqlDataAdapter Class](#) , [MySql.Data.MySqlClient Namespace](#)

## 18.2.4.1.3.1.2.2.1.1.5. UpdateCommand Property

**Syntax: Visual Basic**

```
Overloads Public Property UpdateCommand As MySqlCommand
```

**Syntax: C#**

```
new public MySqlCommand UpdateCommand {get; set;}
```

**See Also**

[MySqlDataAdapter Class](#) , [MySql.Data.MySqlClient Namespace](#)

## 18.2.4.1.3.1.2.2.1.1.6. MySqlDataAdapter.RowUpdated Event

Occurs during Update after a command is executed against the data source. The attempt to update is made, so the event fires.

**Syntax: Visual Basic**

```
Public Event RowUpdated As MySqlRowUpdatedEventHandler
```

**Syntax: C#**

```
public event MySqlRowUpdatedEventHandler RowUpdated;
```

**Event Data**

The event handler receives an argument of type [MySqlRowUpdatedEventArgs](#) containing data related to

this event. The following `MySqlRowUpdatedEventArgs` properties provide information specific to this event.

Property	Description
<a href="#">Command</a>	Gets or sets the <code>MySqlCommand</code> executed when <code>Update</code> is called.
Errors	Gets any errors generated by the .NET Framework data provider when the <code>Command</code> was executed.
RecordsAffected	Gets the number of rows changed, inserted, or deleted by execution of the SQL statement.
Row	Gets the <code>DataRow</code> sent through an <code>Update</code> .
RowCount	Gets the number of rows processed in a batch of updated records.
StatementType	Gets the type of SQL statement executed.
Status	Gets the <code>UpdateStatus</code> of the <code>Command</code> property.
TableMapping	Gets the <code>DataTableMapping</code> sent through an <code>Update</code> .

#### See Also

[MySqlDataAdapter Class](#) , [MySql.Data.MySqlClient Namespace](#)

#### 18.2.4.1.3.1.2.2.1.1.6.1. MySqlRowUpdatedEventHandler Delegate

Represents the method that will handle the `RowUpdated` event of a [MySqlDataAdapter](#) .

#### Syntax: Visual Basic

```
Public Delegate Sub MySqlRowUpdatedEventHandler( _
    ByVal sender As Object, _
    ByVal e As MySqlRowUpdatedEventArgs _
)
```

#### Syntax: C#

```
public delegate void MySqlRowUpdatedEventHandler(
    object sender,
    MySqlRowUpdatedEventArgs
);
```

#### Requirements

Namespace: [MySql.Data.MySqlClient](#)

Assembly: `MySql.Data` (in `MySql.Data.dll`)

#### See Also

[MySql.Data.MySqlClient Namespace](#)

#### 18.2.4.1.3.1.2.2.1.1.6.1.1. MySqlRowUpdatedEventArgs Class

Provides data for the `RowUpdated` event. This class cannot be inherited.

For a list of all members of this type, see [MySqlRowUpdatedEventArgs Members](#) .

**Syntax: Visual Basic**

```
NotInheritable Public Class MySqlRowUpdatedEventArgs_
    Inherits RowUpdatedEventArgs
```

**Syntax: C#**

```
public sealed class MySqlRowUpdatedEventArgs : RowUpdatedEventArgs
```

**Thread Safety**

Public static (Shared in Visual Basic) members of this type are safe for multithreaded operations. Instance members are not guaranteed to be thread-safe.

**Requirements**

Namespace: [MySql.Data.MySqlClient](#)

Assembly: MySql.Data (in MySql.Data.dll)

**See Also**

[MySqlRowUpdatedEventArgs Members](#) , [MySql.Data.MySqlClient Namespace](#)

## 18.2.4.1.3.1.2.2.1.1.6.1.1.1. MySqlRowUpdatedEventArgs Members

[MySqlRowUpdatedEventArgs overview](#)

**Public Instance Constructors**

<a href="#">MySqlRowUpdatedEventArgs Constructor</a>	Initializes a new instance of the MySqlRowUpdatedEventArgs class.
--	---

**Public Instance Properties**

<a href="#">Command</a>	Overloaded. Gets or sets the MySqlCommand executed when Update is called.
Errors(inherited from RowUpdatedEventArgs)	Gets any errors generated by the .NET Framework data provider when the Command was executed.
RecordsAffected(inherited from RowUpdatedEventArgs)	Gets the number of rows changed, inserted, or deleted by execution of the SQL statement.
Row(inherited from RowUpdatedEventArgs)	Gets the DataRow sent through an Update.
RowCount(inherited from RowUpdatedEventArgs)	Gets the number of rows processed in a batch of updated records.
StatementType(inherited from RowUpdatedEventArgs)	Gets the type of SQL statement executed.
Status(inherited from RowUpdatedEventArgs)	Gets the UpdateStatus of the Command property.
TableMapping(inherited from RowUpdatedEventArgs)	Gets the DataTableMapping sent through an Update.

**Public Instance Methods**

CopyToRows(inherited from RowUpdatedEventArgs)	Overloaded. Copies references to the modified
--	---

Args)	rows into the provided array.
Equals(inherited from Object)	Determines whether the specified Object is equal to the current Object.
GetHashCode(inherited from Object)	Serves as a hash function for a particular type. GetHashCode is suitable for use in hashing algorithms and data structures like a hash table.
GetType(inherited from Object)	Gets the Type of the current instance.
ToString(inherited from Object)	Returns a String that represents the current Object.

**See Also**

[MySQLRowUpdatedEventArgs Class](#) , [MySQL.Data.MySqlCommand Namespace](#)

#### 18.2.4.1.3.1.2.2.1.1.6.1.1.1.1. MySQLRowUpdatedEventArgs Constructor

Initializes a new instance of the MySQLRowUpdatedEventArgs class.

**Syntax: Visual Basic**

```
Public Sub New( _
    ByVal row As DataRow, _
    ByVal command As IDbCommand, _
    ByVal statementType As StatementType, _
    ByVal tableMapping As DataTableMapping _
)
```

**Syntax: C#**

```
public MySQLRowUpdatedEventArgs(
    DataRow row,
    IDbCommand command,
    StatementType statementType,
    DataTableMapping tableMapping
);
```

**Parameters**

- **row**: The DataRow sent through an Update.
- **command**: The IDbCommand executed when Update is called.
- **statementType**: One of the StatementType values that specifies the type of query executed.
- **tableMapping**: The DataTableMapping sent through an Update.

**See Also**

[MySQLRowUpdatedEventArgs Class](#) , [MySQL.Data.MySqlCommand Namespace](#)

#### 18.2.4.1.3.1.2.2.1.1.6.1.1.1.2. Command Property

Gets or sets the MySqlCommand executed when Update is called.

**Syntax: Visual Basic**

```
Overloads Public ReadOnly Property Command As MySqlCommand
```

**Syntax: C#**

```
new public MySqlCommand Command {get;}
```

**See Also**

[MySqlRowUpdatedEventArgs Class](#) , [MySql.Data.MySqlClient Namespace](#)

## 18.2.4.1.3.1.2.2.1.1.7. MySqlDataAdapter.RowUpdating Event

Occurs during Update before a command is executed against the data source. The attempt to update is made, so the event fires.

**Syntax: Visual Basic**

```
Public Event RowUpdating As MySqlRowUpdatingEventHandler
```

**Syntax: C#**

```
public event MySqlRowUpdatingEventHandler RowUpdating;
```

**Event Data**

The event handler receives an argument of type [MySqlRowUpdatingEventArgs](#) containing data related to this event. The following [MySqlRowUpdatingEventArgs](#) properties provide information specific to this event.

Property	Description
<a href="#">Command</a>	Gets or sets the MySqlCommand to execute when performing the Update.
Errors	Gets any errors generated by the .NET Framework data provider when the Command executes.
Row	Gets the DataRow that will be sent to the server as part of an insert, update, or delete operation.
StatementType	Gets the type of SQL statement to execute.
Status	Gets or sets the UpdateStatus of the Command-property.
TableMapping	Gets the DataTableMapping to send through the Update.

**See Also**

[MySqlDataAdapter Class](#) , [MySql.Data.MySqlClient Namespace](#)

## 18.2.4.1.3.1.2.2.1.1.7.1. MySqlRowUpdatingEventHandler Delegate

Represents the method that will handle the RowUpdating event of a [MySqlDataAdapter](#) .

**Syntax: Visual Basic**

```
Public Delegate Sub MySqlRowUpdatingEventHandler( _
    ByVal sender As Object, _
    ByVal e As MySqlRowUpdatingEventArgs _
)
```



**Syntax: C#**

```
public delegate void MySqlRowUpdatingEventHandler(
    object sender,
    MySqlRowUpdatingEventArgs
);
```

**Requirements**

Namespace: [MySql.Data.MySqlClient](#)

Assembly: MySql.Data (in MySql.Data.dll)

**See Also**

[MySql.Data.MySqlClient Namespace](#)

## 18.2.4.1.3.1.2.2.1.1.7.1.1. MySqlRowUpdatingEventArgs Class

Provides data for the RowUpdating event. This class cannot be inherited.

For a list of all members of this type, see [MySqlRowUpdatingEventArgs Members](#) .

**Syntax: Visual Basic**

```
NotInheritable Public Class MySqlRowUpdatingEventArgs_
    Inherits RowUpdatingEventArgs
```

**Syntax: C#**

```
public sealed class MySqlRowUpdatingEventArgs : RowUpdatingEventArgs
```

**Thread Safety**

Public static (Shared in Visual Basic) members of this type are safe for multithreaded operations. Instance members are not guaranteed to be thread-safe.

**Requirements**

Namespace: [MySql.Data.MySqlClient](#)

Assembly: MySql.Data (in MySql.Data.dll)

**See Also**

[MySqlRowUpdatingEventArgs Members](#) , [MySql.Data.MySqlClient Namespace](#)

## 18.2.4.1.3.1.2.2.1.1.7.1.1.1. MySqlRowUpdatingEventArgs Members

[MySqlRowUpdatingEventArgs overview](#)

**Public Instance Constructors**

<a href="#">MySqlRowUpdatingEventArgs Constructor</a>	Initializes a new instance of the MySqlRowUpdatingEventArgs class.
---	--

**Public Instance Properties**

<a href="#">Command</a>	Overloaded. Gets or sets the MySqlCommand to execute when performing the Update.
Errors(inherited from RowUpdatingEventArgs)	Gets any errors generated by the .NET Framework data provider when the Command executes.
Row(inherited from RowUpdatingEventArgs)	Gets the DataRow that will be sent to the server as part of an insert, update, or delete operation.
StatementType(inherited from RowUpdatingEventArgs)	Gets the type of SQL statement to execute.
Status(inherited from RowUpdatingEventArgs)	Gets or sets the UpdateStatus of the Command property.
TableMapping(inherited from RowUpdatingEventArgs)	Gets the DataTableMapping to send through the Update.

### Public Instance Methods

Equals(inherited from Object)	Determines whether the specified Object is equal to the current Object.
GetHashCode(inherited from Object)	Serves as a hash function for a particular type. GetHashCode is suitable for use in hashing algorithms and data structures like a hash table.
GetType(inherited from Object)	Gets the Type of the current instance.
ToString(inherited from Object)	Returns a String that represents the current Object.

### See Also

[MySQLRowUpdatingEventArgs Class](#) , [MySQL.Data.MySqlClient Namespace](#)

#### 18.2.4.1.3.1.2.2.1.1.7.1.1.1.1. MySQLRowUpdatingEventArgs Constructor

Initializes a new instance of the MySQLRowUpdatingEventArgs class.

#### Syntax: Visual Basic

```
Public Sub New( _
    ByVal row As DataRow, _
    ByVal command As IDbCommand, _
    ByVal statementType As StatementType, _
    ByVal tableMapping As DataTableMapping _
)
```

#### Syntax: C#

```
public MySQLRowUpdatingEventArgs(
    DataRow row,
    IDbCommand command,
    StatementType statementType,
    DataTableMapping tableMapping
);
```

### Parameters

- row**: The DataRow to Update.
- command**: The IDbCommand to execute during Update.

- `statementType`: One of the `StatementType` values that specifies the type of query executed.
- `tableMapping`: The `DataTableMapping` sent through an `Update`.

**See Also**

[MySQLRowUpdatingEventArgs Class](#) , [MySQL.Data.MySqlCommand Namespace](#)

## 18.2.4.1.3.1.2.2.1.1.7.1.1.1.2. Command Property

Gets or sets the `SqlCommand` to execute when performing the `Update`.

**Syntax: Visual Basic**

```
Overloads Public Property Command As MySqlCommand
```

**Syntax: C#**

```
new public MySqlCommand Command {get; set;}
```

**See Also**

[MySQLRowUpdatingEventArgs Class](#) , [MySQL.Data.MySqlCommand Namespace](#)

## 18.2.4.1.3.1.2.3. MySqlCommandBuilder Constructor (MySqlDataAdapter, Boolean)

**Syntax: Visual Basic**

```
Overloads Public Sub New( _  
    ByVal adapter As MySqlDataAdapter, _  
    ByVal lastOneWins As Boolean _  
)
```

**Syntax: C#**

```
public MySqlCommandBuilder(  
    MySqlDataAdapter adapter,  
    boollastOneWins  
);
```

**See Also**

[MySQLCommandBuilder Class](#) , [MySQL.Data.MySqlCommand Namespace](#) , [MySQLCommandBuilder Constructor Overload List](#)

## 18.2.4.1.3.1.2.4. MySqlCommandBuilder Constructor (Boolean)

**Syntax: Visual Basic**

```
Overloads Public Sub New( _  
    ByVal lastOneWins As Boolean _  
)
```

**Syntax: C#**

```
public MySqlCommandBuilder(  
    boollastOneWins  
);
```

**See Also**

[MySQLCommandBuilder Class](#) , [MySQL.Data.MySqlClient Namespace](#) , [MySQLCommandBuilder Constructor Overload List](#)

## 18.2.4.1.3.1.3. DataAdapter Property

**Syntax: Visual Basic**

```
Public Property DataAdapter As MySqlDataAdapter
```

**Syntax: C#**

```
public MySqlDataAdapter DataAdapter {get; set;}
```

**See Also**

[MySQLCommandBuilder Class](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.3.1.4. QuotePrefix Property

**Syntax: Visual Basic**

```
Public Property QuotePrefix As String
```

**Syntax: C#**

```
public string QuotePrefix {get; set;}
```

**See Also**

[MySQLCommandBuilder Class](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.3.1.5. QuoteSuffix Property

**Syntax: Visual Basic**

```
Public Property QuoteSuffix As String
```

**Syntax: C#**

```
public string QuoteSuffix {get; set;}
```

**See Also**

[MySQLCommandBuilder Class](#) , [MySQL.Data.MySqlClient Namespace](#)

## 18.2.4.1.3.1.6. MySQLCommandBuilder.GetDeleteCommand Method

**Syntax: Visual Basic**

```
Public Function GetDeleteCommand() As MySqlCommand
```

**Syntax: C#**

```
public MySqlCommand GetDeleteCommand();
```

**See Also**

[MySqlCommandBuilder Class](#) ,  [MySql.Data.MySqlClient Namespace](#)

## 18.2.4.1.3.1.7. MySqlCommandBuilder.GetInsertCommand Method

**Syntax: Visual Basic**

```
Public Function GetInsertCommand() As MySqlCommand
```

**Syntax: C#**

```
public MySqlCommand GetInsertCommand();
```

**See Also**

[MySqlCommandBuilder Class](#) ,  [MySql.Data.MySqlClient Namespace](#)

## 18.2.4.1.3.1.8. MySqlCommandBuilder.GetUpdateCommand Method

**Syntax: Visual Basic**

```
Public Function GetUpdateCommand() As MySqlCommand
```

**Syntax: C#**

```
public MySqlCommand GetUpdateCommand();
```

**See Also**

[MySqlCommandBuilder Class](#) ,  [MySql.Data.MySqlClient Namespace](#)

## 18.2.4.1.3.1.9. MySqlCommandBuilder.RefreshSchema Method

**Syntax: Visual Basic**

```
Public Sub RefreshSchema()
```

**Syntax: C#**

```
public void RefreshSchema();
```

**See Also**

[MySqlCommandBuilder Class](#) ,  [MySql.Data.MySqlClient Namespace](#)

**18.2.4.1.4. MySqlConnectionException Class**

The exception that is thrown when MySQL returns an error. This class cannot be inherited.

For a list of all members of this type, see  [MySqlConnectionException Members](#) .

**Syntax: Visual Basic**

```
NotInheritable Public Class MySqlConnectionException_  
    Inherits SystemException
```

**Syntax: C#**

```
public sealed class MySqlConnection : SystemException
```

**Thread Safety**

Public static (Shared in Visual Basic) members of this type are safe for multithreaded operations. Instance members are not guaranteed to be thread-safe.

**Requirements**

Namespace: [MySql.Data.MySqlClient](#)

Assembly: MySql.Data (in MySql.Data.dll)

**See Also**

[MySqlConnection Members](#) , [MySql.Data.MySqlClient Namespace](#)

**18.2.4.1.4.1. MySqlConnection Members**

[MySqlConnection overview](#)

**Public Instance Properties**

Data(inherited from Exception)	Gets a collection of key/value pairs that provide additional, user-defined information about the exception.
HelpLink(inherited from Exception)	Gets or sets a link to the help file associated with this exception.
InnerException(inherited from Exception)	Gets the Exception instance that caused the current exception.
Message(inherited from Exception)	Gets a message that describes the current exception.
<a href="#">Number</a>	Gets a number that identifies the type of error.
Source(inherited from Exception)	Gets or sets the name of the application or the object that causes the error.
StackTrace(inherited from Exception)	Gets a string representation of the frames on the call stack at the time the current exception was thrown.
TargetSite(inherited from Exception)	Gets the method that throws the current exception.

**Public Instance Methods**

Equals(inherited from Object)	Determines whether the specified Object is equal to the current Object.
GetBaseException(inherited from Exception)	When overridden in a derived class, returns the Exception that is the root cause of one or more subsequent exceptions.
GetHashCode(inherited from Object)	Serves as a hash function for a particular type. GetHashCode is suitable for use in hashing algorithms and data structures like a hash table.
GetObjectData(inherited from Exception)	When overridden in a derived class, sets the Serial-

	izationInfo with information about the exception.
GetType(inherited from Exception)	Gets the runtime type of the current instance.
ToString(inherited from Exception)	Creates and returns a string representation of the current exception.

**See Also**

[MySqlException Class](#) , [MySql.Data.MySqlClient Namespace](#)

**18.2.4.1.4.1.1. Number Property**

Gets a number that identifies the type of error.

**Syntax: Visual Basic**

```
Public ReadOnly Property Number As Integer
```

**Syntax: C#**

```
public int Number {get;}
```

**See Also**

[MySqlException Class](#) , [MySql.Data.MySqlClient Namespace](#)

**18.2.4.1.5. MySqlHelper Class**

Helper class that makes it easier to work with the provider.

For a list of all members of this type, see [MySqlHelper Members](#) .

**Syntax: Visual Basic**

```
NotInheritable Public Class MySqlHelper
```

**Syntax: C#**

```
public sealed class MySqlHelper
```

**Thread Safety**

Public static (Shared in Visual Basic) members of this type are safe for multithreaded operations. Instance members are not guaranteed to be thread-safe.

**Requirements**

Namespace: [MySql.Data.MySqlClient](#)

Assembly: MySql.Data (in MySql.Data.dll)

**See Also**

[MySqlHelper Members](#) , [MySql.Data.MySqlClient Namespace](#)

**18.2.4.1.5.1. MySqlHelper Members**

[MySqlHelper overview](#)**Public Static (Shared) Methods**

<a href="#">ExecuteDataRow</a>	Executes a single SQL command and returns the first row of the resultset. A new MySqlConnection object is created, opened, and closed during this method.
<a href="#">ExecuteDataset</a>	Overloaded. Executes a single SQL command and returns the resultset in a DataSet. A new MySqlConnection object is created, opened, and closed during this method.
<a href="#">ExecuteNonQuery</a>	Overloaded. Executes a single command against a MySQL database. The <a href="#">MySqlConnection</a> is assumed to be open when the method is called and remains open after the method completes.
<a href="#">ExecuteReader</a>	Overloaded. Executes a single command against a MySQL database.
<a href="#">ExecuteScalar</a>	Overloaded. Execute a single command against a MySQL database.
<a href="#">UpdateDataSet</a>	Updates the given table with data from the given DataSet

**Public Instance Methods**

Equals(inherited from Object)	Determines whether the specified Object is equal to the current Object.
GetHashCode(inherited from Object)	Serves as a hash function for a particular type. GetHashCode is suitable for use in hashing algorithms and data structures like a hash table.
GetType(inherited from Object)	Gets the Type of the current instance.
ToString(inherited from Object)	Returns a String that represents the current Object.

**See Also**[MySqlHelper Class](#) , [MySql.Data.MySqlClient Namespace](#)18.2.4.1.5.1.1. [MySqlHelper.ExecuteDataRow Method](#)

Executes a single SQL command and returns the first row of the resultset. A new MySqlConnection object is created, opened, and closed during this method.

**Syntax: Visual Basic**

```
Public Shared Function ExecuteDataRow( _
    ByVal connectionString As String, _
    ByVal commandText As String, _
    ParamArray parms As MySqlParameter() _
) As DataRow
```

**Syntax: C#**

```
public static DataRow ExecuteDataRow(
    string connectionString,
```



```
stringcommandText,  
    params MySqlParameter[]parms  
);
```

### Parameters

- `connectionString`: Settings to be used for the connection
- `commandText`: Command to execute
- `parms`: Parameters to use for the command

### Return Value

DataRow containing the first row of the resultset

### See Also

[MySqlHelper Class](#) , [MySql.Data.MySqlClient Namespace](#)

#### 18.2.4.1.5.1.2. ExecuteDataset Method

Executes a single SQL command and returns the resultset in a DataSet. The state of the [MySqlConnection](#) object remains unchanged after execution of this method.

#### Overload List

Executes a single SQL command and returns the resultset in a DataSet. The state of the [MySqlConnection](#) object remains unchanged after execution of this method.

- `public static DataSet ExecuteDataset(MySqlConnection,string);`

Executes a single SQL command and returns the resultset in a DataSet. The state of the [MySqlConnection](#) object remains unchanged after execution of this method.

- `public static DataSet ExecuteDataset(MySqlConnection,string,params MySqlParameter[]);`

Executes a single SQL command and returns the resultset in a DataSet. A new MySqlConnection object is created, opened, and closed during this method.

- `public static DataSet ExecuteDataset(string,string);`

Executes a single SQL command and returns the resultset in a DataSet. A new MySqlConnection object is created, opened, and closed during this method.

- `public static DataSet ExecuteDataset(string,string,params MySqlParameter[]);`

### See Also

[MySqlHelper Class](#) , [MySql.Data.MySqlClient Namespace](#)

#### 18.2.4.1.5.1.2.1. MySqlConnection.ExecuteDataset Method (MySqlConnection, String)

Executes a single SQL command and returns the resultset in a DataSet. The state of the [MySqlConnection](#) object remains unchanged after execution of this method.

##### Syntax: Visual Basic

```
Overloads Public Shared Function ExecuteDataset( _  
    ByVal connection As MySqlConnection, _  
    ByVal commandText As String _  
) As DataSet
```

##### Syntax: C#

```
public static DataSet ExecuteDataset(  
    MySqlConnection connection,  
    string commandText  
);
```

##### Parameters

- `connection`: [MySqlConnection](#) object to use
- `commandText`: Command to execute

##### Return Value

DataSet containing the resultset

##### See Also

[MySqlHelper Class](#) , [MySql.Data.MySqlClient Namespace](#) , [MySqlHelper.ExecuteDataset Overload List](#)

#### 18.2.4.1.5.1.2.2. MySqlConnection.ExecuteDataset Method (MySqlConnection, String, MySqlParameter[])

Executes a single SQL command and returns the resultset in a DataSet. The state of the [MySqlConnection](#) object remains unchanged after execution of this method.

##### Syntax: Visual Basic

```
Overloads Public Shared Function ExecuteDataset( _  
    ByVal connection As MySqlConnection, _  
    ByVal commandText As String, _  
    ParamArray commandParameters As MySqlParameter() _  
) As DataSet
```

##### Syntax: C#

```
public static DataSet ExecuteDataset(  
    MySqlConnection connection,  
    string commandText,  
    params MySqlParameter[] commandParameters  
);
```

##### Parameters

- `connection`: [MySqlConnection](#) object to use

- `commandText`: Command to execute
- `commandParameters`: Parameters to use for the command

**Return Value**

DataSet containing the resultset

**See Also**

[MySQLHelper Class](#) , [MySQL.Data.MySqlClient Namespace](#) , [MySQLHelper.ExecuteDataset Overload List](#)

#### 18.2.4.1.5.1.2.3. MySQLHelper.ExecuteDataset Method (String, String)

Executes a single SQL command and returns the resultset in a DataSet. A new MySqlConnection object is created, opened, and closed during this method.

**Syntax: Visual Basic**

```
Overloads Public Shared Function ExecuteDataset( _  
    ByVal connectionString As String, _  
    ByVal commandText As String _  
) As DataSet
```

**Syntax: C#**

```
public static DataSet ExecuteDataset(  
    string connectionString,  
    string commandText  
);
```

**Parameters**

- `connectionString`: Settings to be used for the connection
- `commandText`: Command to execute

**Return Value**

DataSet containing the resultset

**See Also**

[MySQLHelper Class](#) , [MySQL.Data.MySqlClient Namespace](#) , [MySQLHelper.ExecuteDataset Overload List](#)

#### 18.2.4.1.5.1.2.4. MySQLHelper.ExecuteDataset Method (String, String, MySqlParameter[])

Executes a single SQL command and returns the resultset in a DataSet. A new MySqlConnection object is created, opened, and closed during this method.

**Syntax: Visual Basic**

```
Overloads Public Shared Function ExecuteDataset( _  
    ByVal connectionString As String, _  
    ByVal commandText As String, _  
    ParamArray commandParameters As MySqlParameter() _  
) As DataSet
```

**Syntax: C#**

```
public static DataSet ExecuteDataset(
    string connectionString,
    string commandText,
    params MySqlParameter[] commandParameters
);
```

**Parameters**

- `connectionString`: Settings to be used for the connection
- `commandText`: Command to execute
- `commandParameters`: Parameters to use for the command

**Return Value**

DataSet containing the resultset

**See Also**

[MySqlHelper Class](#) , [MySql.Data.MySqlClient Namespace](#) , [MySqlHelper.ExecuteDataset Overload List](#)

## 18.2.4.1.5.1.3. ExecuteNonQuery Method

Executes a single command against a MySQL database. The [MySqlConnection](#) is assumed to be open when the method is called and remains open after the method completes.

**Overload List**

Executes a single command against a MySQL database. The [MySqlConnection](#) is assumed to be open when the method is called and remains open after the method completes.

- `public static int ExecuteNonQuery(MySqlConnection, string, params MySqlParameter[]);`

Executes a single command against a MySQL database. A new [MySqlConnection](#) is created using the [ConnectionString](#) given.

- `public static int ExecuteNonQuery(string, string, params MySqlParameter[]);`

**See Also**

[MySqlHelper Class](#) , [MySql.Data.MySqlClient Namespace](#)

## 18.2.4.1.5.1.3.1. MySqlHelper.ExecuteNonQuery Method (MySqlConnection, String, MySqlParameter[])

Executes a single command against a MySQL database. The [MySqlConnection](#) is assumed to be open when the method is called and remains open after the method completes.

**Syntax: Visual Basic**

```
Overloads Public Shared Function ExecuteNonQuery( _
    ByVal connection As MySqlConnection, _
    ByVal commandText As String, _
```

```
ParamArray commandParameters As MySqlParameter() _
) As Integer
```

**Syntax: C#**

```
public static int ExecuteNonQuery(
    MySqlConnection connection,
    string commandText,
    params MySqlParameter[] commandParameters
);
```

**Parameters**

- **connection**: [MySqlConnection](#) object to use
- **commandText**: SQL command to be executed
- **commandParameters**: Array of [MySqlParameter](#) objects to use with the command.

**Return Value****See Also**

[MySqlHelper Class](#) , [MySql.Data.MySqlClient Namespace](#) , [MySqlHelper.ExecuteNonQuery Overload List](#)

#### 18.2.4.1.5.1.3.2. MySqlHelper.ExecuteNonQuery Method (String, String, MySqlParameter[])

Executes a single command against a MySQL database. A new [MySqlConnection](#) is created using the [ConnectionString](#) given.

**Syntax: Visual Basic**

```
Overloads Public Shared Function ExecuteNonQuery( _
    ByVal connectionString As String, _
    ByVal commandText As String, _
    ParamArray parms As MySqlParameter() _
) As Integer
```

**Syntax: C#**

```
public static int ExecuteNonQuery(
    string connectionString,
    string commandText,
    params MySqlParameter[] parms
);
```

**Parameters**

- **connectionString**: [ConnectionString](#) to use
- **commandText**: SQL command to be executed
- **parms**: Array of [MySqlParameter](#) objects to use with the command.

**Return Value****See Also**

[MySQLHelper Class](#) , [MySQL.Data.MySqlClient Namespace](#) , [MySQLHelper.ExecuteNonQuery Overload List](#)

#### 18.2.4.1.5.1.4. ExecuteReader Method

Executes a single command against a MySQL database.

##### Overload List

Executes a single command against a MySQL database.

- `public static MySqlDataReader ExecuteReader(string,string);`

Executes a single command against a MySQL database.

- `public static MySqlDataReader ExecuteReader(string,string,params MySqlParameter[]);`

##### See Also

[MySQLHelper Class](#) , [MySQL.Data.MySqlClient Namespace](#)

#### 18.2.4.1.5.1.4.1. MySQLHelper.ExecuteReader Method (String, String)

Executes a single command against a MySQL database.

##### Syntax: Visual Basic

```
Overloads Public Shared Function ExecuteReader( _  
    ByVal connectionString As String, _  
    ByVal commandText As String _  
) As MySqlDataReader
```

##### Syntax: C#

```
public static MySqlDataReader ExecuteReader(  
    stringconnectionString,  
    stringcommandText  
);
```

##### Parameters

- `connectionString`: Settings to use for this command
- `commandText`: Command text to use

##### Return Value

[MySqlDataReader](#) object ready to read the results of the command

##### See Also

[MySQLHelper Class](#) , [MySQL.Data.MySqlClient Namespace](#) , [MySQLHelper.ExecuteReader Overload List](#)

#### 18.2.4.1.5.1.4.2. MySQLHelper.ExecuteReader Method (String, String, MySqlParameter[])

Executes a single command against a MySQL database.

**Syntax: Visual Basic**

```
Overloads Public Shared Function ExecuteReader( _  
    ByVal connectionString As String, _  
    ByVal commandText As String, _  
    ParamArray commandParameters As MySqlParameter() _  
) As MySqlDataReader
```

**Syntax: C#**

```
public static MySqlDataReader ExecuteReader(  
    string connectionString,  
    string commandText,  
    params MySqlParameter[] commandParameters  
);
```

**Parameters**

- `connectionString`: Settings to use for this command
- `commandText`: Command text to use
- `commandParameters`: Array of [MySqlParameter](#) objects to use with the command

**Return Value**

[MySqlDataReader](#) object ready to read the results of the command

**See Also**

[MySqlHelper Class](#) , [MySql.Data.MySqlClient Namespace](#) , [MySqlHelper.ExecuteReader Overload List](#)

#### 18.2.4.1.5.1.5. ExecuteScalar Method

Execute a single command against a MySQL database.

**Overload List**

Execute a single command against a MySQL database.

- [public static object ExecuteScalar\(MySqlConnection,string\);](#)

Execute a single command against a MySQL database.

- [public static object ExecuteScalar\(MySqlConnection,string,params MySqlParameter\[\]\);](#)

Execute a single command against a MySQL database.

- [public static object ExecuteScalar\(string,string\);](#)

Execute a single command against a MySQL database.

- [public static object ExecuteScalar\(string,string,params MySqlParameter\[\]\);](#)

#### See Also

[MySqlHelper Class](#) , [MySql.Data.MySqlClient Namespace](#)

#### 18.2.4.1.5.1.5.1. MySqlHelper.ExecuteScalar Method (MySqlConnection, String)

Execute a single command against a MySQL database.

#### Syntax: Visual Basic

```
Overloads Public Shared Function ExecuteScalar( _
    ByVal connection As MySqlConnection, _
    ByVal commandText As String _
) As Object
```

#### Syntax: C#

```
public static object ExecuteScalar(
    MySqlConnectionconnection,
    stringcommandText
);
```

#### Parameters

- **connection**: [MySqlConnection](#) object to use
- **commandText**: Command text to use for the command

#### Return Value

The first column of the first row in the result set, or a null reference if the result set is empty.

#### See Also

[MySqlHelper Class](#) , [MySql.Data.MySqlClient Namespace](#) , [MySqlHelper.ExecuteScalar Overload List](#)

#### 18.2.4.1.5.1.5.2. MySqlHelper.ExecuteScalar Method (MySqlConnection, String, MySqlParameter[])

Execute a single command against a MySQL database.

#### Syntax: Visual Basic

```
Overloads Public Shared Function ExecuteScalar( _
    ByVal connection As MySqlConnection, _
    ByVal commandText As String, _
    ParamArray commandParameters As MySqlParameter() _
) As Object
```

#### Syntax: C#

```
public static object ExecuteScalar(
    MySqlConnectionconnection,
    stringcommandText,
    params MySqlParameter[]commandParameters
);
```



**Parameters**

- `connection`: [MySQLConnection](#) object to use
- `commandText`: Command text to use for the command
- `commandParameters`: Parameters to use for the command

**Return Value**

The first column of the first row in the result set, or a null reference if the result set is empty.

**See Also**

[MySQLHelper Class](#) , [MySQL.Data.MySqlClient Namespace](#) , [MySQLHelper.ExecuteScalar Overload List](#)

## 18.2.4.1.5.1.5.3. MySQLHelper.ExecuteScalar Method (String, String)

Execute a single command against a MySQL database.

**Syntax: Visual Basic**

```
Overloads Public Shared Function ExecuteScalar( _
    ByVal connectionString As String, _
    ByVal commandText As String _
) As Object
```

**Syntax: C#**

```
public static object ExecuteScalar(
    string connectionString,
    string commandText
);
```

**Parameters**

- `connectionString`: Settings to use for the update
- `commandText`: Command text to use for the update

**Return Value**

The first column of the first row in the result set, or a null reference if the result set is empty.

**See Also**

[MySQLHelper Class](#) , [MySQL.Data.MySqlClient Namespace](#) , [MySQLHelper.ExecuteScalar Overload List](#)

## 18.2.4.1.5.1.5.4. MySQLHelper.ExecuteScalar Method (String, String, MySQLParameter[])

Execute a single command against a MySQL database.

**Syntax: Visual Basic**

```
Overloads Public Shared Function ExecuteScalar( _
    ByVal connectionString As String, _
    ByVal commandText As String, _
    ParamArray commandParameters As MySQLParameter() _
```

```
) As Object
```

**Syntax: C#**

```
public static object ExecuteScalar(  
    string connectionString,  
    string commandText,  
    params MySqlParameter[] commandParameters  
);
```

**Parameters**

- `connectionString`: Settings to use for the command
- `commandText`: Command text to use for the command
- `commandParameters`: Parameters to use for the command

**Return Value**

The first column of the first row in the result set, or a null reference if the result set is empty.

**See Also**

[MySqlHelper Class](#) , [MySql.Data.MySqlClient Namespace](#) , [MySqlHelper.ExecuteScalar Overload List](#)

#### 18.2.4.1.5.1.6. MySqlHelper.UpdateDataSet Method

Updates the given table with data from the given DataSet

**Syntax: Visual Basic**

```
Public Shared Sub UpdateDataSet( _  
    ByVal connectionString As String, _  
    ByVal commandText As String, _  
    ByVal ds As DataSet, _  
    ByVal tablename As String _  
)
```

**Syntax: C#**

```
public static void UpdateDataSet(  
    string connectionString,  
    string commandText,  
    DataSet ds,  
    string tablename  
);
```

**Parameters**

- `connectionString`: Settings to use for the update
- `commandText`: Command text to use for the update
- `ds`: DataSet containing the new data to use in the update
- `tablename`: Tablename in the dataset to update

**See Also**

[MySqlHelper Class](#) , [MySql.Data.MySqlClient Namespace](#)

**18.2.4.1.6. MySqlConnection ErrorCode Enumeration****Syntax: Visual Basic**

```
Public Enum MySqlConnectionErrorCode
```

**Syntax: C#**

```
public enum MySqlConnectionErrorCode
```

**Members**

Member Name	Description
PacketTooLarge	
PasswordNotAllowed	
DuplicateKeyEntry	
HostNotPrivileged	
PasswordNoMatch	
AnonymousUser	
DuplicateKey	
KeyNotFound	
DuplicateKeyName	

**Requirements**

Namespace: [MySql.Data.MySqlClient](#)

Assembly: MySql.Data (in MySql.Data.dll)

**See Also**

[MySql.Data.MySqlClient Namespace](#)

**18.2.4.2. MySql.Data.Types**

[Namespace hierarchy](#)

**Classes**

Class	Description
<a href="#">MySqlConnectionException</a>	Summary description for MySqlConnectionException.
<a href="#">MySqlConnectionDateTime</a>	Summary description for MySqlConnectionDateTime.
<a href="#">MySqlConnectionValue</a>	

**18.2.4.2.1. MySqlConnection.TypesHierarchy**

**See Also**[MySQL.Data.Types Namespace](#)**18.2.4.2.2. MySqlConnectionException Class**

Summary description for MySqlConnectionException.

For a list of all members of this type, see [MySQLConversionException Members](#) .

**Syntax: Visual Basic**

```
Public Class MySqlConnectionException_
    Inherits ApplicationException
```

**Syntax: C#**

```
public class MySqlConnectionException : ApplicationException
```

**Thread Safety**

Public static (Shared in Visual Basic) members of this type are safe for multithreaded operations. Instance members are not guaranteed to be thread-safe.

**Requirements**

Namespace: [MySQL.Data.Types](#)

Assembly: MySQL.Data (in MySQL.Data.dll)

**See Also**

[MySQLConversionException Members](#) , [MySQL.Data.Types Namespace](#)

**18.2.4.2.2.1. MySqlConnectionException Members**

[MySQLConversionException overview](#)

**Public Instance Constructors**

<a href="#">MySQLConversionException Constructor</a>	Ctor
--	------

**Public Instance Properties**

Data(inherited from Exception)	Gets a collection of key/value pairs that provide additional, user-defined information about the exception.
HelpLink(inherited from Exception)	Gets or sets a link to the help file associated with this exception.
InnerException(inherited from Exception)	Gets the Exception instance that caused the current exception.
Message(inherited from Exception)	Gets a message that describes the current exception.
Source(inherited from Exception)	Gets or sets the name of the application or the object that causes the error.

StackTrace(inherited from Exception)	Gets a string representation of the frames on the call stack at the time the current exception was thrown.
TargetSite(inherited from Exception)	Gets the method that throws the current exception.

### Public Instance Methods

Equals(inherited from Object)	Determines whether the specified Object is equal to the current Object.
GetBaseException(inherited from Exception)	When overridden in a derived class, returns the Exception that is the root cause of one or more subsequent exceptions.
GetHashCode(inherited from Object)	Serves as a hash function for a particular type. GetHashCode is suitable for use in hashing algorithms and data structures like a hash table.
GetObjectData(inherited from Exception)	When overridden in a derived class, sets the SerializationInfo with information about the exception.
GetType(inherited from Exception)	Gets the runtime type of the current instance.
ToString(inherited from Exception)	Creates and returns a string representation of the current exception.

### Protected Instance Properties

HResult(inherited from Exception)	Gets or sets HRESULT, a coded numerical value that is assigned to a specific exception.
-----------------------------------	---

### Protected Instance Methods

Finalize(inherited from Object)	Allows an Object to attempt to free resources and perform other cleanup operations before the Object is reclaimed by garbage collection.
MemberwiseClone(inherited from Object)	Creates a shallow copy of the current Object.

### See Also

[MySQLConversionException Class](#) , [MySQL.Data.Types Namespace](#)

#### 18.2.4.2.2.1.1. MySQLConversionException Constructor

##### Syntax: Visual Basic

```
Public Sub New( _
    ByVal msg As String _
)
```

##### Syntax: C#

```
public MySQLConversionException(
    string msg
);
```

**See Also**

[MySQLConversionException Class](#) , [MySQL.Data.Types Namespace](#)

**18.2.4.2.3. MySqlDateTime Class**

Summary description for MySqlDateTime.

For a list of all members of this type, see [MySqlDateTime Members](#) .

**Syntax: Visual Basic**

```
Public Class MySqlDateTime_
    Inherits MySqlValue_
    Implements IConvertible, IComparable
```

**Syntax: C#**

```
public class MySqlDateTime : MySqlValue, IConvertible, IComparable
```

**Thread Safety**

Public static (Shared in Visual Basic) members of this type are safe for multithreaded operations. Instance members are not guaranteed to be thread-safe.

**Requirements**

Namespace: [MySQL.Data.Types](#)

Assembly: MySQL.Data (in MySQL.Data.dll)

**See Also**

[MySqlDateTime Members](#) , [MySQL.Data.Types Namespace](#)

**18.2.4.2.3.1. MySqlDateTime Members**

[MySqlDateTime overview](#)

**Public Static (Shared) Type Conversions**

<a href="#">Explicit MySqlDateTime to DateTime Conversion</a>	
---	--

**Public Instance Properties**

<a href="#">Day</a>	Returns the day portion of this datetime
<a href="#">Hour</a>	Returns the hour portion of this datetime
<a href="#">IsNull</a> (inherited from <a href="#">MySqlValue</a> )	
<a href="#">IsValidDateTime</a>	Indicates if this object contains a value that can be represented as a <a href="#">DateTime</a>
<a href="#">Minute</a>	Returns the minute portion of this datetime
<a href="#">Month</a>	Returns the month portion of this datetime
<a href="#">Second</a>	Returns the second portion of this datetime
<a href="#">ValueAsObject</a> (inherited from <a href="#">MySqlValue</a> )	Returns the value of this field as an object

<a href="#">Year</a>	Returns the year portion of this datetime
----------------------	---

**Public Instance Methods**

<a href="#">Equals</a> (inherited from Object)	Determines whether the specified Object is equal to the current Object.
<a href="#">GetDateTime</a>	Returns this value as a DateTime
<a href="#">GetHashCode</a> (inherited from Object)	Serves as a hash function for a particular type. GetHashCode is suitable for use in hashing algorithms and data structures like a hash table.
<a href="#">GetType</a> (inherited from Object)	Gets the Type of the current instance.
<a href="#">ToString</a>	Returns a MySQL specific string representation of this value

**Protected Instance Fields**

<a href="#">classType</a> (inherited from MySqlValue)	The system type represented by this value
<a href="#">dbType</a> (inherited from MySqlValue)	The generic dbtype of this value
<a href="#">isNull</a> (inherited from MySqlValue)	Is this value null
<a href="#">mysqlDbType</a> (inherited from MySqlValue)	The specific MySQL db type
<a href="#">mysqlTypeName</a> (inherited from MySqlValue)	The MySQL specific typename of this value
<a href="#">objectValue</a> (inherited from MySqlValue)	

**Protected Instance Methods**

<a href="#">Finalize</a> (inherited from Object)	Allows an Object to attempt to free resources and perform other cleanup operations before the Object is reclaimed by garbage collection.
<a href="#">MemberwiseClone</a> (inherited from Object)	Creates a shallow copy of the current Object.

**See Also**

[MySqlDateTime Class](#) , [MySql.Data.Types Namespace](#)

## 18.2.4.2.3.1.1. MySqlDateTime Explicit MySqlDateTime to DateTime Conversion

**Syntax: Visual Basic**

```
MySqlDateTime.op_Explicit(val)
```

**Syntax: C#**

```
public static explicit operator DateTime(
    MySqlDateTime val
);
```

**Parameters**

- `val:`

**Return Value**

**See Also**

[MySQLDateTime Class](#) , [MySQL.Data.Types Namespace](#)

18.2.4.2.3.1.2. Day Property

Returns the day portion of this datetime

**Syntax: Visual Basic**

```
Public Property Day As Integer
```

**Syntax: C#**

```
public int Day {get; set;}
```

**See Also**

[MySQLDateTime Class](#) , [MySQL.Data.Types Namespace](#)

18.2.4.2.3.1.3. Hour Property

Returns the hour portion of this datetime

**Syntax: Visual Basic**

```
Public Property Hour As Integer
```

**Syntax: C#**

```
public int Hour {get; set;}
```

**See Also**

[MySQLDateTime Class](#) , [MySQL.Data.Types Namespace](#)

18.2.4.2.3.1.4. IsNull Property

**Syntax: Visual Basic**

```
Public Property IsNull As Boolean
```

**Syntax: C#**

```
public bool IsNull {get; set;}
```

**See Also**

[MySQLValue Class](#) , [MySQL.Data.Types Namespace](#)

18.2.4.2.3.1.4.1. MySQLValue Class



For a list of all members of this type, see [MySqlValue Members](#) .

### Syntax: Visual Basic

```
MustInherit Public Class MySqlValue
```

### Syntax: C#

```
public abstract class MySqlValue
```

### Thread Safety

Public static (Shared in Visual Basic) members of this type are safe for multithreaded operations. Instance members are not guaranteed to be thread-safe.

### Requirements

Namespace: [MySql.Data.Types](#)

Assembly: MySql.Data (in MySql.Data.dll)

### See Also

[MySqlValue Members](#) , [MySql.Data.Types Namespace](#)

#### 18.2.4.2.3.1.4.1.1. MySqlValue Members

[MySqlValue overview](#)

### Protected Static (Shared) Fields

<a href="#">numberFormat</a>	
------------------------------	--

### Public Instance Constructors

<a href="#">MySqlValue Constructor</a>	Initializes a new instance of the <a href="#">MySqlValue</a> class.
--	---

### Public Instance Properties

<a href="#">IsNull</a>	
<a href="#">ValueAsObject</a>	Returns the value of this field as an object

### Public Instance Methods

<a href="#">Equals</a> (inherited from Object)	Determines whether the specified Object is equal to the current Object.
<a href="#">GetHashCode</a> (inherited from Object)	Serves as a hash function for a particular type. <a href="#">GetHashCode</a> is suitable for use in hashing algorithms and data structures like a hash table.
<a href="#">GetType</a> (inherited from Object)	Gets the Type of the current instance.
<a href="#">ToString</a>	Returns a string representation of this value

**Protected Instance Fields**

<a href="#">classType</a>	The system type represented by this value
<a href="#">dbType</a>	The generic dbtype of this value
<a href="#">isNull</a>	Is this value null
<a href="#">mysqlDbType</a>	The specific MySQL db type
<a href="#">mysqlTypeName</a>	The MySQL specific typename of this value
<a href="#">objectValue</a>	

**Protected Instance Methods**

<a href="#">Finalize</a> (inherited from Object)	Allows an Object to attempt to free resources and perform other cleanup operations before the Object is reclaimed by garbage collection.
<a href="#">MemberwiseClone</a> (inherited from Object)	Creates a shallow copy of the current Object.

**See Also**

[MySqlValue Class](#) , [MySql.Data.Types Namespace](#)

18.2.4.2.3.1.4.1.1.1. [MySqlValue.numberFormat](#) Field**Syntax: Visual Basic**

```
Protected Shared numberFormat As NumberFormatInfo
```

**Syntax: C#**

```
protected static NumberFormatInfo numberFormat;
```

**See Also**

[MySqlValue Class](#) , [MySql.Data.Types Namespace](#)

18.2.4.2.3.1.4.1.1.2. [MySqlValue](#) Constructor

Initializes a new instance of the [MySqlValue](#) class.

**Syntax: Visual Basic**

```
Public Sub New()
```

**Syntax: C#**

```
public MySqlValue();
```

**See Also**

[MySqlValue Class](#) , [MySql.Data.Types Namespace](#)

18.2.4.2.3.1.4.1.1.3. [ValueAsObject](#) Property

Returns the value of this field as an object

**Syntax: Visual Basic**

```
Public ReadOnly Property ValueAsObject As Object
```

**Syntax: C#**

```
public object ValueAsObject {get;}
```

**See Also**

[MySQLValue Class](#) , [MySQL.Data.Types Namespace](#)

## 18.2.4.2.3.1.4.1.1.4. MySQLValue.ToString Method

Returns a string representation of this value

**Syntax: Visual Basic**

```
Overrides Public Function ToString() As String
```

**Syntax: C#**

```
public override string ToString();
```

**See Also**

[MySQLValue Class](#) , [MySQL.Data.Types Namespace](#)

## 18.2.4.2.3.1.4.1.1.5. MySQLValue.classType Field

The system type represented by this value

**Syntax: Visual Basic**

```
Protected classType As Type
```

**Syntax: C#**

```
protected Type classType;
```

**See Also**

[MySQLValue Class](#) , [MySQL.Data.Types Namespace](#)

## 18.2.4.2.3.1.4.1.1.6. MySQLValue.dbType Field

The generic dbtype of this value

**Syntax: Visual Basic**

```
Protected dbType As DbType
```

**Syntax: C#**

```
protected DbType dbType;
```

**See Also**

[MySqlValue Class](#) , [MySql.Data.Types Namespace](#)

## 18.2.4.2.3.1.4.1.1.7. MySqlValue.mySqlDbType Field

The specific MySQL db type

**Syntax: Visual Basic**

```
Protected MySqlDbType As MySqlDbType
```

**Syntax: C#**

```
protected MySqlDbType mySqlDbType;
```

**See Also**

[MySqlValue Class](#) , [MySql.Data.Types Namespace](#)

## 18.2.4.2.3.1.4.1.1.8. MySqlValue.mySqlTypeName Field

The MySQL specific typename of this value

**Syntax: Visual Basic**

```
Protected MySqlTypeName As String
```

**Syntax: C#**

```
protected string mySqlTypeName;
```

**See Also**

[MySqlValue Class](#) , [MySql.Data.Types Namespace](#)

## 18.2.4.2.3.1.4.1.1.9. MySqlValue.objectValue Field

**Syntax: Visual Basic**

```
Protected objectValue As Object
```

**Syntax: C#**

```
protected object objectValue;
```

**See Also**

[MySqlValue Class](#) , [MySql.Data.Types Namespace](#)

## 18.2.4.2.3.1.5. IsValidDateTime Property

Indicates if this object contains a value that can be represented as a DateTime

**Syntax: Visual Basic**

```
Public ReadOnly Property IsValidDateTime As Boolean
```

**Syntax: C#**

```
public bool IsValidDateTime {get;}
```

**See Also**

[MySqlDateTime Class](#) , [MySql.Data.Types Namespace](#)

## 18.2.4.2.3.1.6. Minute Property

Returns the minute portion of this datetime

**Syntax: Visual Basic**

```
Public Property Minute As Integer
```

**Syntax: C#**

```
public int Minute {get; set;}
```

**See Also**

[MySqlDateTime Class](#) , [MySql.Data.Types Namespace](#)

## 18.2.4.2.3.1.7. Month Property

Returns the month portion of this datetime

**Syntax: Visual Basic**

```
Public Property Month As Integer
```

**Syntax: C#**

```
public int Month {get; set;}
```

**See Also**

[MySqlDateTime Class](#) , [MySql.Data.Types Namespace](#)

## 18.2.4.2.3.1.8. Second Property

Returns the second portion of this datetime

**Syntax: Visual Basic**

```
Public Property Second As Integer
```

**Syntax: C#**

```
public int Second {get; set;}
```

**See Also**

[MySqlDateTime Class](#) , [MySql.Data.Types Namespace](#)

## 18.2.4.2.3.1.9. Year Property

Returns the year portion of this datetime

**Syntax: Visual Basic**

```
Public Property Year As Integer
```

**Syntax: C#**

```
public int Year {get; set;}
```

**See Also**

[MySQLDateTime Class](#) , [MySQL.Data.Types Namespace](#)

#### 18.2.4.2.3.1.10. MySQLDateTime.GetDateTime Method

Returns this value as a DateTime

**Syntax: Visual Basic**

```
Public Function GetDateTime() As Date
```

**Syntax: C#**

```
public DateTime GetDateTime();
```

**See Also**

[MySQLDateTime Class](#) , [MySQL.Data.Types Namespace](#)

#### 18.2.4.2.3.1.11. MySQLDateTime.ToString Method

Returns a MySQL specific string representation of this value

**Syntax: Visual Basic**

```
Overrides Public Function ToString() As String
```

**Syntax: C#**

```
public override string ToString();
```

**See Also**

[MySQLDateTime Class](#) , [MySQL.Data.Types Namespace](#)

## 18.2.5. Connector/NET Notes and Tips

In this section we will cover some of the more common use cases for Connector/NET, including BLOB handling, date handling, and using Connector/NET with common tools such as Crystal Reports.

### 18.2.5.1. Connecting to MySQL Using Connector/NET

#### 18.2.5.1.1. Introduction

All interaction between a .NET application and the MySQL server is routed through a [MySQLConnec-](#)

---

tion object. Before your application can interact with the server, a `MySQLConnection` object must be instantiated, configured, and opened.

Even when using the `MySQLHelper` class, a `MySQLConnection` object is created by the helper class.

In this section, we will describe how to connect to MySQL using the `MySQLConnection` object.

### 18.2.5.1.2. Creating a Connection String

The `MySQLConnection` object is configured using a connection string. A connection string contains sever key/value pairs, separated by semicolons. Each key/value pair is joined with an equals sign.

The following is a sample connection string:

```
Server=127.0.0.1;Uid=root;Pwd=12345;Database=test;
```

In this example, the `MySQLConnection` object is configured to connect to a MySQL server at `127.0.0.1`, with a username of `root` and a password of `12345`. The default database for all statements will be the `test` database.

The following options are typically used (a full list of options is available in the API documentation for [Section 18.2.3.3.15, "ConnectionString"](#)):

- **Server:** The name or network address of the instance of MySQL to which to connect. The default is `localhost`. Aliases include `host`, `Data Source`, `DataSource`, `Address`, `Addr` and `Network Address`.
- **Uid:** The MySQL user account to use when connecting. Aliases include `User Id`, `Username` and `User name`.
- **Pwd:** The password for the MySQL account being used. Alias `Password` can also be used.
- **Database:** The default database that all statements are applied to. Default is `mysql`. Alias `Initial Catalog` can also be used.
- **Port:** The port MySQL is using to listen for connections. Default is `3306`. Specify `-1` for this value to use a named-pipe connection.

### 18.2.5.1.3. Opening a Connection

Once you have created a connection string it can be used to open a connection to the MySQL server.

The following code is used to create a `MySQLConnection` object, assign the connection string, and open the connection.

Visual Basic Example

```
Dim conn As New MySql.Data.MySqlClient.MySqlConnection
Dim myConnectionString as String

myConnectionString = "server=127.0.0.1;" _
    & "uid=root;" _
    & "pwd=12345;" _
    & "database=test;"

Try
    conn.ConnectionString = myConnectionString
    conn.Open()
```

```
Catch ex As MySql.Data.MySqlClient.MySqlException
    MessageBox.Show(ex.Message)
End Try
```

### C# Example

```
MySql.Data.MySqlClient.MySqlConnection conn;
string myConnectionString;

myConnectionString = "server=127.0.0.1;uid=root;" +
    "pwd=12345;database=test;";

try
{
    conn = new MySql.Data.MySqlClient.MySqlConnection();
    conn.ConnectionString = myConnectionString;
    conn.Open();
}
catch (MySql.Data.MySqlClient.MySqlException ex)
{
    MessageBox.Show(ex.Message);
}
```

You can also pass the connection string to the constructor of the [MySqlConnection](#) class:

### Visual Basic Example

```
Dim myConnectionString as String

myConnectionString = "server=127.0.0.1;" _
    & "uid=root;" _
    & "pwd=12345;" _
    & "database=test;"

Try
    Dim conn As New MySql.Data.MySqlClient.MySqlConnection(myConnectionString)
    conn.Open()
Catch ex As MySql.Data.MySqlClient.MySqlException
    MessageBox.Show(ex.Message)
End Try
```

### C# Example

```
MySql.Data.MySqlClient.MySqlConnection conn;
string myConnectionString;

myConnectionString = "server=127.0.0.1;uid=root;" +
    "pwd=12345;database=test;";

try
{
    conn = new MySql.Data.MySqlClient.MySqlConnection(myConnectionString);
    conn.Open();
}
catch (MySql.Data.MySqlClient.MySqlException ex)
{
    MessageBox.Show(ex.Message);
}
```

Once the connection is open it can be used by the other Connector/.NET classes to communicate with the MySQL server.

#### 18.2.5.1.4. Handling Connection Errors

Because connecting to an external server is unpredictable, it is important to add error handling to your .NET application. When there is an error connecting, the [MySqlConnection](#) class will return a [MySqlException](#) object. This object has two properties that are of interest when handling errors:



- **Message:** A message that describes the current exception.
- **Number:** The MySQL error number.

When handling errors, you can your application's response based on the error number. The two most common error numbers when connecting are as follows:

- **0:** Cannot connect to server.
- **1045:** Invalid username and/or password.

The following code shows how to adapt the application's response based on the actual error:

#### Visual Basic Example

```
Dim myConnectionString as String

myConnectionString = "server=127.0.0.1;" _
    & "uid=root;" _
    & "pwd=12345;" _
    & "database=test;"

Try
    Dim conn As New MySql.Data.MySqlClient.MySqlConnection(myConnectionString)
    conn.Open()
Catch ex As MySql.Data.MySqlClient.MySqlException
    Select Case ex.Number
        Case 0
            MessageBox.Show("Cannot connect to server. Contact administrator")
        Case 1045
            MessageBox.Show("Invalid username/password, please try again")
    End Select
End Try
```

#### C# Example

```
MySql.Data.MySqlClient.MySqlConnection conn;
string myConnectionString;

myConnectionString = "server=127.0.0.1;uid=root;" +
    "pwd=12345;database=test;";

try
{
    conn = new MySql.Data.MySqlClient.MySqlConnection(myConnectionString);
    conn.Open();
}
catch (MySql.Data.MySqlClient.MySqlException ex)
{
    switch (ex.Number)
    {
        case 0:
            MessageBox.Show("Cannot connect to server. Contact administrator");
        case 1045:
            MessageBox.Show("Invalid username/password, please try again");
    }
}
```

**Important:** Note that if you are using multilanguage databases you must specify the character set in the connection string. If you do not specify the character set, the connection defaults to the `latin1` charset. You can specify the character set as part of the connection string, for example:

```
MySqlConnection myConnection = new MySqlConnection("server=127.0.0.1;uid=root;" +
    "pwd=12345;database=test;Charset=latin1;");
```

## 18.2.5.2. Using the Connector/NET with Prepared Statements

### 18.2.5.2.1. Introduction

As of MySQL 4.1, it is possible to use prepared statements with Connector/NET. Use of prepared statements can provide significant performance improvements on queries that are executed more than once.

Prepared execution is faster than direct execution for statements executed more than once, primarily because the query is parsed only once. In the case of direct execution, the query is parsed every time it is executed. Prepared execution also can provide a reduction of network traffic because for each execution of the prepared statement, it is necessary only to send the data for the parameters.

Another advantage of prepared statements is that it uses a binary protocol that makes data transfer between client and server more efficient.

### 18.2.5.2.2. Preparing Statements in Connector/NET

To prepare a statement, create a command object and set the `.CommandText` property to your query.

After entering your statement, call the `.Prepare` method of the `MySQLCommand` object. After the statement is prepared, add parameters for each of the dynamic elements in the query.

After you enter your query and enter parameters, execute the statement using the `.ExecuteNonQuery()`, `.ExecuteScalar()`, or `.ExecuteReader` methods.

For subsequent executions, you need only modify the values of the parameters and call the execute method again, there is no need to set the `.CommandText` property or redefine the parameters.

#### Visual Basic Example

```
Dim conn As New MySqlConnection
Dim cmd As New MySqlCommand

conn.ConnectionString = strConnection

Try
    conn.Open()
    cmd.Connection = conn

    cmd.CommandText = "INSERT INTO myTable VALUES(NULL, ?number, ?text)"
    cmd.Prepare()

    cmd.Parameters.Add("?number", 1)
    cmd.Parameters.Add("?text", "One")

    For i = 1 To 1000
        cmd.Parameters["?number"].Value = i
        cmd.Parameters["?text"].Value = "A string value"

        cmd.ExecuteNonQuery()
    Next
Catch ex As MySqlException
    MessageBox.Show("Error " & ex.Number & " has occurred: " & ex.Message, "Error", MessageBoxButtons.OK)
End Try
```

#### C# Example

```
MySQL.Data.MySqlClient.MySqlConnection conn;
MySQL.Data.MySqlClient.MySqlCommand cmd;

conn = new MySQL.Data.MySqlClient.MySqlConnection();
cmd = new MySQL.Data.MySqlClient.MySqlCommand();

conn.ConnectionString = strConnection;

try
```

```
{
    conn.Open();
    cmd.Connection = conn;

    cmd.CommandText = "INSERT INTO myTable VALUES(NULL, ?number, ?text)";
    cmd.Prepare();

    cmd.Parameters.Add("?number", 1);
    cmd.Parameters.Add("?text", "One");

    for (int i=1; i <= 1000; i++)
    {
        cmd.Parameters["?number"].Value = i;
        cmd.Parameters["?text"].Value = "A string value";

        cmd.ExecuteNonQuery();
    }
}
catch (MySql.Data.MySqlClient.MySqlException ex)
{
    MessageBox.Show("Error " + ex.Number + " has occurred: " + ex.Message,
        "Error", MessageBoxButtons.OK, MessageBoxIcon.Error);
}
```

### 18.2.5.3. Accessing Stored Procedures with Connector/NET

#### 18.2.5.3.1. Introduction

With the release of MySQL version 5 the MySQL server now supports stored procedures with the SQL 2003 stored procedure syntax.

A stored procedure is a set of SQL statements that can be stored in the server. Once this has been done, clients don't need to keep reissuing the individual statements but can refer to the stored procedure instead.

Stored procedures can be particularly useful in situations such as the following:

- When multiple client applications are written in different languages or work on different platforms, but need to perform the same database operations.
- When security is paramount. Banks, for example, use stored procedures for all common operations. This provides a consistent and secure environment, and procedures can ensure that each operation is properly logged. In such a setup, applications and users would not get any access to the database tables directly, but can only execute specific stored procedures.

Connector/NET supports the calling of stored procedures through the `MySqlCommand` object. Data can be passed in and out of a MySQL stored procedure through use of the `MySqlCommand.Parameters` collection.

**Note:** When you call a stored procedure, the command object makes an additional `SELECT` call to determine the parameters of the stored procedure. You must ensure that the user calling the procedure has the `SELECT` privilege on the `mysql.proc` table to enable them to verify the parameters. Failure to do this will result in an error when calling the procedure.

This section will not provide in-depth information on creating Stored Procedures. For such information, please refer to <http://dev.mysql.com/doc/mysql/en/stored-procedures.html>.

A sample application demonstrating how to use stored procedures with Connector/NET can be found in the `Samples` directory of your Connector/NET installation.

#### 18.2.5.3.2. Creating Stored Procedures from Connector/NET

Stored procedures in MySQL can be created using a variety of tools. First, stored procedures can be created using the [mysql](#) command-line client. Second, stored procedures can be created using the [MySQL Query Browser](#) GUI client. Finally, stored procedures can be created using the [.ExecuteNonQuery](#) method of the [MySqlCommand](#) object:

#### Visual Basic Example

```
Dim conn As New MySqlConnection
Dim cmd As New MySqlCommand

conn.ConnectionString = "server=127.0.0.1;" & _
    & "uid=root;" & _
    & "pwd=12345;" & _
    & "database=test"

Try
    conn.Open()
    cmd.Connection = conn

    cmd.CommandText = "CREATE PROCEDURE add_emp(" & _
        & "IN fname VARCHAR(20), IN lname VARCHAR(20), IN bday DATETIME, OUT empno INT) " & _
        & "BEGIN INSERT INTO emp(first_name, last_name, birthdate) " & _
        & "VALUES(fname, lname, DATE(bday)); SET empno = LAST_INSERT_ID(); END"

    cmd.ExecuteNonQuery()
Catch ex As MySqlException
    MessageBox.Show("Error " & ex.Number & " has occurred: " & ex.Message, "Error", MessageBoxButtons.OK, MessageBoxIcon.Error)
End Try
```

#### C# Example

```
MySql.Data.MySqlClient.MySqlConnection conn;
MySql.Data.MySqlClient.MySqlCommand cmd;

conn = new MySql.Data.MySqlClient.MySqlConnection();
cmd = new MySql.Data.MySqlClient.MySqlCommand();

conn.ConnectionString = "server=127.0.0.1;uid=root;" +
    "pwd=12345;database=test;";

try
{
    conn.Open();
    cmd.Connection = conn;

    cmd.CommandText = "CREATE PROCEDURE add_emp(" +
        "IN fname VARCHAR(20), IN lname VARCHAR(20), IN bday DATETIME, OUT empno INT) " +
        "BEGIN INSERT INTO emp(first_name, last_name, birthdate) " +
        "VALUES(fname, lname, DATE(bday)); SET empno = LAST_INSERT_ID(); END";

    cmd.ExecuteNonQuery();
}
catch (MySql.Data.MySqlClient.MySqlException ex)
{
    MessageBox.Show("Error " + ex.Number + " has occurred: " + ex.Message,
        "Error", MessageBoxButtons.OK, MessageBoxIcon.Error);
}
```

It should be noted that, unlike the command-line and GUI clients, you are not required to specify a special delimiter when creating stored procedures in Connector/NET.

### 18.2.5.3.3. Calling a Stored Procedure from Connector/NET

To call a stored procedure using Connector/NET, create a [MySqlCommand](#) object and pass the stored procedure name as the [.CommandText](#) property. Set the [.CommandType](#) property to [CommandType.StoredProcedure](#).

After the stored procedure is named, create one [MySqlCommand](#) parameter for every parameter in the stored procedure. [IN](#) parameters are defined with the parameter name and the object containing the value, [OUT](#) parameters are defined with the parameter name and the datatype that is expected to be re-

turned. All parameters need the parameter direction defined.

After defining parameters, call the stored procedure by using the `MySqlCommand.ExecuteNonQuery()` method:

### Visual Basic Example

```
Dim conn As New MySqlConnection
Dim cmd As New MySqlCommand

conn.ConnectionString = "server=127.0.0.1;" _
    & "uid=root;" _
    & "pwd=12345;" _
    & "database=test"

Try
    conn.Open()
    cmd.Connection = conn

    cmd.CommandText = "add_emp"
    cmd.CommandType = CommandType.StoredProcedure

    cmd.Parameters.Add("?lname", 'Jones')
    cmd.Parameters["?lname"].Direction = ParameterDirection.Input

    cmd.Parameters.Add("?fname", 'Tom')
    cmd.Parameters["?fname"].Direction = ParameterDirection.Input

    cmd.Parameters.Add("?bday", #12/13/1977 2:17:36 PM#)
    cmd.Parameters["?bday"].Direction = ParameterDirection.Input

    cmd.Parameters.Add("?empno", MySqlDbType.Int32)
    cmd.Parameters["?empno"].Direction = ParameterDirection.Output

    cmd.ExecuteNonQuery()

    MessageBox.Show(cmd.Parameters["?empno"].Value)
Catch ex As MySqlException
    MessageBox.Show("Error " & ex.Number & " has occurred: " & ex.Message, "Error", MessageBoxButtons.YesNo)
End Try
```

### C# Example

```
MySql.Data.MySqlClient.MySqlConnection conn;
MySql.Data.MySqlClient.MySqlCommand cmd;

conn = new MySql.Data.MySqlClient.MySqlConnection();
cmd = new MySql.Data.MySqlClient.MySqlCommand();

conn.ConnectionString = "server=127.0.0.1;uid=root;" +
    "pwd=12345;database=test;";

try
{
    conn.Open();
    cmd.Connection = conn;

    cmd.CommandText = "add_emp";
    cmd.CommandType = CommandType.StoredProcedure;

    cmd.Parameters.Add("?lname", "Jones");
    cmd.Parameters["?lname"].Direction = ParameterDirection.Input;

    cmd.Parameters.Add("?fname", "Tom");
    cmd.Parameters["?fname"].Direction = ParameterDirection.Input;

    cmd.Parameters.Add("?bday", DateTime.Parse("12/13/1977 2:17:36 PM"));
    cmd.Parameters["?bday"].Direction = ParameterDirection.Input;

    cmd.Parameters.Add("?empno", MySqlDbType.Int32);
    cmd.Parameters["?empno"].Direction = ParameterDirection.Output;

    cmd.ExecuteNonQuery();

    MessageBox.Show(cmd.Parameters["?empno"].Value);
}
catch (MySql.Data.MySqlClient.MySqlException ex)
```

```
{
    MessageBox.Show("Error " + ex.Number + " has occurred: " + ex.Message,
        "Error", MessageBoxButtons.OK, MessageBoxIcon.Error);
}
```

Once the stored procedure is called, the values of output parameters can be retrieved by using the `.Value` property of the `MySQLConnector.Parameters` collection.

## 18.2.5.4. Handling BLOB Data With Connector/NET

### 18.2.5.4.1. Introduction

One common use for MySQL is the storage of binary data in `BLOB` columns. MySQL supports four different BLOB datatypes: `TINYBLOB`, `BLOB`, `MEDIUMBLOB`, and `LONGBLOB`.

Data stored in a BLOB column can be accessed using Connector/NET and manipulated using client-side code. There are no special requirements for using Connector/NET with BLOB data.

Simple code examples will be presented within this section, and a full sample application can be found in the `Samples` directory of the Connector/NET installation.

### 18.2.5.4.2. Preparing the MySQL Server

The first step in using MySQL with BLOB data is to configure the server. Let's start by creating a table to be accessed. In my file tables, I usually have four columns: an `AUTO_INCREMENT` column of appropriate size (`UNSIGNED SMALLINT`) to serve as a primary key to identify the file, a `VARCHAR` column that stores the filename, an `UNSIGNED MEDIUMINT` column that stores the size of the file, and a `MEDIUMBLOB` column that stores the file itself. For this example, I will use the following table definition:

```
CREATE TABLE file(
file_id SMALLINT UNSIGNED AUTO_INCREMENT NOT NULL PRIMARY KEY,
file_name VARCHAR(64) NOT NULL,
file_size MEDIUMINT UNSIGNED NOT NULL,
file MEDIUMBLOB NOT NULL);
```

After creating a table, you may need to modify the `max_allowed_packet` system variable. This variable determines how large of a packet (i.e. a single row) can be sent to the MySQL server. By default, the server will only accept a maximum size of 1 meg from our client application. If you do not intend to exceed 1 meg, this should be fine. If you do intend to exceed 1 meg in your file transfers, this number has to be increased.

The `max_allowed_packet` option can be modified using MySQL Administrator's Startup Variables screen. Adjust the Maximum allowed option in the Memory section of the Networking tab to an appropriate setting. After adjusting the value, click the `Apply Changes` button and restart the server using the `Service Control` screen of MySQL Administrator. You can also adjust this value directly in the `my.cnf` file (add a line that reads `max_allowed_packet=xxM`), or use the `SET max_allowed_packet=xxM`; syntax from within MySQL.

Try to be conservative when setting `max_allowed_packet`, as transfers of BLOB data can take some time to complete. Try to set a value that will be adequate for your intended use and increase the value if necessary.

### 18.2.5.4.3. Writing a File to the Database

To write a file to a database we need to convert the file to a byte array, then use the byte array as a parameter to an `INSERT` query.

The following code opens a file using a `FileStream` object, reads it into a byte array, and inserts it into

the [file](#) table:

### Visual Basic Example

```

Dim conn As New MySqlConnection
Dim cmd As New MySqlCommand

Dim SQL As String

Dim FileSize As UInt32
Dim rawData() As Byte
Dim fs As FileStream

conn.ConnectionString = "server=127.0.0.1;" _
    & "uid=root;" _
    & "pwd=12345;" _
    & "database=test"

Try
    fs = New FileStream("c:\image.png", FileMode.Open, FileAccess.Read)
    FileSize = fs.Length

    rawData = New Byte(FileSize) {}
    fs.Read(rawData, 0, FileSize)
    fs.Close()

    conn.Open()

    SQL = "INSERT INTO file VALUES(NULL, ?FileName, ?FileSize, ?File)"

    cmd.Connection = conn
    cmd.CommandText = SQL
    cmd.Parameters.Add("?FileName", strFileName)
    cmd.Parameters.Add("?FileSize", FileSize)
    cmd.Parameters.Add("?File", rawData)

    cmd.ExecuteNonQuery()

    MessageBox.Show("File Inserted into database successfully!", _
        "Success!", MessageBoxButtons.OK, MessageBoxIcon.Asterisk)

    conn.Close()
Catch ex As Exception
    MessageBox.Show("There was an error: " & ex.Message, "Error", _
        MessageBoxButtons.OK, MessageBoxIcon.Error)
End Try

```

### C# Example

```

MySQL.Data.MySqlClient.MySqlConnection conn;
MySQL.Data.MySqlClient.MySqlCommand cmd;

conn = new MySQL.Data.MySqlClient.MySqlConnection();
cmd = new MySQL.Data.MySqlClient.MySqlCommand();

string SQL;
UInt32 FileSize;
byte[] rawData;
FileStream fs;

conn.ConnectionString = "server=127.0.0.1;uid=root;" +
    "pwd=12345;database=test;";

try
{
    fs = new FileStream(@"c:\image.png", FileMode.Open, FileAccess.Read);
    FileSize = fs.Length;

    rawData = new byte[FileSize];
    fs.Read(rawData, 0, FileSize);
    fs.Close();

    conn.Open();

    SQL = "INSERT INTO file VALUES(NULL, ?FileName, ?FileSize, ?File)";
}

```

```

cmd.Connection = conn;
cmd.CommandText = SQL;
cmd.Parameters.Add("?FileName", strFileName);
cmd.Parameters.Add("?FileSize", FileSize);
cmd.Parameters.Add("?File", rawData);

cmd.ExecuteNonQuery();

MessageBox.Show("File Inserted into database successfully!",
    "Success!", MessageBoxButtons.OK, MessageBoxIcon.Asterisk);

conn.Close();
}
catch (MySql.Data.MySqlClient.MySqlException ex)
{
    MessageBox.Show("Error " + ex.Number + " has occurred: " + ex.Message,
        "Error", MessageBoxButtons.OK, MessageBoxIcon.Error);
}

```

The `Read` method of the `FileStream` object is used to load the file into a byte array which is sized according to the `Length` property of the `FileStream` object.

After assigning the byte array as a parameter of the `MySqlCommand` object, the `ExecuteNonQuery` method is called and the BLOB is inserted into the `file` table.

#### 18.2.5.4.4. Reading a BLOB from the Database to a File on Disk

Once a file is loaded into the `file` table, we can use the `MySqlDataReader` class to retrieve it.

The following code retrieves a row from the `file` table, then loads the data into a `FileStream` object to be written to disk:

##### Visual Basic Example

```

Dim conn As New MySqlConnection
Dim cmd As New MySqlCommand
Dim myData As MySqlDataReader
Dim SQL As String
Dim rawData() As Byte
Dim FileSize As UInt32
Dim fs As FileStream

conn.ConnectionString = "server=127.0.0.1;" _
    & "uid=root;" _
    & "pwd=12345;" _
    & "database=test"

SQL = "SELECT file_name, file_size, file FROM file"

Try
    conn.Open()

    cmd.Connection = conn
    cmd.CommandText = SQL

    myData = cmd.ExecuteReader

    If Not myData.HasRows Then Throw New Exception("There are no BLOBs to save")

    myData.Read()

    FileSize = myData.GetUInt32(myData.GetOrdinal("file_size"))
    rawData = New Byte(FileSize) {}

    myData.GetBytes(myData.GetOrdinal("file"), 0, rawData, 0, FileSize)

    fs = New FileStream("C:\newfile.png", FileMode.OpenOrCreate, FileAccess.Write)
    fs.Write(rawData, 0, FileSize)
    fs.Close()

    MessageBox.Show("File successfully written to disk!", "Success!", MessageBoxButtons.OK, MessageBoxIcon.Asterisk)

    myData.Close()

```



```

    conn.Close()
Catch ex As Exception
    MessageBox.Show("There was an error: " & ex.Message, "Error", MessageBoxButtons.OK, MessageBoxIcon.Error)
End Try

```

### C# Example

```

MySQL.Data.MySqlClient.MySqlConnection conn;
MySQL.Data.MySqlClient.MySqlCommand cmd;
MySQL.Data.MySqlClient.MySqlDataReader myData;

conn = new MySQL.Data.MySqlClient.MySqlConnection();
cmd = new MySQL.Data.MySqlClient.MySqlCommand();

string SQL;
UInt32 FileSize;
byte[] rawData;
FileStream fs;

conn.ConnectionString = "server=127.0.0.1;uid=root;" +
    "pwd=12345;database=test;";

SQL = "SELECT file_name, file_size, file FROM file";

try
{
    conn.Open();

    cmd.Connection = conn;
    cmd.CommandText = SQL;

    myData = cmd.ExecuteReader();

    if (! myData.HasRows)
        throw new Exception("There are no BLOBs to save");

    myData.Read();

    FileSize = myData.GetUInt32(myData.GetOrdinal("file_size"));
    rawData = new byte[FileSize];

    myData.GetBytes(myData.GetOrdinal("file"), 0, rawData, 0, FileSize);

    fs = new FileStream(@"C:\newfile.png", FileMode.OpenOrCreate, FileAccess.Write);
    fs.Write(rawData, 0, FileSize);
    fs.Close();

    MessageBox.Show("File successfully written to disk!",
        "Success!", MessageBoxButtons.OK, MessageBoxIcon.Asterisk);

    myData.Close();
    conn.Close();
}
catch (MySQL.Data.MySqlClient.MySqlException ex)
{
    MessageBox.Show("Error " + ex.Number + " has occurred: " + ex.Message,
        "Error", MessageBoxButtons.OK, MessageBoxIcon.Error);
}

```

After connecting, the contents of the `file` table are loaded into a `MySqlDataReader` object. The `GetBytes` method of the `MySqlDataReader` is used to load the BLOB into a byte array, which is then written to disk using a `FileStream` object.

The `GetOrdinal` method of the `MySqlDataReader` can be used to determine the integer index of a named column. Use of the `GetOrdinal` method prevents errors if the column order of the `SELECT` query is changed.

## 18.2.5.5. Using Connector/NET with Crystal Reports

### 18.2.5.5.1. Introduction

Crystal Reports is a common tool used by Windows application developers to perform reporting and document generation. In this section we will show how to use Crystal Reports XI with MySQL and Connector/NET.

### 18.2.5.5.2. Creating a Data Source

When creating a report in Crystal Reports there are two options for accessing the MySQL data while designing your report.

The first option is to use Connector/ODBC as an ADO data source when designing your report. You will be able to browse your database and choose tables and fields using drag and drop to build your report. The disadvantage of this approach is that additional work must be performed within your application to produce a dataset that matches the one expected by your report.

The second option is to create a dataset in VB.NET and save it as XML. This XML file can then be used to design a report. This works quite well when displaying the report in your application, but is less versatile at design time because you must choose all relevant columns when creating the dataset. If you forget a column you must re-create the dataset before the column can be added to the report.

The following code can be used to create a dataset from a query and write it to disk:

#### Visual Basic Example

```
Dim myData As New DataSet
Dim conn As New MySqlConnection
Dim cmd As New MySqlCommand
Dim myAdapter As New MySqlDataAdapter

conn.ConnectionString = "server=127.0.0.1;" _
    & "uid=root;" _
    & "pwd=12345;" _
    & "database=world"

Try
    conn.Open()
    cmd.CommandText = "SELECT city.name AS cityName, city.population AS CityPopulation, " _
        & "country.name, country.population, country.continent " _
        & "FROM country, city ORDER BY country.continent, country.name"
    cmd.Connection = conn

    myAdapter.SelectCommand = cmd
    myAdapter.Fill(myData)

    myData.WriteXml("C:\dataset.xml", XmlWriteMode.WriteSchema)
Catch ex As Exception
    MessageBox.Show(ex.Message, "Report could not be created", MessageBoxButtons.OK, MessageBoxIcon.Error)
End Try
```

#### C# Example

```
DataSet myData = new DataSet();
MySql.Data.MySqlClient.MySqlConnection conn;
MySql.Data.MySqlClient.MySqlCommand cmd;
MySql.Data.MySqlClient.MySqlDataAdapter myAdapter;

conn = new MySql.Data.MySqlClient.MySqlConnection();
cmd = new MySql.Data.MySqlClient.MySqlCommand();
myAdapter = new MySql.Data.MySqlClient.MySqlDataAdapter();

conn.ConnectionString = "server=127.0.0.1;uid=root;" +
    "pwd=12345;database=test;";

try
{
    cmd.CommandText = "SELECT city.name AS cityName, city.population AS CityPopulation, " +
        "country.name, country.population, country.continent " +
        "FROM country, city ORDER BY country.continent, country.name";
    cmd.Connection = conn;
```

```

myAdapter.SelectCommand = cmd;
myAdapter.Fill(myData);

myData.WriteXml(@"C:\dataset.xml", XmlWriteMode.WriteSchema);
}
catch (MySQL.Data.MySqlClient.MySqlException ex)
{
    MessageBox.Show(ex.Message, "Report could not be created",
        MessageBoxButtons.OK, MessageBoxIcon.Error);
}

```

The resulting XML file can be used as an ADO.NET XML datasource when designing your report.

If you choose to design your reports using Connector/ODBC, it can be downloaded from [dev.mysql.com](http://dev.mysql.com/downloads/connector/odbc/3.51.html) [http://dev.mysql.com/downloads/connector/odbc/3.51.html].

### 18.2.5.5.3. Creating the Report

For most purposes the Standard Report wizard should help with the initial creation of a report. To start the wizard, open Crystal Reports and choose the New > Standard Report option from the File menu.

The wizard will first prompt you for a data source. If you are using Connector/ODBC as your data source, use the OLEDB provider for ODBC option from the OLE DB (ADO) tree instead of the ODBC (RDO) tree when choosing a data source. If using a saved dataset, choose the ADO.NET (XML) option and browse to your saved dataset.

The remainder of the report creation process is done automatically by the wizard.

After the report is created, choose the Report Options... entry of the File menu. Un-check the Save Data With Report option. This prevents saved data from interfering with the loading of data within our application.

### 18.2.5.5.4. Displaying the Report

To display a report we first populate a dataset with the data needed for the report, then load the report and bind it to the dataset. Finally we pass the report to the crViewer control for display to the user.

The following references are needed in a project that displays a report:

- CrystalDecisions.CrystalReports.Engine
- CrystalDecisions.ReportSource
- CrystalDecisions.Shared
- CrystalDecisions.Windows.Forms

The following code assumes that you created your report using a dataset saved using the code shown in [Section 18.2.5.5.2, "Creating a Data Source"](#), and have a crViewer control on your form named `myViewer`.

Visual Basic Example

```

Imports CrystalDecisions.CrystalReports.Engine
Imports System.Data
Imports MySQL.Data.MySqlClient

Dim myReport As New ReportDocument
Dim myData As New DataSet
Dim conn As New MySqlConnection
Dim cmd As New MySqlCommand

```

```

Dim myAdapter As New MySqlDataAdapter

conn.ConnectionString = _
    "server=127.0.0.1;" _
    & "uid=root;" _
    & "pwd=12345;" _
    & "database=test"

Try
    conn.Open()

    cmd.CommandText = "SELECT city.name AS cityName, city.population AS CityPopulation, " _
        & "country.name, country.population, country.continent " _
        & "FROM country, city ORDER BY country.continent, country.name"
    cmd.Connection = conn

    myAdapter.SelectCommand = cmd
    myAdapter.Fill(myData)

    myReport.Load(".\world_report.rpt")
    myReport.SetDataSource(myData)
    myViewer.ReportSource = myReport
Catch ex As Exception
    MessageBox.Show(ex.Message, "Report could not be created", MessageBoxButtons.OK, MessageBoxIcon.Error)
End Try

```

### C# Example

```

using CrystalDecisions.CrystalReports.Engine;
using System.Data;
using MySql.Data.MySqlClient;

ReportDocument myReport = new ReportDocument();
DataSet myData = new DataSet();
MySql.Data.MySqlClient.MySqlConnection conn;
MySql.Data.MySqlClient.MySqlCommand cmd;
MySql.Data.MySqlClient.MySqlDataAdapter myAdapter;

conn = new MySql.Data.MySqlClient.MySqlConnection();
cmd = new MySql.Data.MySqlClient.MySqlCommand();
myAdapter = new MySql.Data.MySqlClient.MySqlDataAdapter();

conn.ConnectionString = "server=127.0.0.1;uid=root;" +
    "pwd=12345;database=test;";

try
{
    cmd.CommandText = "SELECT city.name AS cityName, city.population AS CityPopulation, " +
        "country.name, country.population, country.continent " +
        "FROM country, city ORDER BY country.continent, country.name";
    cmd.Connection = conn;

    myAdapter.SelectCommand = cmd;
    myAdapter.Fill(myData);

    myReport.Load(@".\world_report.rpt");
    myReport.SetDataSource(myData);
    myViewer.ReportSource = myReport;
}
catch (MySql.Data.MySqlClient.MySqlException ex)
{
    MessageBox.Show(ex.Message, "Report could not be created",
        MessageBoxButtons.OK, MessageBoxIcon.Error);
}

```

A new dataset is generated using the same query used to generate the previously saved dataset. Once the dataset is filled, a ReportDocument is used to load the report file and bind it to the dataset. The ReportDocument is then passed as the ReportSource of the crViewer.

This same approach is taken when a report is created from a single table using Connector/ODBC. The dataset replaces the table used in the report and the report is displayed properly.

When a report is created from multiple tables using Connector/ODBC, a dataset with multiple tables must be created in our application. This allows each table in the report data source to be replaced with a

report in the dataset.

We populate a dataset with multiple tables by providing multiple `SELECT` statements in our MySQL-Command object. These `SELECT` statements are based on the SQL query shown in Crystal Reports in the Database menu's Show SQL Query option. Assume the following query:

```
SELECT `country`.`Name`, `country`.`Continent`, `country`.`Population`, `city`.`Name`, `city`.`Populat
FROM `world`.`country` `country` LEFT OUTER JOIN `world`.`city` `city` ON `country`.`Code`=`city`.`Cour
ORDER BY `country`.`Continent`, `country`.`Name`, `city`.`Name`
```

This query is converted to two `SELECT` queries and displayed with the following code:

### Visual Basic Example

```
Imports CrystalDecisions.CrystalReports.Engine
Imports System.Data
Imports MySql.Data.MySqlClient

Dim myReport As New ReportDocument
Dim myData As New DataSet
Dim conn As New MySqlConnection
Dim cmd As New MySqlCommand
Dim myAdapter As New MySqlDataAdapter

conn.ConnectionString = "server=127.0.0.1;" _
    & "uid=root;" _
    & "pwd=12345;" _
    & "database=world"

Try
    conn.Open()
    cmd.CommandText = "SELECT name, population, countrycode FROM city ORDER BY countrycode, name; " _
        & "SELECT name, population, code, continent FROM country ORDER BY continent, name"
    cmd.Connection = conn

    myAdapter.SelectCommand = cmd
    myAdapter.Fill(myData)

    myReport.Load(".\world_report.rpt")
    myReport.Database.Tables(0).SetDataSource(myData.Tables(0))
    myReport.Database.Tables(1).SetDataSource(myData.Tables(1))
    myViewer.ReportSource = myReport
Catch ex As Exception
    MessageBox.Show(ex.Message, "Report could not be created", MessageBoxButtons.OK, MessageBoxIcon.Err
End Try
```

### C# Example

```
using CrystalDecisions.CrystalReports.Engine;
using System.Data;
using MySql.Data.MySqlClient;

ReportDocument myReport = new ReportDocument();
DataSet myData = new DataSet();
MySql.Data.MySqlClient.MySqlConnection conn;
MySql.Data.MySqlClient.MySqlCommand cmd;
MySql.Data.MySqlClient.MySqlDataAdapter myAdapter;

conn = new MySql.Data.MySqlClient.MySqlConnection();
cmd = new MySql.Data.MySqlClient.MySqlCommand();
myAdapter = new MySql.Data.MySqlClient.MySqlDataAdapter();

conn.ConnectionString = "server=127.0.0.1;uid=root;" +
    "pwd=12345;database=test;";

try
{
    cmd.CommandText = "SELECT name, population, countrycode FROM city ORDER " +
        "BY countrycode, name; SELECT name, population, code, continent FROM " +
        "country ORDER BY continent, name";
    cmd.Connection = conn;

    myAdapter.SelectCommand = cmd;
    myAdapter.Fill(myData);
}
```

```
myReport.Load(@".\world_report.rpt");
myReport.Database.Tables(0).SetDataSource(myData.Tables(0));
myReport.Database.Tables(1).SetDataSource(myData.Tables(1));
myViewer.ReportSource = myReport;
}
catch (MySQL.Data.MySqlClient.MySqlException ex)
{
    MessageBox.Show(ex.Message, "Report could not be created",
        MessageBoxButtons.OK, MessageBoxIcon.Error);
}
```

It is important to order the [SELECT](#) queries in alphabetical order, as this is the order the report will expect its source tables to be in. One `SetDataSource` statement is needed for each table in the report.

This approach can cause performance problems because Crystal Reports must bind the tables together on the client-side, which will be slower than using a pre-saved dataset.

## 18.2.5.6. Handling Date and Time Information in Connector/NET

### 18.2.5.6.1. Introduction

MySQL and the .NET languages handle date and time information differently, with MySQL allowing dates that cannot be represented by a .NET data type, such as '0000-00-00 00:00:00'. These differences can cause problems if not properly handled.

In this section we will demonstrate how to properly handle date and time information when using Connector/NET.

### 18.2.5.6.2. Problems when Using Invalid Dates

The differences in date handling can cause problems for developers who use invalid dates. Invalid MySQL dates cannot be loaded into native .NET `DateTime` objects, including `NULL` dates.

Because of this issue, .NET `DataSet` objects cannot be populated by the `Fill` method of the `MySqlDataAdapter` class as invalid dates will cause a `System.ArgumentOutOfRangeException` exception to occur.

### 18.2.5.6.3. Restricting Invalid Dates

The best solution to the date problem is to restrict users from entering invalid dates. This can be done on either the client or the server side.

Restricting invalid dates on the client side is as simple as always using the .NET `DateTime` class to handle dates. The `DateTime` class will only allow valid dates, ensuring that the values in your database are also valid. The disadvantage of this is that it is not useful in a mixed environment where .NET and non .NET code are used to manipulate the database, as each application must perform its own date validation.

Users of MySQL 5.0.2 and higher can use the new `traditional` SQL mode to restrict invalid date values. For information on using the `traditional` SQL mode, see [Section 5.2.5, “SQL Modes”](#).

### 18.2.5.6.4. Handling Invalid Dates

Although it is strongly recommended that you avoid the use of invalid dates within your .NET application, it is possible to use invalid dates by means of the `MySqlDateTime` datatype.

The `MySqlDateTime` datatype supports the same date values that are supported by the MySQL server. The default behavior of Connector/NET is to return a .NET `DateTime` object for valid date values, and return an error for invalid dates. This default can be modified to cause Connector/NET to return

`MySqlDateTime` objects for invalid dates.

To instruct Connector/NET to return a `MySqlDateTime` object for invalid dates, add the following line to your connection string:

```
Allow Zero Datetime=True
```

Please note that the use of the `MySqlDateTime` class can still be problematic. The following are some known issues:

1. Data binding for invalid dates can still cause errors (zero dates like 0000-00-00 do not seem to have this problem).
2. The `ToString` method return a date formatted in the standard MySQL format (for example, 2005-02-23 08:50:25). This differs from the `ToString` behavior of the .NET `DateTime` class.
3. The `MySqlDateTime` class supports NULL dates, while the .NET `DateTime` class does not. This can cause errors when trying to convert a `MySQLDateTime` to a `DateTime` if you do not check for NULL first.

Because of the known issues, the best recommendation is still to use only valid dates in your application.

#### 18.2.5.6.5. Handling NULL Dates

The .NET `DateTime` datatype cannot handle `NULL` values. As such, when assigning values from a query to a `DateTime` variable, you must first check whether the value is in fact `NULL`.

When using a `MySqlDataReader`, use the `.IsDBNull` method to check whether a value is `NULL` before making the assignment:

Visual Basic Example

```
If Not myReader.IsDBNull(myReader.GetOrdinal("mytime")) Then
    myTime = myReader.GetDateTime(myReader.GetOrdinal("mytime"))
Else
    myTime = DateTime.MinValue
End If
```

C# Example

```
if (! myReader.IsDBNull(myReader.GetOrdinal("mytime")))
    myTime = myReader.GetDateTime(myReader.GetOrdinal("mytime"));
else
    myTime = DateTime.MinValue;
```

`NULL` values will work in a dataset and can be bound to form controls without special handling.

## 18.2.6. Connector/NET Support

The developers of Connector/NET greatly value the input of our users in the software development process. If you find Connector/NET lacking some feature important to you, or if you discover a bug and need to file a bug report, please use the instructions in [Section 1.8, "How to Report Bugs or Problems"](#).

### 18.2.6.1. Connector/NET Community Support

- Community support for Connector/NET can be found through the forums at <http://forums.mysql.com>.
- Community support for Connector/NET can also be found through the mailing lists at <http://lists.mysql.com>.
- Paid support is available from MySQL AB. Additional information is available at <http://www.mysql.com/support/>.

### 18.2.6.2. How to report Connector/NET Problems or Bugs

If you encounter difficulties or problems with Connector/NET, contact the Connector/NET community [Section 18.2.6.1, “Connector/NET Community Support”](#).

You should first try to execute the same SQL statements and commands from the `mysql` client program or from `admindemo`. This helps you determine whether the error is in Connector/NET or MySQL.

If reporting a problem, you should ideally include the following information with the email:

- Operating system and version
- Connector/NET version
- MySQL server version
- Copies of error messages or other unexpected output
- Simple reproducible sample

Remember that the more information you can supply to us, the more likely it is that we can fix the problem.

If you believe the problem to be a bug, then you must report the bug through <http://bugs.mysql.com/>.

### 18.2.6.3. Connector/NET Change History

The Connector/NET Change History (Changelog) is located with the main Changelog for MySQL. See [Section D.7, “MySQL Connector/NET Change History”](#).

## 18.3. MySQL Connector/J

MySQL provides connectivity for client applications developed in the Java programming language via a JDBC driver, which is called MySQL Connector/J.

MySQL Connector/J is a JDBC-3.0 Type 4 driver, which means that is pure Java, implements version 3.0 of the JDBC specification, and communicates directly with the MySQL server using the MySQL protocol.

Although JDBC is useful by itself, we would hope that if you are not familiar with JDBC that after reading the first few sections of this manual, that you would avoid using naked JDBC for all but the most trivial problems and consider using one of the popular persistence frameworks such as [Hibernate](http://www.hibernate.org/) [http://www.hibernate.org/], [Spring's JDBC templates](http://www.springframework.org/) [http://www.springframework.org/] or [Ibatis SQL Maps](http://ibatis.apache.org/) [http://ibatis.apache.org/] to do the majority of repetitive work and heavier lifting that is sometimes required with JDBC.



This section is not designed to be a complete JDBC tutorial. If you need more information about using JDBC you might be interested in the following online tutorials that are more in-depth than the information presented here:

- [JDBC Basics](http://java.sun.com/docs/books/tutorial/jdbc/basics/index.html) [http://java.sun.com/docs/books/tutorial/jdbc/basics/index.html] — A tutorial from Sun covering beginner topics in JDBC
- [JDBC Short Course](http://java.sun.com/developer/onlineTraining/Database/JDBCShortCourse/index.html) [http://java.sun.com/developer/onlineTraining/Database/JDBCShortCourse/index.html] — A more in-depth tutorial from Sun and JGuru

## 18.3.1. Connector/J Versions

There are currently three version of MySQL Connector/J available:

- Connector/J 3.0 provides core functionality and was designed with connectivity to MySQL 3.x or MySQL 4.1 servers, although it will provide basic compatibility with later versions of MySQL. Connector/J 3.0 does not support server-side prepared statements, and does not support any of the features in versions of MySQL later than 4.1.
- Connector/J 3.1 was designed for connectivity to MySQL 4.1 and MySQL 5.0 servers and provides support for all the functionality in MySQL 5.0 except distributed transaction (XA) support.
- Connector/J 5.0 provides support for all the functionality offered by Connector/J 3.1 and includes distributed transaction (XA) support.

The current recommended version for Connector/J is 5.0. This guide covers all three connector versions, with specific notes given where a setting applies to a specific option.

### 18.3.1.1. Java Versions Supported

MySQL Connector/J supports Java-2 JVMs, including:

- JDK 1.2.x
- JDK 1.3.x
- JDK 1.4.x
- JDK 1.5.x

If you are building Connector/J from source using the source distribution (see [Section 18.3.2.4, “Installing from the Development Source Tree”](#)) then you must use JDK 1.4.x or newer to compile the Connector package.

MySQL Connector/J does not support JDK-1.1.x or JDK-1.0.x

Because of the implementation of `java.sql.Savepoint`, Connector/J 3.1.0 and newer will not run on JDKs older than 1.4 unless the class verifier is turned off (by setting the `-Xverify:none` option to the Java runtime). This is because the class verifier will try to load the class definition for `java.sql.Savepoint` even though it is not accessed by the driver unless you actually use savepoint functionality.

Caching functionality provided by Connector/J 3.1.0 or newer is also not available on JVMs older than 1.4.x, as it relies on `java.util.LinkedHashMap` which was first available in JDK-1.4.0.

## 18.3.2. Installing Connector/J

You can install the Connector/J package using two methods, using either the binary or source distribution. The binary distribution provides the easiest methods for installation; the source distribution enables you to customize your installation further. With either solution, you must

### 18.3.2.1. Installing Connector/J from a Binary Distribution

The easiest method of installation is to use the binary distribution of the Connector/J package. The binary distribution is available either as a Tar/Gzip or Zip file which you must extract to a suitable location and then optionally make the information about the package available by changing your `CLASSPATH` (see Section 18.3.2.2, “Installing the Driver and Configuring the `CLASSPATH`”).

MySQL Connector/J is distributed as a `.zip` or `.tar.gz` archive containing the sources, the class files, and the JAR archive named `mysql-connector-java-[version]-bin.jar`, and starting with Connector/J 3.1.8 a debug build of the driver in a file named `mysql-connector-java-[version]-bin-g.jar`.

Starting with Connector/J 3.1.9, the `.class` files that constitute the JAR files are only included as part of the driver JAR file.

You should not use the debug build of the driver unless instructed to do so when reporting a problem or bug to MySQL AB, as it is not designed to be run in production environments, and will have adverse performance impact when used. The debug binary also depends on the Aspect/J runtime library, which is located in the `src/lib/aspectjrt.jar` file that comes with the Connector/J distribution.

You will need to use the appropriate graphical or command-line utility to un-archive the distribution (for example, WinZip for the `.zip` archive, and `tar` for the `.tar.gz` archive). Because there are potentially long filenames in the distribution, we use the GNU tar archive format. You will need to use GNU tar (or an application that understands the GNU tar archive format) to unpack the `.tar.gz` variant of the distribution.

### 18.3.2.2. Installing the Driver and Configuring the `CLASSPATH`

Once you have extracted the distribution archive, you can install the driver by placing `mysql-connector-java-[version]-bin.jar` in your classpath, either by adding the full path to it to your `CLASSPATH` environment variable, or by directly specifying it with the command line switch `-cp` when starting your JVM.

If you are going to use the driver with the JDBC DriverManager, you would use `com.mysql.jdbc.Driver` as the class that implements `java.sql.Driver`.

You can set the `CLASSPATH` environment variable under UNIX, Linux or Mac OS X either locally for a user within their `.profile`, `.login` or other login file. You can also set it globally by editing the global `/etc/profile` file.

For example, under a C shell (`csh`, `tcsh`) you would add the Connector/J driver to your `CLASSPATH` using the following:

```
shell> setenv CLASSPATH /path/to/mysql-connector-java-[version]-bin.jar:$CLASSPATH
```

Or with a Bourne-compatible shell (`sh`, `ksh`, `bash`):

```
export set CLASSPATH=/path/to/mysql-connector-java-[version]-bin.jar:$CLASSPATH
```

Within Windows 2000, Windows XP and Windows Server 2003, you must set the environment variable through the System control panel.

If you want to use MySQL Connector/J with an application server such as Tomcat or JBoss, you will have to read your vendor's documentation for more information on how to configure third-party class libraries, as most application servers ignore the `CLASSPATH` environment variable. For configuration examples for some J2EE application servers, see [Section 18.3.5.2, “Using Connector/J with J2EE and Other Java Frameworks”](#). However, the authoritative source for JDBC connection pool configuration information for your particular application server is the documentation for that application server.

If you are developing servlets or JSPs, and your application server is J2EE-compliant, you can put the driver's .jar file in the `WEB-INF/lib` subdirectory of your webapp, as this is a standard location for third party class libraries in J2EE web applications.

You can also use the `MysqlDataSource` or `MysqlConnectionPoolDataSource` classes in the `com.mysql.jdbc.jdbc2.optional` package, if your J2EE application server supports or requires them. Starting with Connector/J 5.0.0, the `javax.sql.XADataSource` interface is implemented via the `com.mysql.jdbc.jdbc2.optional.MysqlXADataSource` class, which supports XA distributed transactions when used in combination with MySQL server version 5.0.

The various `MysqlDataSource` classes support the following parameters (through standard set mutators):

- user
- password
- serverName (see the previous section about fail-over hosts)
- databaseName
- port

### 18.3.2.3. Upgrading from an Older Version

MySQL AB tries to keep the upgrade process as easy as possible, however as is the case with any software, sometimes changes need to be made in new versions to support new features, improve existing functionality, or comply with new standards.

This section has information about what users who are upgrading from one version of Connector/J to another (or to a new version of the MySQL server, with respect to JDBC functionality) should be aware of.

#### 18.3.2.3.1. Upgrading from MySQL Connector/J 3.0 to 3.1

Connector/J 3.1 is designed to be backward-compatible with Connector/J 3.0 as much as possible. Major changes are isolated to new functionality exposed in MySQL-4.1 and newer, which includes Unicode character sets, server-side prepared statements, `SQLState` codes returned in error messages by the server and various performance enhancements that can be enabled or disabled via configuration properties.

- **Unicode Character Sets** — See the next section, as well as [Chapter 10, \*Character Set Support\*](#), for information on this new feature of MySQL. If you have something misconfigured, it will usually show up as an error with a message similar to `Illegal mix of collations`.
- **Server-side Prepared Statements** — Connector/J 3.1 will automatically detect and use server-side prepared statements when they are available (MySQL server version 4.1.0 and newer).

Starting with version 3.1.7, the driver scans SQL you are preparing via all variants of `Connec-`

`tion.prepareStatement()` to determine if it is a supported type of statement to prepare on the server side, and if it is not supported by the server, it instead prepares it as a client-side emulated prepared statement. You can disable this feature by passing `emulateUnsupportedPstmts=false` in your JDBC URL.

If your application encounters issues with server-side prepared statements, you can revert to the older client-side emulated prepared statement code that is still presently used for MySQL servers older than 4.1.0 with the connection property `useServerPrepStmts=false`

- **Datetimes** with all-zero components (`0000-00-00 ...`) — These values can not be represented reliably in Java. Connector/J 3.0.x always converted them to NULL when being read from a ResultSet.

Connector/J 3.1 throws an exception by default when these values are encountered as this is the most correct behavior according to the JDBC and SQL standards. This behavior can be modified using the `zeroDateTimeBehavior` configuration property. The allowable values are:

- `exception` (the default), which throws an `SQLException` with an `SQLState` of `S1009`.
- `convertToNull`, which returns `NULL` instead of the date.
- `round`, which rounds the date to the nearest closest value which is `0001-01-01`.

Starting with Connector/J 3.1.7, `ResultSet.getString()` can be decoupled from this behavior via `noDatetimeStringSync=true` (the default value is `false`) so that you can get retrieve the unaltered all-zero value as a `String`. It should be noted that this also precludes using any time zone conversions, therefore the driver will not allow you to enable `noDatetimeStringSync` and `useTimeZone` at the same time.

- **New SQLState Codes** — Connector/J 3.1 uses SQL:1999 SQLState codes returned by the MySQL server (if supported), which are different from the legacy X/Open state codes that Connector/J 3.0 uses. If connected to a MySQL server older than MySQL-4.1.0 (the oldest version to return SQLStates as part of the error code), the driver will use a built-in mapping. You can revert to the old mapping by using the configuration property `useSqlStateCodes=false`.
- **ResultSet.getString()** — Calling `ResultSet.getString()` on a BLOB column will now return the address of the `byte[]` array that represents it, instead of a `String` representation of the BLOB. BLOBs have no character set, so they can't be converted to `java.lang.Strings` without data loss or corruption.

To store strings in MySQL with LOB behavior, use one of the TEXT types, which the driver will treat as a `java.sql.Clob`.

- **Debug builds** — Starting with Connector/J 3.1.8 a debug build of the driver in a file named `mysql-connector-java-[version]-bin-g.jar` is shipped alongside the normal binary jar file that is named `mysql-connector-java-[version]-bin.jar`.

Starting with Connector/J 3.1.9, we don't ship the `.class` files unbundled, they are only available in the JAR archives that ship with the driver.

You should not use the debug build of the driver unless instructed to do so when reporting a problem or bug to MySQL AB, as it is not designed to be run in production environments, and will have adverse performance impact when used. The debug binary also depends on the Aspect/J runtime library, which is located in the `src/lib/aspectjrt.jar` file that comes with the Connector/J distribution.

### 18.3.2.3.2. JDBC-Specific Issues When Upgrading to MySQL Server 4.1 or Newer

- *Using the UTF-8 Character Encoding* - Prior to MySQL server version 4.1, the UTF-8 character encoding was not supported by the server, however the JDBC driver could use it, allowing storage of multiple character sets in latin1 tables on the server.

Starting with MySQL-4.1, this functionality is deprecated. If you have applications that rely on this functionality, and can not upgrade them to use the official Unicode character support in MySQL server version 4.1 or newer, you should add the following property to your connection URL:

```
useOldUTF8Behavior=true
```

- *Server-side Prepared Statements* - Connector/J 3.1 will automatically detect and use server-side prepared statements when they are available (MySQL server version 4.1.0 and newer). If your application encounters issues with server-side prepared statements, you can revert to the older client-side emulated prepared statement code that is still presently used for MySQL servers older than 4.1.0 with the following connection property:

```
useServerPrepStmts=false
```

### 18.3.2.4. Installing from the Development Source Tree

**Caution.** You should read this section only if you are interested in helping us test our new code. If you just want to get MySQL Connector/J up and running on your system, you should use a standard release distribution.

To install MySQL Connector/J from the development source tree, make sure that you have the following prerequisites:

- Subversion, to check out the sources from our repository (available from <http://subversion.tigris.org/>).
- Apache Ant version 1.6 or newer (available from <http://ant.apache.org/>).
- JDK-1.4.2 or later. Although MySQL Connector/J can be installed on older JDKs, to compile it from source you must have at least JDK-1.4.2.

The Subversion source code repository for MySQL Connector/J is located at <http://svn.mysql.com/svnpublic/connector-j>. In general, you should not check out the entire repository because it contains every branch and tag for MySQL Connector/J and is quite large.

To check out and compile a specific branch of MySQL Connector/J, follow these steps:

1. At the time of this writing, there are three active branches of Connector/J: `branch_3_0`, `branch_3_1` and `branch_5_0`. Check out the latest code from the branch that you want with the following command (replacing `[major]` and `[minor]` with appropriate version numbers):

```
shell> svn co »  
http://svn.mysql.com/svnpublic/connector-j/branches/branch\_\[major\]\_\[minor\]/connector-j
```

This creates a `connector-j` subdirectory in the current directory that contains the latest sources for the requested branch.

2. Change location to the `connector-j` directory to make it your current working directory:

```
shell> cd connector-j
```

3. Issue the following command to compile the driver and create a `.jar` file suitable for installation:

```
shell> ant dist
```

This creates a `build` directory in the current directory, where all build output will go. A directory is created in the `build` directory that includes the version number of the sources you are building from. This directory contains the sources, compiled `.class` files, and a `.jar` file suitable for deployment. For other possible targets, including ones that will create a fully packaged distribution, issue the following command:

```
shell> ant --projecthelp
```

4. A newly created `.jar` file containing the JDBC driver will be placed in the directory `build/mysql-connector-java-[version]`.

Install the newly created JDBC driver as you would a binary `.jar` file that you download from MySQL by following the instructions in [Section 18.3.2.2, “Installing the Driver and Configuring the CLASSPATH”](#).

### 18.3.3. Connector/J Examples

Examples of using Connector/J are located throughout this document, this section provides a summary and links to these examples.

- [Example 18.1, “Obtaining a connection from the DriverManager”](#)
- [Example 18.2, “Using java.sql.Statement to execute a SELECT query”](#)
- [Example 18.3, “Stored Procedures”](#)
- [Example 18.4, “Using Connection.prepareStatement\(\)”](#)
- [Example 18.5, “Registering output parameters”](#)
- [Example 18.6, “Setting CallableStatement input parameters”](#)
- [Example 18.7, “Retrieving results and output parameter values”](#)
- [Example 18.8, “Retrieving AUTO\\_INCREMENT column values using Statement.getGeneratedKeys\(\)”](#)
- [Example 18.9, “Retrieving AUTO\\_INCREMENT column values using SELECT LAST\\_INSERT\\_ID\(\)”](#)
- [Example 18.10, “Retrieving AUTO\\_INCREMENT column values in Updatable ResultSets”](#)
- [Example 18.11, “Using a connection pool with a J2EE application server”](#)
- [Example 18.12, “Example of transaction with retry logic”](#)

### 18.3.4. Connector/J (JDBC) Reference

This section of the manual contains reference material for MySQL Connector/J, some of which is automatically generated during the Connector/J build process.

### 18.3.4.1. Driver/Datasource Class Names, URL Syntax and Configuration Properties for Connector/J

The name of the class that implements `java.sql.Driver` in MySQL Connector/J is `com.mysql.jdbc.Driver`. The `org.gjt.mm.mysql.Driver` class name is also usable to remain backward-compatible with MM.MySQL. You should use this class name when registering the driver, or when otherwise configuring software to use MySQL Connector/J.

The JDBC URL format for MySQL Connector/J is as follows, with items in square brackets ([, ]) being optional:

```
jdbc:mysql://[host][,failoverhost...][:port]/[database] »
[?propertyName1[=propertyValue1][&propertyName2[=propertyValue2]...
```

If the hostname is not specified, it defaults to 127.0.0.1. If the port is not specified, it defaults to 3306, the default port number for MySQL servers.

```
jdbc:mysql://[host:port],[host:port].../[database] »
[?propertyName1[=propertyValue1][&propertyName2[=propertyValue2]...
```

If the database is not specified, the connection will be made with no default database. In this case, you will need to either call the `setCatalog()` method on the `Connection` instance or fully-specify table names using the database name (i.e. `SELECT dbname.tablename.colname FROM dbname.tablename...`) in your SQL. Not specifying the database to use upon connection is generally only useful when building tools that work with multiple databases, such as GUI database managers.

MySQL Connector/J has fail-over support. This allows the driver to fail-over to any number of slave hosts and still perform read-only queries. Fail-over only happens when the connection is in an `autoCommit(true)` state, because fail-over can not happen reliably when a transaction is in progress. Most application servers and connection pools set `autoCommit` to `true` at the end of every transaction/connection use.

The fail-over functionality has the following behavior:

- If the URL property `autoReconnect` is `false`: Failover only happens at connection initialization, and failback occurs when the driver determines that the first host has become available again.
- If the URL property `autoReconnect` is `true`: Failover happens when the driver determines that the connection has failed (before *every* query), and falls back to the first host when it determines that the host has become available again (after `queriesBeforeRetryMaster` queries have been issued).

In either case, whenever you are connected to a "failed-over" server, the connection will be set to read-only state, so queries that would modify data will have exceptions thrown (the query will **never** be processed by the MySQL server).

Configuration properties define how Connector/J will make a connection to a MySQL server. Unless otherwise noted, properties can be set for a `DataSource` object or for a `Connection` object.

Configuration Properties can be set in one of the following ways:

- Using the `set*()` methods on MySQL implementations of `java.sql.DataSource` (which is the preferred method when using implementations of `java.sql.DataSource`):
  - `com.mysql.jdbc.jdbc2.optional.MysqlDataSource`

- `com.mysql.jdbc.jdbc2.optional.MysqlConnectionPoolDataSource`
- As a key/value pair in the `java.util.Properties` instance passed to `DriverManager.getConnection()` or `Driver.connect()`
- As a JDBC URL parameter in the URL given to `java.sql.DriverManager.getConnection()`, `java.sql.Driver.connect()` or the MySQL implementations of the `javax.sql.DataSource.setURL()` method.

**Note.** If the mechanism you use to configure a JDBC URL is XML-based, you will need to use the XML character literal `&#amp;`; to separate configuration parameters, as the ampersand is a reserved character for XML.

The properties are listed in the following tables.

#### Connection/Authentication.

Property Name	Definition	Default Value	Since Version
user	The user to connect as		all
password	The password to use when connecting		all
socketFactory	The name of the class that the driver should use for creating socket connections to the server. This class must implement the interface <code>com.mysql.jdbc.SocketFactory</code> and have public no-args constructor.	<code>com.mysql.jdbc.StandardSocketFactory</code>	3.0.3
connectTimeout	Timeout for socket connect (in milliseconds), with 0 being no timeout. Only works on JDK-1.4 or newer. Defaults to 0.	0	3.0.1
socketTimeout	Timeout on network socket operations (0, the default means no timeout).	0	3.0.1
useConfigs	Load the comma-delimited list of configuration properties before parsing the URL or applying user-specified properties. See <a href="#">Section 18.3.4.1, “Driver/Datasource Class Names, URL Syntax and Configuration Properties for Connector/J”</a>		3.1.5
interactiveClient	Set the <code>CLIENT_INTERACTIVE</code> flag, which tells MySQL to timeout connections based on <code>INTERACTIVE_TIMEOUT</code> instead of <code>WAIT_TIMEOUT</code>	false	3.1.0
propertiesTransform	An implementation of <code>com.mysql.jdbc.ConnectionPropertiesTransform</code> that the driver will use to modify URL properties passed to the driver before attempting a connection		3.1.4
useCompression	Use zlib compression when communicating with the server (true/false)? Defaults to <code>false</code> .	false	3.0.17

#### High Availability and Clustering.

Property Name	Definition	Default	Since
---------------	------------	---------	-------



		Value	Version
autoReconnect	Should the driver try to re-establish stale and/or dead connections? If enabled the driver will throw an exception for a queries issued on a stale or dead connection, which belong to the current transaction, but will attempt reconnect before the next query issued on the connection in a new transaction. The use of this feature is not recommended, because it has side effects related to session state and data consistency when applications don'thandle SQLExceptions properly, and is only designed to be used when you are unable to configure your application to handle SQLExceptions resulting from dead andstale connections properly. Alternatively, investigate setting the MySQL server variable "wait_timeout"to some high value rather than the default of 8 hours.	false	1.1
autoReconnectForPools	Use a reconnection strategy appropriate for connection pools (defaults to 'false')	false	3.1.3
failOverReadOnly	When failing over in autoReconnect mode, should the connection be set to 'read-only'?	true	3.0.12
reconnectAtTxEnd	If autoReconnect is set to true, should the driver attempt reconnectionsat the end of every transaction?	false	3.0.10
roundRobinLoadBalance	When autoReconnect is enabled, and failover-ReadOnly is false, should we pick hosts to connect to on a round-robin basis?	false	3.1.2
queriesBeforeRetryMaster	Number of queries to issue before falling back to master when failed over (when using multi-host failover). Whichever condition is met first, 'queriesBeforeRetryMaster' or 'secondsBeforeRetryMaster' will cause an attempt to be made to reconnect to the master. Defaults to 50.	50	3.0.2
secondsBeforeRetryMaster	How long should the driver wait, when failed over, before attempting to reconnect to the master server? Whichever condition is met first, 'queriesBeforeRetryMaster' or 'secondsBeforeRetryMaster' will cause an attempt to be made to reconnect to the master. Time in seconds, defaults to 30	30	3.0.2
resourceId	A globally unique name that identifies the resource that this datasource or connection is connected to, used for <code>XAResource.isSameRM()</code> when the driver can't determine this value based on host-names used in the URL		5.0.1

**Security.**

Property Name	Definition	Default Value	Since Version
allowMultiQueries	Allow the use of ';' to delimit multiple queries during one statement (true/false, defaults to 'false')	false	3.1.1
useSSL	Use SSL when communicating with the server	false	3.0.2

	(true/false), defaults to 'false'		
requireSSL	Require SSL connection if useSSL=true? (defaults to 'false').	false	3.1.0
allowUrlInLocalInfile	Should the driver allow URLs in 'LOAD DATA LOCAL INFILE' statements?	false	3.1.4
paranoid	Take measures to prevent exposure sensitive information in error messages and clear data structures holding sensitive data when possible? (defaults to 'false')	false	3.0.1

**Performance Extensions.**

Property Name	Definition	Default Value	Since Version
metadataCacheSize	The number of queries to cacheResultSetMetadata for if cacheResultSetMetaData is set to 'true' (default 50)	50	3.1.1
prepStmtCacheSize	If prepared statement caching is enabled, how many prepared statements should be cached?	25	3.0.10
prepStmtCacheSqlLimit	If prepared statement caching is enabled, what's the largest SQL the driver will cache the parsing for?	256	3.0.10
useCursorFetch	If connected to MySQL > 5.0.2, and <code>setFetchSize() &gt; 0</code> on a statement, should that statement use cursor-based fetching to retrieve rows?	false	5.0.0
blobSendChunkSize	Chunk to use when sending BLOB/CLOBs via ServerPreparedStatements	1048576	3.1.9
cacheCallableStmts	Should the driver cache the parsing stage of CallableStatements	false	3.1.2
cachePrepStmts	Should the driver cache the parsing stage of PreparedStatements of client-side prepared statements, the check for suitability of server-side prepared and server-side prepared statements themselves?	false	3.0.10
cacheResultSetMetadata	Should the driver cache ResultSetMetaData for Statements and PreparedStatements? (Req. JDK-1.4+, true/false, default 'false')	false	3.1.1
cacheServerConfiguration	Should the driver cache the results of <code>SHOW VARIABLES</code> and <code>SHOW COLLATION</code> on a per-URL basis?	false	3.1.5
defaultFetchSize	The driver will call <code>setFetchSize(n)</code> with this value on all newly-created Statements	0	3.1.9
dontTrackOpenResources	The JDBC specification requires the driver to automatically track and close resources, however if your application doesn't do a good job of explicitly calling <code>close()</code> on statements or result sets, this can cause memory leakage. Setting this property to true relaxes this constraint, and can be more memory efficient for some applications.	false	3.1.7
dynamicCalendars	Should the driver retrieve the default calendar	false	3.1.5

	when required, or cache it per connection/session?		
elideSetAutoCommits	If using MySQL-4.1 or newer, should the driver only issue 'set autocommit=n' queries when the server's state doesn't match the requested state by <code>Connection.setAutoCommit(boolean)</code> ?	false	3.1.3
holdResultsOpenOverStatementClose	Should the driver close result sets on <code>Statement.close()</code> as required by the JDBC specification?	false	3.1.7
locatorFetchBufferSize	If 'emulateLocators' is configured to 'true', what size buffer should be used when fetching BLOB data for <code>getBinaryInputStream</code> ?	1048576	3.2.1
rewriteBatchedStatements	Should the driver use multiqueries (irregardless of the setting of <code>allowMultiQueries</code> ) as well as rewriting of prepared statements for <code>INSERT</code> into multi-value inserts when <code>executeBatch()</code> is called? Notice that this has the potential for SQL injection if using plain <code>java.sql.Statements</code> and your code doesn't sanitize input correctly. Notice that for prepared statements, server-side prepared statements can not currently take advantage of this rewrite option, and that if you don't specify stream lengths when using <code>PreparedStatement.set*Stream()</code> , the driver won't be able to determine the optimum number of parameters per batch and you might receive an error from the driver that the resultant packet is too large. <code>Statement.getGeneratedKeys()</code> for these rewritten statements only works when the entire batch includes <code>INSERT</code> statements.	false	3.1.13
useFastIntParsing	Use internal String->Integer conversion routines to avoid excessive object creation?	true	3.1.4
useJvmCharsetConverters	Always use the character encoding routines built into the JVM, rather than using lookup tables for single-byte character sets? (The default of "true" for this is appropriate for newer JVMs)	true	5.0.1
useLocalSessionState	Should the driver refer to the internal values of autocommit and transaction isolation that are set by <code>Connection.setAutoCommit()</code> and <code>Connection.setTransactionIsolation()</code> , rather than querying the database?	false	3.1.7
useReadAheadInput	Use newer, optimized non-blocking, buffered input stream when reading from the server?	true	3.1.5

**Debugging/Profiling.**

Property Name	Definition	Default Value	Since Version
logger	The name of a class that implements 'com.mysql.jdbc.log.Log' that will be used to log messages to. (default is 'com.mysql.jdbc.log.StandardLogger', which logs to <code>STDERR</code> )	com.mysql.jdbc.log.StandardLogger	3.1.1
profileSQL	Trace queries and their execution/fetch times to the	false	3.1.0

	configured logger (true/false) defaults to 'false'		
reportMetricsIntervalMillis	If 'gatherPerfMetrics' is enabled, how often should they be logged (in ms)?	30000	3.1.2
maxQuerySizeToLog	Controls the maximum length/size of a query that will get logged when profiling or tracing	2048	3.1.3
packetDebugBufferSize	The maximum number of packets to retain when 'enablePacketDebug' is true	20	3.1.3
slowQueryThresholdMillis	If 'logSlowQueries' is enabled, how long should a query (in ms) before it is logged as 'slow'?	2000	3.1.2
useUsageAdvisor	Should the driver issue 'usage' warnings advising proper and efficient usage of JDBC and MySQL Connector/J to the log (true/false, defaults to 'false')?	false	3.1.1
autoGenerateTestcaseScript	Should the driver dump the SQL it is executing, including server-side prepared statements to STDERR?	false	3.1.9
dumpMetadataOnColumnNotFound	Should the driver dump the field-level metadata of a result set into the exception message when ResultSet.findColumn() fails?	false	3.1.13
dumpQueriesOnException	Should the driver dump the contents of the query sent to the server in the message for SQLExceptions?	false	3.1.3
enablePacketDebug	When enabled, a ring-buffer of 'packetDebugBufferSize' packets will be kept, and dumped when exceptions are thrown in key areas in the driver's code	false	3.1.3
explainSlowQueries	If 'logSlowQueries' is enabled, should the driver automatically issue an 'EXPLAIN' on the server and send the results to the configured log at a WARN level?	false	3.1.2
logSlowQueries	Should queries that take longer than 'slowQueryThresholdMillis' be logged?	false	3.1.2
traceProtocol	Should trace-level network protocol be logged?	false	3.1.2

**Miscellaneous.**

Property Name	Definition	Default Value	Since Version
useUnicode	Should the driver use Unicode character encodings when handling strings? Should only be used when the driver can't determine the character set mapping, or you are trying to 'force' the driver to use a character set that MySQL either doesn't natively support (such as UTF-8), true/false, defaults to 'true'	true	1.1g
characterEncoding	If 'useUnicode' is set to true, what character encoding should the driver use when dealing with strings? (defaults is to 'autodetect')		1.1g
characterSetResults	Character set to tell the server to return results as.		3.0.13
connectionCollation	If set, tells the server to use this collation via 'set'		3.0.13

	collation_connection'		
sessionVariables	A comma-separated list of name/value pairs to be sent as SET SESSION ... to the server when the driver connects.		3.1.8
allowNanAndInf	Should the driver allow NaN or +/- INF values in PreparedStatement.setDouble()?	false	3.1.5
autoClosePstmtStreams	Should the driver automatically call .close() on streams/readers passed as arguments via set*() methods?	false	3.1.12
autoDeserialize	Should the driver automatically detect and de-serialize objects stored in BLOB fields?	false	3.1.5
capitalizeTypeNames	Capitalize type names in DatabaseMetaData? (usually only useful when using WebObjects, true/false, defaults to 'false')	false	2.0.7
clobCharacterEncoding	The character encoding to use for sending and retrieving TEXT, MEDIUMTEXT and LONGTEXT values instead of the configured connection characterEncoding		5.0.0
clobberStreamingResults	This will cause a 'streaming' ResultSet to be automatically closed, and any outstanding data still streaming from the server to be discarded if another query is executed before all the data has been read from the server.	false	3.0.9
continueBatchOnError	Should the driver continue processing batch commands if one statement fails. The JDBC spec allows either way (defaults to 'true').	true	3.0.3
createDatabaseIfNotExist	Creates the database given in the URL if it doesn't yet exist. Assumes the configured user has permissions to create databases.	false	3.1.9
emptyStringsConvertToZero	Should the driver allow conversions from empty string fields to numeric values of '0'?	true	3.1.8
emulateLocators	N/A	false	3.1.0
emulateUnsupportedPstmts	Should the driver detect prepared statements that are not supported by the server, and replace them with client-side emulated versions?	true	3.1.7
ignoreNonTxTables	Ignore non-transactional table warning for roll-back? (defaults to 'false').	false	3.0.9
jdbcCompliantTruncation	Should the driver throw java.sql.DataTruncation exceptions when data is truncated as is required by the JDBC specification when connected to a server that supports warnings(MySQL 4.1.0 and newer)?	true	3.1.2
maxRows	The maximum number of rows to return (0, the default means return all rows).	-1	all versions
noAccessToProcedureBodies	When determining procedure parameter types for CallableStatements, and the connected user can't access procedure bodies through "SHOW CREATE PROCEDURE" or select on mysql.proc should the driver instead create basic metadata (all parameters reported as INOUT VARCHARs) instead of throwing an exception?	false	5.0.3
noDatetimeStringSync	Don't ensure that Result-	false	3.1.7

	<code>Set.getDatetimeType().toString().equals(ResultSet.getString())</code>		
noTimezoneConversionForTimeType	Don't convert TIME values using the server timezone if 'useTimezone'='true'	false	5.0.0
nullCatalogMeansCurrent	When DatabaseMetadataMethods ask for a 'catalog' parameter, does the value null mean use the current catalog? (this is not JDBC-compliant, but follows legacy behavior from earlier versions of the driver)	true	3.1.8
nullNamePatternMatchesAll	Should DatabaseMetaData methods that accept *pattern parameters treat null the same as '%' (this is not JDBC-compliant, however older versions of the driver accepted this departure from the specification)	true	3.1.8
overrideSupportsIntegrityEnhancementFacility	Should the driver return "true" for DatabaseMetaData.supportsIntegrityEnhancementFacility() even if the database doesn't support it to work-around applications that require this method to return "true" to signal support of foreign keys, even though the SQL specification states that this facility contains much more than just foreign key support (one such application being OpenOffice)?	false	3.1.12
pedantic	Follow the JDBC spec to the letter.	false	3.0.0
pinGlobalTxToPhysicalConnection	When using XAConnections, should the driver ensure that operations on a given XID are always routed to the same physical connection? This allows the XAConnection to support "XA START ... JOIN" after "XA END" has been called	false	5.0.1
processEscapeCodesForPrepStmts	Should the driver process escape codes in queries that are prepared?	true	3.1.12
relaxAutoCommit	If the version of MySQL the driver connects to does not support transactions, still allow calls to commit(), rollback() and setAutoCommit() (true/false, defaults to 'false')?	false	2.0.13
retainStatementAfterResultSetClose	Should the driver retain the Statement reference in a ResultSet after ResultSet.close() has been called. This is not JDBC-compliant after JDBC-4.0.	false	3.1.11
rollbackOnPooledClose	Should the driver issue a rollback() when the logical connection in a pool is closed?	true	3.0.15
runningCTS13	Enables workarounds for bugs in Sun's JDBC compliance testsuite version 1.3	false	3.1.7
serverTimezone	Override detection/mapping of timezone. Used when timezone from server doesn't map to Java timezone		3.0.2
strictFloatingPoint	Used only in older versions of compliance test	false	3.0.0
strictUpdates	Should the driver do strict checking (all primary keys selected) of updatable result sets (true, false, defaults to 'true')?	true	3.0.4
tinyInt1isBit	Should the driver treat the datatype TINYINT(1) as the BIT type (because the server silently converts BIT -> TINYINT(1) when creating tables)?	true	3.0.16

transformedBitIsBoolean	If the driver converts TINYINT(1) to a different type, should it use BOOLEAN instead of BIT for future compatibility with MySQL-5.0, as MySQL-5.0 has a BIT type?	false	3.1.9
ultraDevHack	Create PreparedStatements for prepareCall() when required, because UltraDev is broken and issues a prepareCall() for _all_ statements? (true/false, defaults to 'false')	false	2.0.3
useGmtMillisForDatetimes	Convert between session timezone and GMT before creating Date and Timestamp instances (value of "false" is legacy behavior, "true" leads to more JDBC-compliant behavior.	false	3.1.12
useHostsInPrivileges	Add '@hostname' to users in DatabaseMetaData.getColumn/TablePrivileges() (true/false), defaults to 'true'.	true	3.0.2
useInformationSchema	When connected to MySQL-5.0.7 or newer, should the driver use the INFORMATION_SCHEMA to derive information used by DatabaseMetaData?	false	5.0.0
useJDBCCompliant-TimezoneShift	Should the driver use JDBC-compliant rules when converting TIME/TIMESTAMP/DATETIME values' timezone information for those JDBC arguments which take a java.util.Calendar argument? (Notice that this option is exclusive of the "use-Timezone=true" configuration option.)	false	5.0.0
useOldUTF8Behavior	Use the UTF-8 behavior the driver did when communicating with 4.0 and older servers	false	3.1.6
useOnlyServerErrorMessages	Don't prepend 'standard' SQLState error messages to error messages returned by the server.	true	3.0.15
useServerPrepStmts	Use server-side prepared statements if the server supports them? (defaults to 'true').	true	3.1.0
useSqlStateCodes	Use SQL Standard state codes instead of 'legacy' X/Open/SQL state codes (true/false), default is 'true'	true	3.1.3
useStreamLengthsInPrepStmts	Honor stream length parameter in PreparedStatement/ResultSet.setXXXStream() method calls (true/false, defaults to 'true')?	true	3.0.2
useTimezone	Convert time/date types between client and server timezones (true/false, defaults to 'false')?	false	3.0.2
useUnbufferedInput	Don't use BufferedInputStream for reading data from the server	true	3.0.11
yearIsDateType	Should the JDBC driver treat the MySQL type "YEAR" as a java.sql.Date, or as a SHORT?	true	3.1.9
zeroDateTimeBehavior	What should happen when the driver encounters DATETIME values that are composed entirely of zeroes (used by MySQL to represent invalid dates)? Valid values are 'exception', 'round' and 'convertToNull'.	exception	3.1.4

Connector/J also supports access to MySQL via named pipes on Windows NT/2000/XP using the NamedPipeSocketFactory as a plugin-socket factory via the socketFactory property. If you don't use a namedPipePath property, the default of '\\.\pipe\MySQL' will be used. If you use the [Named-](#)

`PipeSocketFactory`, the hostname and port number values in the JDBC url will be ignored. You can enable this feature using:

```
socketFactory=com.mysql.jdbc.NamedPipeSocketFactory
```

Named pipes only work when connecting to a MySQL server on the same physical machine as the one the JDBC driver is being used on. In simple performance tests, it appears that named pipe access is between 30%-50% faster than the standard TCP/IP access.

You can create your own socket factories by following the example code in `com.mysql.jdbc.NamedPipeSocketFactory`, or `com.mysql.jdbc.StandardSocketFactory`.

### 18.3.4.2. JDBC API Implementation Notes

MySQL Connector/J passes all of the tests in the publicly-available version of Sun's JDBC compliance test suite. However, in many places the JDBC specification is vague about how certain functionality should be implemented, or the specification allows leeway in implementation.

This section gives details on a interface-by-interface level about how certain implementation decisions may affect how you use MySQL Connector/J.

- Blob

The Blob implementation does not allow in-place modification (they are copies, as reported by the `DatabaseMetaData.locatorsUpdateCopies()` method). Because of this, you should use the corresponding `PreparedStatement.setBlob()` or `ResultSet.updateBlob()` (in the case of updatable result sets) methods to save changes back to the database.

Starting with Connector/J version 3.1.0, you can emulate Blobs with locators by adding the property 'emulateLocators=true' to your JDBC URL. You must then use a column alias with the value of the column set to the actual name of the Blob column in the `SELECT` that you write to retrieve the Blob. The `SELECT` must also reference only one table, the table must have a primary key, and the `SELECT` must cover all columns that make up the primary key. The driver will then delay loading the actual Blob data until you retrieve the Blob and call retrieval methods (`getInputStream()`, `getBytes()`, and so forth) on it.

- CallableStatement

Starting with Connector/J 3.1.1, stored procedures are supported when connecting to MySQL version 5.0 or newer via the `CallableStatement` interface. Currently, the `getParameterMetaData()` method of `CallableStatement` is not supported.

- Clob

The Clob implementation does not allow in-place modification (they are copies, as reported by the `DatabaseMetaData.locatorsUpdateCopies()` method). Because of this, you should use the `PreparedStatement.setClob()` method to save changes back to the database. The JDBC API does not have a `ResultSet.updateClob()` method.

- Connection

Unlike older versions of MM.MySQL the `isClosed()` method does not ping the server to determine if it is alive. In accordance with the JDBC specification, it only returns true if `closed()` has been called on the connection. If you need to determine if the connection is still valid, you should issue a simple query, such as `SELECT 1`. The driver will throw an exception if the connection is no longer valid.



- DatabaseMetaData

Foreign Key information (`getImportedKeys()`/`getExportedKeys()` and `getCrossReference()`) is only available from InnoDB tables. However, the driver uses `SHOW CREATE TABLE` to retrieve this information, so when other storage engines support foreign keys, the driver will transparently support them as well.

- PreparedStatement

PreparedStatements are implemented by the driver, as MySQL does not have a prepared statement feature. Because of this, the driver does not implement `getParameterMetaData()` or `getMetaData()` as it would require the driver to have a complete SQL parser in the client.

Starting with version 3.1.0 MySQL Connector/J, server-side prepared statements and binary-encoded result sets are used when the server supports them.

Take care when using a server-side prepared statement with **large** parameters that are set via `setBinaryStream()`, `setAsciiStream()`, `setUnicodeStream()`, `setBlob()`, or `setClob()`. If you want to re-execute the statement with any large parameter changed to a non-large parameter, it is necessary to call `clearParameters()` and set all parameters again. The reason for this is as follows:

- The driver streams the large data out-of-band to the prepared statement on the server side when the parameter is set (before execution of the prepared statement).
- Once that has been done, the stream used to read the data on the client side is closed (as per the JDBC spec), and can't be read from again.
- If a parameter changes from large to non-large, the driver must reset the server-side state of the prepared statement to allow the parameter that is being changed to take the place of the prior large value. This removes all of the large data that has already been sent to the server, thus requiring the data to be re-sent, via the `setBinaryStream()`, `setAsciiStream()`, `setUnicodeStream()`, `setBlob()` or `setClob()` methods.

Consequently, if you want to change the type of a parameter to a non-large one, you must call `clearParameters()` and set all parameters of the prepared statement again before it can be re-executed.

- ResultSet

By default, ResultSets are completely retrieved and stored in memory. In most cases this is the most efficient way to operate, and due to the design of the MySQL network protocol is easier to implement. If you are working with ResultSets that have a large number of rows or large values, and can not allocate heap space in your JVM for the memory required, you can tell the driver to stream the results back one row at a time.

To enable this functionality, you need to create a Statement instance in the following manner:

```
stmt = conn.createStatement( java.sql.ResultSet.TYPE_FORWARD_ONLY,
                             java.sql.ResultSet.CONCUR_READ_ONLY );
stmt.setFetchSize( Integer.MIN_VALUE );
```

The combination of a forward-only, read-only result set, with a fetch size of `Integer.MIN_VALUE` serves as a signal to the driver to stream result sets row-by-row. After this any result sets created with the statement will be retrieved row-by-row.

There are some caveats with this approach. You will have to read all of the rows in the result set (or close it) before you can issue any other queries on the connection, or an exception will be thrown.

The earliest the locks these statements hold can be released (whether they be [MyISAM](#) table-level locks or row-level locks in some other storage engine such as [InnoDB](#)) is when the statement completes.

If the statement is within scope of a transaction, then locks are released when the transaction completes (which implies that the statement needs to complete first). As with most other databases, statements are not complete until all the results pending on the statement are read or the active result set for the statement is closed.

Therefore, if using streaming results, you should process them as quickly as possible if you want to maintain concurrent access to the tables referenced by the statement producing the result set.

- `ResultSetMetaData`

The `isAutoIncrement()` method only works when using MySQL servers 4.0 and newer.

- `Statement`

When using versions of the JDBC driver earlier than 3.2.1, and connected to server versions earlier than 5.0.3, the `setFetchSize()` method has no effect, other than to toggle result set streaming as described above.

MySQL does not support SQL cursors, and the JDBC driver doesn't emulate them, so `setCursorName()` has no effect.

### 18.3.4.3. Java, JDBC and MySQL Types

MySQL Connector/J is flexible in the way it handles conversions between MySQL data types and Java data types.

In general, any MySQL data type can be converted to a `java.lang.String`, and any numerical type can be converted to any of the Java numerical types, although round-off, overflow, or loss of precision may occur.

Starting with Connector/J 3.1.0, the JDBC driver will issue warnings or throw `DataTruncation` exceptions as is required by the JDBC specification unless the connection was configured not to do so by using the property `jdbcCompliantTruncation` and setting it to `false`.

The conversions that are always guaranteed to work are listed in the following table:

#### Connection Properties - Miscellaneous.

These MySQL Data Types	Can always be converted to these Java types
<code>CHAR</code> , <code>VARCHAR</code> , <code>BLOB</code> , <code>TEXT</code> , <code>ENUM</code> , and <code>SET</code>	<code>java.lang.String</code> , <code>java.io.InputStream</code> , <code>java.io.Reader</code> , <code>java.sql.Blob</code> , <code>java.sql.Clob</code>
<code>FLOAT</code> , <code>REAL</code> , <code>DOUBLE PRECISION</code> , <code>NUMERIC</code> , <code>DECIMAL</code> , <code>TINYINT</code> , <code>SMALLINT</code> , <code>MEDIUMINT</code> , <code>INTEGER</code> , <code>BIGINT</code>	<code>java.lang.String</code> , <code>java.lang.Short</code> , <code>java.lang.Integer</code> , <code>java.lang.Long</code> , <code>java.lang.Double</code> , <code>java.math.BigDecimal</code>
<code>DATE</code> , <code>TIME</code> , <code>DATETIME</code> , <code>TIMESTAMP</code>	<code>java.lang.String</code> , <code>java.sql.Date</code> , <code>java.sql.Timestamp</code>

**Note:** round-off, overflow or loss of precision may occur if you choose a Java numeric data type that has less precision or capacity than the MySQL data type you are converting to/from.

The `ResultSet.getObject()` method uses the type conversions between MySQL and Java types, following the JDBC specification where appropriate. The value returned by `ResultSet-  
MetaData.getColumnClassName()` is also shown below. For more information on the `java.sql.Types` classes see [Java 2 Platform Types](http://java.sun.com/j2se/1.4.2/docs/api/java/sql/Types.html) [http://java.sun.com/j2se/1.4.2/docs/api/java/sql/Types.html].

#### MySQL Types to Java Types for `ResultSet.getObject()`.

MySQL Type Name	Return value of <code>getColumnClassName</code>
TINYINT	TINYINT
BOOL, BOOLEAN	TINYINT
SMALLINT[(M)] [UNSIGNED]	SMALLINT [UNSIGNED]
MEDIUMINT[(M)] [UNSIGNED]	MEDIUMINT [UNSIGNED]
INT,INTEGER[(M)] [UNSIGNED]	INTEGER [UNSIGNED]
BIGINT[(M)] [UNSIGNED]	BIGINT [UNSIGNED]
FLOAT[(M,D)]	FLOAT
DOUBLE[(M,B)]	DOUBLE
DECIMAL[(M[,D])]	DECIMAL
DATE	DATE
DATETIME	DATETIME
TIMESTAMP[(M)]	TIMESTAMP
TIME	TIME
YEAR[(2 4)]	YEAR
CHAR(M)	CHAR
VARCHAR(M) [BINARY]	VARCHAR
BINARY(M)	BINARY
VARBINARY(M)	VARBINARY
TINYBLOB	TINYBLOB
TINYTEXT	VARCHAR
BLOB	BLOB
TEXT	VARCHAR
MEDIUMBLOB	MEDIUMBLOB
MEDIUMTEXT	VARCHAR
LOBLOB	LOBLOB
LONGTEXT	VARCHAR
ENUM('value1','value2',...)	CHAR
SET('value1','value2',...)	CHAR

#### 18.3.4.4. Using Character Sets and Unicode

All strings sent from the JDBC driver to the server are converted automatically from native Java Unicode form to the client character encoding, including all queries sent via `Statement.execute()`,

`Statement.executeUpdate()`, `Statement.executeQuery()` as well as all `PreparedStatement` and `CallableStatement` parameters with the exclusion of parameters set using `setBytes()`, `setBinaryStream()`, `setAsciiStream()`, `setUnicodeStream()` and `setBlob()`.

Prior to MySQL Server 4.1, Connector/J supported a single character encoding per connection, which could either be automatically detected from the server configuration, or could be configured by the user through the `useUnicode` and "`characterEncoding`" properties.

Starting with MySQL Server 4.1, Connector/J supports a single character encoding between client and server, and any number of character encodings for data returned by the server to the client in `ResultSet`.

The character encoding between client and server is automatically detected upon connection. The encoding used by the driver is specified on the server via the `character_set` system variable for server versions older than 4.1.0 and `character_set_server` for server versions 4.1.0 and newer. For more information, see [Section 10.3.1, "Server Character Set and Collation"](#).

To override the automatically-detected encoding on the client side, use the `characterEncoding` property in the URL used to connect to the server.

When specifying character encodings on the client side, Java-style names should be used. The following table lists Java-style names for MySQL character sets:

#### MySQL to Java Encoding Name Translations.

MySQL Character Set Name	Java-Style Character Encoding Name
ascii	US-ASCII
big5	Big5
gbk	GBK
sjis	SJIS (or Cp932 or MS932 for MySQL Server < 4.1.11)
cp932	Cp932 or MS932 (MySQL Server > 4.1.11)
gb2312	EUC_CN
ujis	EUC_JP
euckr	EUC_KR
latin1	ISO8859_1
latin2	ISO8859_2
greek	ISO8859_7
hebrew	ISO8859_8
cp866	Cp866
tis620	TIS620
cp1250	Cp1250
cp1251	Cp1251
cp1257	Cp1257
macroman	MacRoman
macce	MacCentralEurope
utf8	UTF-8
ucs2	UnicodeBig

**Warning.** Do not issue the query 'set names' with Connector/J, as the driver will not detect that the character set has changed, and will continue to use the character set detected during the initial connection setup.

To allow multiple character sets to be sent from the client, the UTF-8 encoding should be used, either by configuring `utf8` as the default server character set, or by configuring the JDBC driver to use UTF-8 through the `characterEncoding` property.

### 18.3.4.5. Connecting Securely Using SSL

SSL in MySQL Connector/J encrypts all data (other than the initial handshake) between the JDBC driver and the server. The performance penalty for enabling SSL is an increase in query processing time between 35% and 50%, depending on the size of the query, and the amount of data it returns.

For SSL Support to work, you must have the following:

- A JDK that includes JSSE (Java Secure Sockets Extension), like JDK-1.4.1 or newer. SSL does not currently work with a JDK that you can add JSSE to, like JDK-1.2.x or JDK-1.3.x due to the following JSSE bug: <http://developer.java.sun.com/developer/bugParade/bugs/4273544.html>
- A MySQL server that supports SSL and has been compiled and configured to do so, which is MySQL-4.0.4 or later, see [Section 5.8.7, “Using Secure Connections”](#), for more information.
- A client certificate (covered later in this section)

You will first need to import the MySQL server CA Certificate into a Java truststore. A sample MySQL server CA Certificate is located in the [SSL](#) subdirectory of the MySQL source distribution. This is what SSL will use to determine if you are communicating with a secure MySQL server.

To use Java's `keytool` to create a truststore in the current directory, and import the server's CA certificate (`cacert.pem`), you can do the following (assuming that `keytool` is in your path. The `keytool` should be located in the `bin` subdirectory of your JDK or JRE):

```
shell> keytool -import -alias mysqlServerCACert -file cacert.pem -keystore truststore
```

Keytool will respond with the following information:

```
Enter keystore password: *****
Owner: EMAILADDRESS=walrus@example.com, CN=Walrus, O=MySQL AB, L=Orenburg, ST=Some-
State, C=RU
Issuer: EMAILADDRESS=walrus@example.com, CN=Walrus, O=MySQL AB, L=Orenburg, ST=Som
e-State, C=RU
Serial number: 0
Valid from: Fri Aug 02 16:55:53 CDT 2002 until: Sat Aug 02 16:55:53 CDT 2003
Certificate fingerprints:
    MD5:  61:91:A0:F2:03:07:61:7A:81:38:66:DA:19:C4:8D:AB
    SHA1: 25:77:41:05:D5:AD:99:8C:14:8C:CA:68:9C:2F:B8:89:C3:34:4D:6C
Trust this certificate? [no]: yes
Certificate was added to keystore
```

You will then need to generate a client certificate, so that the MySQL server knows that it is talking to a secure client:

```
shell> keytool -genkey -keyalg rsa -alias mysqlClientCertificate -keystore keystore
```

Keytool will prompt you for the following information, and create a keystore named `keystore` in the current directory.

You should respond with information that is appropriate for your situation:

```

Enter keystore password: *****
What is your first and last name?
  [Unknown]: Matthews
What is the name of your organizational unit?
  [Unknown]: Software Development
What is the name of your organization?
  [Unknown]: MySQL AB
What is the name of your City or Locality?
  [Unknown]: Flossmoor
What is the name of your State or Province?
  [Unknown]: IL
What is the two-letter country code for this unit?
  [Unknown]: US
Is <CN=Matthews, OU=Software Development, O=MySQL AB,
L=Flossmoor, ST=IL, C=US> correct?
  [no]: y

Enter key password for <mysqlClientCertificate>
      (RETURN if same as keystore password):

```

Finally, to get JSSE to use the keystore and truststore that you have generated, you need to set the following system properties when you start your JVM, replacing `path_to_keystore_file` with the full path to the keystore file you created, `path_to_truststore_file` with the path to the truststore file you created, and using the appropriate password values for each property.

```

-Djavax.net.ssl.keyStore=path_to_keystore_file
-Djavax.net.ssl.keyStorePassword=*****
-Djavax.net.ssl.trustStore=path_to_truststore_file
-Djavax.net.ssl.trustStorePassword=*****

```

You will also need to set `useSSL` to `true` in your connection parameters for MySQL Connector/J, either by adding `useSSL=true` to your URL, or by setting the property `useSSL` to `true` in the `java.util.Properties` instance you pass to `DriverManager.getConnection()`.

You can test that SSL is working by turning on JSSE debugging (as detailed below), and look for the following key events:

```

...
*** ClientHello, v3.1
RandomCookie: GMT: 1018531834 bytes = { 199, 148, 180, 215, 74, 12, 54, 244, 0, 168, 55, 103, 215, 64 }
Session ID: {}
Cipher Suites: { 0, 5, 0, 4, 0, 9, 0, 10, 0, 18, 0, 19, 0, 3, 0, 17 }
Compression Methods: { 0 }
***
[write] MD5 and SHA1 hashes: len = 59
0000: 01 00 00 37 03 01 3D B6 90 FA C7 94 B4 D7 4A 0C ...7..=.....J.
0010: 36 F4 00 A8 37 67 D7 40 10 8A E1 BE 84 99 02 D9 6...7g.@.....
0020: DB EF CA 13 79 4E 00 00 10 00 05 00 04 00 09 00 ...yN.....
0030: 0A 00 12 00 13 00 03 00 11 01 00 .....
main, WRITE: SSL v3.1 Handshake, length = 59
main, READ: SSL v3.1 Handshake, length = 74
*** ServerHello, v3.1
RandomCookie: GMT: 1018577560 bytes = { 116, 50, 4, 103, 25, 100, 58, 202, 79, 185, 178, 100, 215, 64 }
Session ID: {163, 227, 84, 53, 81, 127, 252, 254, 178, 179, 68, 63, 182, 158, 30, 11, 150, 79, 170, 7 }
Cipher Suite: { 0, 5 }
Compression Method: 0
***
%% Created: [Session-1, SSL_RSA_WITH_RC4_128_SHA]
** SSL_RSA_WITH_RC4_128_SHA
[read] MD5 and SHA1 hashes: len = 74
0000: 02 00 00 46 03 01 3D B6 43 98 74 32 04 67 19 64 ...F..=.C.t2.g.d
0010: 3A CA 4F B9 B2 64 D7 42 FE 15 53 BB BE 2A AA 03 :.O..d.B..S..*..
0020: 84 6E 52 94 A0 5C 20 A3 E3 54 35 51 7F FC FE B2 .nR..\ ..T5Q....
0030: B3 44 3F B6 9E 1E 0B 96 4F AA 4C FF 5C 0F E2 18 .D?.....O.L.\...
0040: 11 B1 DB 9E B1 BB 8F 00 05 00 .....
main, READ: SSL v3.1 Handshake, length = 1712
...

```

JSSE provides debugging (to STDOUT) when you set the following system property: -

`Djavax.net.debug=all` This will tell you what keystores and truststores are being used, as well as what is going on during the SSL handshake and certificate exchange. It will be helpful when trying to determine what is not working when trying to get an SSL connection to happen.

### 18.3.4.6. Using Master/Slave Replication with ReplicationConnection

Starting with Connector/J 3.1.7, we've made available a variant of the driver that will automatically send queries to a read/write master, or a failover or round-robin loadbalanced set of slaves based on the state of `Connection.getReadOnly()`.

An application signals that it wants a transaction to be read-only by calling `Connection.setReadOnly(true)`, this replication-aware connection will use one of the slave connections, which are load-balanced per-vm using a round-robin scheme (a given connection is sticky to a slave unless that slave is removed from service). If you have a write transaction, or if you have a read that is time-sensitive (remember, replication in MySQL is asynchronous), set the connection to be not read-only, by calling `Connection.setReadOnly(false)` and the driver will ensure that further calls are sent to the master MySQL server. The driver takes care of propagating the current state of auto-commit, isolation level, and catalog between all of the connections that it uses to accomplish this load balancing functionality.

To enable this functionality, use the "`com.mysql.jdbc.ReplicationDriver`" class when configuring your application server's connection pool or when creating an instance of a JDBC driver for your standalone application. Because it accepts the same URL format as the standard MySQL JDBC driver, `ReplicationDriver` does not currently work with `java.sql.DriverManager`-based connection creation unless it is the only MySQL JDBC driver registered with the `DriverManager`.

Here is a short, simple example of how `ReplicationDriver` might be used in a standalone application.

```
import java.sql.Connection;
import java.sql.ResultSet;
import java.util.Properties;

import com.mysql.jdbc.ReplicationDriver;

public class ReplicationDriverDemo {

    public static void main(String[] args) throws Exception {
        ReplicationDriver driver = new ReplicationDriver();

        Properties props = new Properties();

        // We want this for failover on the slaves
        props.put("autoReconnect", "true");

        // We want to load balance between the slaves
        props.put("roundRobinLoadBalance", "true");

        props.put("user", "foo");
        props.put("password", "bar");

        //
        // Looks like a normal MySQL JDBC url, with a comma-separated list
        // of hosts, the first being the 'master', the rest being any number
        // of slaves that the driver will load balance against
        //
        Connection conn =
            driver.connect("jdbc:mysql://master,slave1,slave2,slave3/test",
                props);

        //
        // Perform read/write work on the master
        // by setting the read-only flag to "false"
        //
        conn.setReadOnly(false);
        conn.setAutoCommit(false);
        conn.createStatement().executeUpdate("UPDATE some_table ...");
        conn.commit();

        //
    }
}
```

```

// Now, do a query from a slave, the driver automatically picks one
// from the list
//
conn.setReadOnly(true);

ResultSet rs = conn.createStatement().executeQuery("SELECT a,b,c FROM some_other_table");
    .....
}
}

```

## 18.3.5. Connector/J Notes and Tips

### 18.3.5.1. Basic JDBC Concepts

This section provides some general JDBC background.

#### 18.3.5.1.1. Connecting to MySQL Using the [DriverManager](#) Interface

When you are using JDBC outside of an application server, the [DriverManager](#) class manages the establishment of Connections.

The [DriverManager](#) needs to be told which JDBC drivers it should try to make Connections with. The easiest way to do this is to use [Class.forName\(\)](#) on the class that implements the [java.sql.Driver](#) interface. With MySQL Connector/J, the name of this class is [com.mysql.jdbc.Driver](#). With this method, you could use an external configuration file to supply the driver class name and driver parameters to use when connecting to a database.

The following section of Java code shows how you might register MySQL Connector/J from the [main\(\)](#) method of your application:

```

import java.sql.Connection;
import java.sql.DriverManager;
import java.sql.SQLException;

// Notice, do not import com.mysql.jdbc.*
// or you will have problems!

public class LoadDriver {
    public static void main(String[] args) {
        try {
            // The newInstance() call is a work around for some
            // broken Java implementations

            Class.forName("com.mysql.jdbc.Driver").newInstance();
        } catch (Exception ex) {
            // handle the error
        }
    }
}

```

After the driver has been registered with the [DriverManager](#), you can obtain a [Connection](#) instance that is connected to a particular database by calling [DriverManager.getConnection\(\)](#):

#### Example 18.1. Obtaining a connection from the [DriverManager](#)

This example shows how you can obtain a [Connection](#) instance from the [DriverManager](#). There are a few different signatures for the [getConnection\(\)](#) method. You should see the API documentation that comes with your JDK for more specific information on how to use them.

```

import java.sql.Connection;
import java.sql.DriverManager;
import java.sql.SQLException;

```



```

... try {
    Connection conn = DriverManager.getConnection("jdbc:mysql://localhost/test?user=monty&passw

    // Do something with the Connection

    ....
} catch (SQLException ex) {
    // handle any errors
    System.out.println("SQLException: " + ex.getMessage());
    System.out.println("SQLState: " + ex.getSQLState());
    System.out.println("VendorError: " + ex.getErrorCode());
}

```

Once a [Connection](#) is established, it can be used to create [Statement](#) and [PreparedStatement](#) objects, as well as retrieve metadata about the database. This is explained in the following sections.

### 18.3.5.1.2. Using Statements to Execute SQL

[Statement](#) objects allow you to execute basic SQL queries and retrieve the results through the [ResultSet](#) class which is described later.

To create a [Statement](#) instance, you call the `createStatement()` method on the [Connection](#) object you have retrieved via one of the `DriverManager.getConnection()` or `DataSource.getConnection()` methods described earlier.

Once you have a [Statement](#) instance, you can execute a [SELECT](#) query by calling the `executeQuery(String)` method with the SQL you want to use.

To update data in the database, use the `executeUpdate(String SQL)` method. This method returns the number of rows affected by the update statement.

If you don't know ahead of time whether the SQL statement will be a [SELECT](#) or an [UPDATE/INSERT](#), then you can use the `execute(String SQL)` method. This method will return true if the SQL query was a [SELECT](#), or false if it was an [UPDATE](#), [INSERT](#), or [DELETE](#) statement. If the statement was a [SELECT](#) query, you can retrieve the results by calling the `getResultSet()` method. If the statement was an [UPDATE](#), [INSERT](#), or [DELETE](#) statement, you can retrieve the affected rows count by calling `getUpdateCount()` on the [Statement](#) instance.

#### Example 18.2. Using `java.sql.Statement` to execute a [SELECT](#) query

```

// assume that conn is an already created JDBC connection
Statement stmt = null;
ResultSet rs = null;

try {
    stmt = conn.createStatement();
    rs = stmt.executeQuery("SELECT foo FROM bar");

    // or alternatively, if you don't know ahead of time that
    // the query will be a SELECT...

    if (stmt.execute("SELECT foo FROM bar")) {
        rs = stmt.getResultSet();
    }

    // Now do something with the ResultSet ....
} finally {
    // it is a good idea to release
    // resources in a finally{} block
    // in reverse-order of their creation
    // if they are no-longer needed

    if (rs != null) {
        try {
            rs.close();

```

```

    } catch (SQLException sqlEx) { // ignore }
    rs = null;
}
if (stmt != null) {
    try {
        stmt.close();
    } catch (SQLException sqlEx) { // ignore }
    stmt = null;
}
}

```

### 18.3.5.1.3. Using `CallableStatements` to Execute Stored Procedures

Starting with MySQL server version 5.0 when used with Connector/J 3.1.1 or newer, the `java.sql.CallableStatement` interface is fully implemented with the exception of the `getParameterMetaData()` method.

Connector/J exposes stored procedure functionality through JDBC's `CallableStatement` interface.

**Note.** Current versions of MySQL server do not return enough information for the JDBC driver to provide result set metadata for callable statements. This means that when using `CallableStatement`, `ResultSetMetaData` may return `NULL`.

The following example shows a stored procedure that returns the value of `inOutParam` incremented by 1, and the string passed in via `inputParam` as a `ResultSet`:

#### Example 18.3. Stored Procedures

```

CREATE PROCEDURE demoSp(IN inputParam VARCHAR(255), INOUT inOutParam INT)
BEGIN
    DECLARE z INT;
    SET z = inOutParam + 1;
    SET inOutParam = z;

    SELECT inputParam;

    SELECT CONCAT('zyxw', inputParam);
END

```

To use the `demoSp` procedure with Connector/J, follow these steps:

1. Prepare the callable statement by using `Connection.prepareCall()`.

Notice that you have to use JDBC escape syntax, and that the parentheses surrounding the parameter placeholders are not optional:

#### Example 18.4. Using `Connection.prepareCall()`

```

import java.sql.CallableStatement;
...

//
// Prepare a call to the stored procedure 'demoSp'
// with two parameters
//
// Notice the use of JDBC-escape syntax ({call ...})
//

```

```
CallableStatement cStmt = conn.prepareCall("{call demoSp(?, ?)}");

cStmt.setString(1, "abcdefg");
```

**Note.** `Connection.prepareCall()` is an expensive method, due to the metadata retrieval that the driver performs to support output parameters. For performance reasons, you should try to minimize unnecessary calls to `Connection.prepareCall()` by reusing `CallableStatement` instances in your code.

## 2. Register the output parameters (if any exist)

To retrieve the values of output parameters (parameters specified as `OUT` or `INOUT` when you created the stored procedure), JDBC requires that they be specified before statement execution using the various `registerOutputParameter()` methods in the `CallableStatement` interface:

### Example 18.5. Registering output parameters

```
import java.sql.Types;
...
//
// Connector/J supports both named and indexed
// output parameters. You can register output
// parameters using either method, as well
// as retrieve output parameters using either
// method, regardless of what method was
// used to register them.
//
// The following examples show how to use
// the various methods of registering
// output parameters (you should of course
// use only one registration per parameter).
//
//
// Registers the second parameter as output, and
// uses the type 'INTEGER' for values returned from
// getObject()
//
cStmt.registerOutParameter(2, Types.INTEGER);

//
// Registers the named parameter 'inOutParam', and
// uses the type 'INTEGER' for values returned from
// getObject()
//
cStmt.registerOutParameter("inOutParam", Types.INTEGER);
...
```

## 3. Set the input parameters (if any exist)

Input and in/out parameters are set as for `PreparedStatement` objects. However, `CallableStatement` also supports setting parameters by name:

### Example 18.6. Setting `CallableStatement` input parameters

```
...
//
// Set a parameter by index
```

```

//
cStmt.setString(1, "abcdefg");

//
// Alternatively, set a parameter using
// the parameter name
//
cStmt.setString("inputParameter", "abcdefg");

//
// Set the 'in/out' parameter using an index
//
cStmt.setInt(2, 1);

//
// Alternatively, set the 'in/out' parameter
// by name
//
cStmt.setInt("inOutParam", 1);

...

```

4. Execute the `CallableStatement`, and retrieve any result sets or output parameters.

Although `CallableStatement` supports calling any of the `Statement` execute methods (`executeUpdate()`, `executeQuery()` or `execute()`), the most flexible method to call is `execute()`, as you do not need to know ahead of time if the stored procedure returns result sets:

#### Example 18.7. Retrieving results and output parameter values

```

...
boolean hadResults = cStmt.execute();

//
// Process all returned result sets
//
while (hadResults) {
    ResultSet rs = cStmt.getResultSet();

    // process result set
    ...

    hadResults = rs.getMoreResults();
}

//
// Retrieve output parameters
//
// Connector/J supports both index-based and
// name-based retrieval
//
int outputValue = cStmt.getInt(2); // index-based
outputValue = cStmt.getInt("inOutParam"); // name-based

...

```

#### 18.3.5.1.4. Retrieving `AUTO_INCREMENT` Column Values

Before version 3.0 of the JDBC API, there was no standard way of retrieving key values from databases

that supported auto increment or identity columns. With older JDBC drivers for MySQL, you could always use a MySQL-specific method on the `Statement` interface, or issue the query `SELECT LAST_INSERT_ID()` after issuing an `INSERT` to a table that had an `AUTO_INCREMENT` key. Using the MySQL-specific method call isn't portable, and issuing a `SELECT` to get the `AUTO_INCREMENT` key's value requires another round-trip to the database, which isn't as efficient as possible. The following code snippets demonstrate the three different ways to retrieve `AUTO_INCREMENT` values. First, we demonstrate the use of the new JDBC-3.0 method `getGeneratedKeys()` which is now the preferred method to use if you need to retrieve `AUTO_INCREMENT` keys and have access to JDBC-3.0. The second example shows how you can retrieve the same value using a standard `SELECT LAST_INSERT_ID()` query. The final example shows how updatable result sets can retrieve the `AUTO_INCREMENT` value when using the `insertRow()` method.

### Example 18.8. Retrieving `AUTO_INCREMENT` column values using `Statement.getGeneratedKeys()`

```
Statement stmt = null;
ResultSet rs = null;

try {

    //
    // Create a Statement instance that we can use for
    // 'normal' result sets assuming you have a
    // Connection 'conn' to a MySQL database already
    // available

    stmt = conn.createStatement(java.sql.ResultSet.TYPE_FORWARD_ONLY,
                                java.sql.ResultSet.CONCUR_UPDATABLE);

    //
    // Issue the DDL queries for the table for this example
    //

    stmt.executeUpdate("DROP TABLE IF EXISTS autoIncTutorial");
    stmt.executeUpdate(
        "CREATE TABLE autoIncTutorial ("
        + "priKey INT NOT NULL AUTO_INCREMENT, "
        + "dataField VARCHAR(64), PRIMARY KEY (priKey))");

    //
    // Insert one row that will generate an AUTO INCREMENT
    // key in the 'priKey' field
    //

    stmt.executeUpdate(
        "INSERT INTO autoIncTutorial (dataField) "
        + "values ('Can I Get the Auto Increment Field?')",
        Statement.RETURN_GENERATED_KEYS);

    //
    // Example of using Statement.getGeneratedKeys()
    // to retrieve the value of an auto-increment
    // value
    //

    int autoIncKeyFromApi = -1;

    rs = stmt.getGeneratedKeys();

    if (rs.next()) {
        autoIncKeyFromApi = rs.getInt(1);
    } else {
        // throw an exception from here
    }

    rs.close();

    rs = null;

    System.out.println("Key returned from getGeneratedKeys():"
        + autoIncKeyFromApi);
} finally {
```

```

if (rs != null) {
    try {
        rs.close();
    } catch (SQLException ex) {
        // ignore
    }
}

if (stmt != null) {
    try {
        stmt.close();
    } catch (SQLException ex) {
        // ignore
    }
}
}

```

### Example 18.9. Retrieving `AUTO_INCREMENT` column values using `SELECT LAST_INSERT_ID()`

```

Statement stmt = null;
ResultSet rs = null;

try {
    //
    // Create a Statement instance that we can use for
    // 'normal' result sets.

    stmt = conn.createStatement();

    //
    // Issue the DDL queries for the table for this example
    //

    stmt.executeUpdate("DROP TABLE IF EXISTS autoIncTutorial");
    stmt.executeUpdate(
        "CREATE TABLE autoIncTutorial ("
        + "priKey INT NOT NULL AUTO_INCREMENT, "
        + "dataField VARCHAR(64), PRIMARY KEY (priKey))");

    //
    // Insert one row that will generate an AUTO INCREMENT
    // key in the 'priKey' field
    //

    stmt.executeUpdate(
        "INSERT INTO autoIncTutorial (dataField) "
        + "values ('Can I Get the Auto Increment Field?')");

    //
    // Use the MySQL LAST_INSERT_ID()
    // function to do the same thing as getGeneratedKeys()
    //

    int autoIncKeyFromFunc = -1;
    rs = stmt.executeQuery("SELECT LAST_INSERT_ID()");

    if (rs.next()) {
        autoIncKeyFromFunc = rs.getInt(1);
    } else {
        // throw an exception from here
    }

    rs.close();

    System.out.println("Key returned from " + "SELECT LAST_INSERT_ID()': "
        + autoIncKeyFromFunc);
} finally {
    if (rs != null) {
        try {
            rs.close();

```

```

        } catch (SQLException ex) {
            // ignore
        }
    }

    if (stmt != null) {
        try {
            stmt.close();
        } catch (SQLException ex) {
            // ignore
        }
    }
}

```

### Example 18.10. Retrieving `AUTO_INCREMENT` column values in `Updatable ResultSets`

```

Statement stmt = null;
ResultSet rs = null;

try {

    //
    // Create a Statement instance that we can use for
    // 'normal' result sets as well as an 'updatable'
    // one, assuming you have a Connection 'conn' to
    // a MySQL database already available
    //

    stmt = conn.createStatement(java.sql.ResultSet.TYPE_FORWARD_ONLY,
                               java.sql.ResultSet.CONCUR_UPDATABLE);

    //
    // Issue the DDL queries for the table for this example
    //

    stmt.executeUpdate("DROP TABLE IF EXISTS autoIncTutorial");
    stmt.executeUpdate(
        "CREATE TABLE autoIncTutorial ( "
        + "priKey INT NOT NULL AUTO_INCREMENT, "
        + "dataField VARCHAR(64), PRIMARY KEY (priKey))");

    //
    // Example of retrieving an AUTO INCREMENT key
    // from an updatable result set
    //

    rs = stmt.executeQuery("SELECT priKey, dataField "
        + "FROM autoIncTutorial");

    rs.moveToInsertRow();

    rs.updateString("dataField", "AUTO INCREMENT here?");
    rs.insertRow();

    //
    // the driver adds rows at the end
    //

    rs.last();

    //
    // We should now be on the row we just inserted
    //

    int autoIncKeyFromRS = rs.getInt("priKey");

    rs.close();

    rs = null;

    System.out.println("Key returned for inserted row: "
        + autoIncKeyFromRS);

} finally {

```

```

if (rs != null) {
    try {
        rs.close();
    } catch (SQLException ex) {
        // ignore
    }
}

if (stmt != null) {
    try {
        stmt.close();
    } catch (SQLException ex) {
        // ignore
    }
}
}

```

When you run the preceding example code, you should get the following output: Key returned from `getGeneratedKeys()`: 1 Key returned from `SELECT LAST_INSERT_ID()`: 1 Key returned for inserted row: 2 You should be aware, that at times, it can be tricky to use the `SELECT LAST_INSERT_ID()` query, as that function's value is scoped to a connection. So, if some other query happens on the same connection, the value will be overwritten. On the other hand, the `getGeneratedKeys()` method is scoped by the `Statement` instance, so it can be used even if other queries happen on the same connection, but not on the same `Statement` instance.

## 18.3.5.2. Using Connector/J with J2EE and Other Java Frameworks

This section describes how to use Connector/J in several contexts.

### 18.3.5.2.1. General J2EE Concepts

This section provides general background on J2EE concepts that pertain to use of Connector/J.

#### 18.3.5.2.1.1. Understanding Connection Pooling

Connection pooling is a technique of creating and managing a pool of connections that are ready for use by any thread that needs them.

This technique of pooling connections is based on the fact that most applications only need a thread to have access to a JDBC connection when they are actively processing a transaction, which usually take only milliseconds to complete. When not processing a transaction, the connection would otherwise sit idle. Instead, connection pooling allows the idle connection to be used by some other thread to do useful work.

In practice, when a thread needs to do work against a MySQL or other database with JDBC, it requests a connection from the pool. When the thread is finished using the connection, it returns it to the pool, so that it may be used by any other threads that want to use it.

When the connection is loaned out from the pool, it is used exclusively by the thread that requested it. From a programming point of view, it is the same as if your thread called `DriverManager.getConnection()` every time it needed a JDBC connection, however with connection pooling, your thread may end up using either a new, or already-existing connection.

Connection pooling can greatly increase the performance of your Java application, while reducing overall resource usage. The main benefits to connection pooling are:

- Reduced connection creation time

Although this is not usually an issue with the quick connection setup that MySQL offers compared



to other databases, creating new JDBC connections still incurs networking and JDBC driver overhead that will be avoided if connections are recycled.

- Simplified programming model

When using connection pooling, each individual thread can act as though it has created its own JDBC connection, allowing you to use straight-forward JDBC programming techniques.

- Controlled resource usage

If you don't use connection pooling, and instead create a new connection every time a thread needs one, your application's resource usage can be quite wasteful and lead to unpredictable behavior under load.

Remember that each connection to MySQL has overhead (memory, CPU, context switches, and so forth) on both the client and server side. Every connection limits how many resources there are available to your application as well as the MySQL server. Many of these resources will be used whether or not the connection is actually doing any useful work!

Connection pools can be tuned to maximize performance, while keeping resource utilization below the point where your application will start to fail rather than just run slower.

Luckily, Sun has standardized the concept of connection pooling in JDBC through the JDBC-2.0 Optional interfaces, and all major application servers have implementations of these APIs that work fine with MySQL Connector/J.

Generally, you configure a connection pool in your application server configuration files, and access it via the Java Naming and Directory Interface (JNDI). The following code shows how you might use a connection pool from an application deployed in a J2EE application server:

### Example 18.11. Using a connection pool with a J2EE application server

```
import java.sql.Connection;
import java.sql.SQLException;
import java.sql.Statement;

import javax.naming.InitialContext;
import javax.sql.DataSource;

public class MyServletJspOrEjb {

    public void doSomething() throws Exception {
        /*
         * Create a JNDI Initial context to be able to
         * lookup the DataSource
         *
         * In production-level code, this should be cached as
         * an instance or static variable, as it can
         * be quite expensive to create a JNDI context.
         *
         * Note: This code only works when you are using servlets
         * or EJBs in a J2EE application server. If you are
         * using connection pooling in standalone Java code, you
         * will have to create/configure datasources using whatever
         * mechanisms your particular connection pooling library
         * provides.
         */

        InitialContext ctx = new InitialContext();

        /*
         * Lookup the DataSource, which will be backed by a pool
         * that the application server provides. DataSource instances
         * are also a good candidate for caching as an instance
         */
    }
}
```

```

    * variable, as JNDI lookups can be expensive as well.
    */

    DataSource ds = (DataSource)ctx.lookup("java:comp/env/jdbc/MySQLDB");

    /*
    * The following code is what would actually be in your
    * Servlet, JSP or EJB 'service' method...where you need
    * to work with a JDBC connection.
    */

    Connection conn = null;
    Statement stmt = null;

    try {
        conn = ds.getConnection();

        /*
        * Now, use normal JDBC programming to work with
        * MySQL, making sure to close each resource when you're
        * finished with it, which allows the connection pool
        * resources to be recovered as quickly as possible
        */

        stmt = conn.createStatement();
        stmt.execute("SOME SQL QUERY");

        stmt.close();
        stmt = null;

        conn.close();
        conn = null;
    } finally {
        /*
        * close any jdbc instances here that weren't
        * explicitly closed during normal code path, so
        * that we don't 'leak' resources...
        */

        if (stmt != null) {
            try {
                stmt.close();
            } catch (SQLException sqlex) {
                // ignore -- as we can't do anything about it here
            }

            stmt = null;
        }

        if (conn != null) {
            try {
                conn.close();
            } catch (SQLException sqlex) {
                // ignore -- as we can't do anything about it here
            }

            conn = null;
        }
    }
}

```

As shown in the example above, after obtaining the JNDI InitialContext, and looking up the DataSource, the rest of the code should look familiar to anyone who has done JDBC programming in the past.

The most important thing to remember when using connection pooling is to make sure that no matter what happens in your code (exceptions, flow-of-control, and so forth), connections, and anything created by them (such as statements or result sets) are closed, so that they may be re-used, otherwise they will be stranded, which in the best case means that the MySQL server resources they represent (such as buffers, locks, or sockets) may be tied up for some time, or worst case, may be tied up forever.

What's the Best Size for my Connection Pool?

As with all other configuration rules-of-thumb, the answer is: it depends. Although the optimal size de-

depends on anticipated load and average database transaction time, the optimum connection pool size is smaller than you might expect. If you take Sun's Java Petstore blueprint application for example, a connection pool of 15-20 connections can serve a relatively moderate load (600 concurrent users) using MySQL and Tomcat with response times that are acceptable.

To correctly size a connection pool for your application, you should create load test scripts with tools such as Apache JMeter or The Grinder, and load test your application.

An easy way to determine a starting point is to configure your connection pool's maximum number of connections to be unbounded, run a load test, and measure the largest amount of concurrently used connections. You can then work backward from there to determine what values of minimum and maximum pooled connections give the best performance for your particular application.

### 18.3.5.2.2. Using Connector/J with Tomcat

The following instructions are based on the instructions for Tomcat-5.x, available at <http://jakarta.apache.org/tomcat/tomcat-5.0-doc/jndi-datasource-examples-howto.html> which is current at the time this document was written.

First, install the .jar file that comes with Connector/J in `$_CATALINA_HOME/common/lib` so that it is available to all applications installed in the container.

Next, Configure the JNDI DataSource by adding a declaration resource to `$_CATALINA_HOME/conf/server.xml` in the context that defines your web application:

```
<Context ....>
  ...
  <Resource name="jdbc/MySQLDB"
            auth="Container"
            type="javax.sql.DataSource"/>

  <!-- The name you used above, must match _exactly_ here!

         The connection pool will be bound into JNDI with the name
         "java:/comp/env/jdbc/MySQLDB"
  -->

  <ResourceParams name="jdbc/MySQLDB">
    <parameter>
      <name>factory</name>
      <value>org.apache.commons.dbcp.BasicDataSourceFactory</value>
    </parameter>

    <!-- Don't set this any higher than max_connections on your
         MySQL server, usually this should be a 10 or a few 10's
         of connections, not hundreds or thousands -->

    <parameter>
      <name>maxActive</name>
      <value>10</value>
    </parameter>

    <!-- You don't want to many idle connections hanging around
         if you can avoid it, only enough to soak up a spike in
         the load -->

    <parameter>
      <name>maxIdle</name>
      <value>5</value>
    </parameter>

    <!-- Don't use autoReconnect=true, it's going away eventually
         and it's a crutch for older connection pools that couldn't
         test connections. You need to decide whether your application is
         supposed to deal with SQLExceptions (hint, it should), and
         how much of a performance penalty you're willing to pay
         to ensure 'freshness' of the connection -->

    <parameter>
```

```
<name>validationQuery</name>
<value>SELECT 1</value>
</parameter>

<!-- The most conservative approach is to test connections
before they're given to your application. For most applications
this is okay, the query used above is very small and takes
no real server resources to process, other than the time used
to traverse the network.

If you have a high-load application you'll need to rely on
something else. -->

<parameter>
  <name>testOnBorrow</name>
  <value>true</value>
</parameter>

<!-- Otherwise, or in addition to testOnBorrow, you can test
while connections are sitting idle -->

<parameter>
  <name>testWhileIdle</name>
  <value>true</value>
</parameter>

<!-- You have to set this value, otherwise even though
you've asked connections to be tested while idle,
the idle evicter thread will never run -->

<parameter>
  <name>timeBetweenEvictionRunsMillis</name>
  <value>10000</value>
</parameter>

<!-- Don't allow connections to hang out idle too long,
never longer than what wait_timeout is set to on the
server...A few minutes or even fraction of a minute
is sometimes okay here, it depends on your application
and how much spikey load it will see -->

<parameter>
  <name>minEvictableIdleTimeMillis</name>
  <value>60000</value>
</parameter>

<!-- Username and password used when connecting to MySQL -->

<parameter>
  <name>username</name>
  <value>someuser</value>
</parameter>

<parameter>
  <name>password</name>
  <value>somepass</value>
</parameter>

<!-- Class name for the Connector/J driver -->

<parameter>
  <name>driverClassName</name>
  <value>com.mysql.jdbc.Driver</value>
</parameter>

<!-- The JDBC connection url for connecting to MySQL, notice
that if you want to pass any other MySQL-specific parameters
you should pass them here in the URL, setting them using the
parameter tags above will have no effect, you will also
need to use & to separate parameter values as the
ampersand is a reserved character in XML -->

<parameter>
  <name>url</name>
  <value>jdbc:mysql://localhost:3306/test</value>
</parameter>

</ResourceParams>
</Context>
```

In general, you should follow the installation instructions that come with your version of Tomcat, as the way you configure datasources in Tomcat changes from time-to-time, and unfortunately if you use the wrong syntax in your XML file, you will most likely end up with an exception similar to the following:

```
Error: java.sql.SQLException: Cannot load JDBC driver class 'null ' SQL
state: null
```

### 18.3.5.2.3. Using Connector/J with JBoss

These instructions cover JBoss-4.x. To make the JDBC driver classes available to the application server, copy the .jar file that comes with Connector/J to the `lib` directory for your server configuration (which is usually called `default`). Then, in the same configuration directory, in the subdirectory named `deploy`, create a datasource configuration file that ends with `-ds.xml`, which tells JBoss to deploy this file as a JDBC Datasource. The file should have the following contents:

```
<datasources>
  <local-tx-datasource>
    <!-- This connection pool will be bound into JNDI with the name
         "java:/MySQLDB" -->

    <jndi-name>MySQLDB</jndi-name>
    <connection-url>jdbc:mysql://localhost:3306/dbname</connection-url>
    <driver-class>com.mysql.jdbc.Driver</driver-class>
    <user-name>user</user-name>
    <password>pass</password>

    <min-pool-size>5</min-pool-size>

    <!-- Don't set this any higher than max_connections on your
         MySQL server, usually this should be a 10 or a few 10's
         of connections, not hundreds or thousands -->

    <max-pool-size>20</max-pool-size>

    <!-- Don't allow connections to hang out idle too long,
         never longer than what wait_timeout is set to on the
         server...A few minutes is usually okay here,
         it depends on your application
         and how much spikey load it will see -->

    <idle-timeout-minutes>5</idle-timeout-minutes>

    <!-- If you're using Connector/J 3.1.8 or newer, you can use
         our implementation of these to increase the robustness
         of the connection pool. -->

    <exception-sorter-class-name>com.mysql.jdbc.integration.jboss.ExtendedMysqlExceptionSorter</ex
    <valid-connection-checker-class-name>com.mysql.jdbc.integration.jboss.MysqlValidConnectionCheck

  </local-tx-datasource>
</datasources>
```

### 18.3.5.3. Common Problems and Solutions

There are a few issues that seem to be commonly encountered often by users of MySQL Connector/J. This section deals with their symptoms, and their resolutions.

#### Questions

- [19.3.5.3.1](#): When I try to connect to the database with MySQL Connector/J, I get the following exception:

```
SQLException: Server configuration denies access to data source
SQLState: 08001
VendorError: 0
```

What's going on? I can connect just fine with the MySQL command-line client.

- [19.3.5.3.2](#): My application throws an `SQLException` 'No Suitable Driver'. Why is this happening?
- [19.3.5.3.3](#): I'm trying to use MySQL Connector/J in an applet or application and I get an exception similar to:

```
SQLException: Cannot connect to MySQL server on host:3306.
Is there a MySQL server running on the machine/port you
are trying to connect to?

(java.security.AccessControlException)
SQLState: 08S01
VendorError: 0
```

- [19.3.5.3.4](#): I have a servlet/application that works fine for a day, and then stops working overnight
- [19.3.5.3.5](#): I'm trying to use JDBC-2.0 updatable result sets, and I get an exception saying my result set is not updatable.

## Questions and Answers

### 19.3.5.3.1: When I try to connect to the database with MySQL Connector/J, I get the following exception:

```
SQLException: Server configuration denies access to data source
SQLState: 08001
VendorError: 0
```

#### What's going on? I can connect just fine with the MySQL command-line client.

MySQL Connector/J must use TCP/IP sockets to connect to MySQL, as Java does not support Unix Domain Sockets. Therefore, when MySQL Connector/J connects to MySQL, the security manager in MySQL server will use its grant tables to determine whether the connection should be allowed.

You must add the necessary security credentials to the MySQL server for this to happen, using the [GRANT](#) statement to your MySQL Server. See [Section 13.5.1.2, "GRANT Syntax"](#), for more information.

**Note.** Testing your connectivity with the `mysql` command-line client will not work unless you add the `--host` flag, and use something other than `localhost` for the host. The `mysql` command-line client will use Unix domain sockets if you use the special hostname `localhost`. If you are testing connectivity to `localhost`, use `127.0.0.1` as the hostname instead.

**Warning.** Changing privileges and permissions improperly in MySQL can potentially cause your server installation to not have optimal security properties.

### 19.3.5.3.2: My application throws an `SQLException` 'No Suitable Driver'. Why is this happening?

There are three possible causes for this error:

- The Connector/J driver is not in your `CLASSPATH`, see [Section 18.3.2, "Installing Connector/J"](#).
- The format of your connection URL is incorrect, or you are referencing the wrong JDBC driver.
- When using `DriverManager`, the `jdbc.drivers` system property has not been populated with the location of the Connector/J driver.

### 19.3.5.3.3: I'm trying to use MySQL Connector/J in an applet or application and I get an exception similar to:

```
SQLException: Cannot connect to MySQL server on host:3306.
```

```
Is there a MySQL server running on the machine/port you
are trying to connect to?
```

```
(java.security.AccessControlException)
SQLState: 08S01
VendorError: 0
```

Either you're running an Applet, your MySQL server has been installed with the "--skip-networking" option set, or your MySQL server has a firewall sitting in front of it.

Applets can only make network connections back to the machine that runs the web server that served the .class files for the applet. This means that MySQL must run on the same machine (or you must have some sort of port re-direction) for this to work. This also means that you will not be able to test applets from your local file system, you must always deploy them to a web server.

MySQL Connector/J can only communicate with MySQL using TCP/IP, as Java does not support Unix domain sockets. TCP/IP communication with MySQL might be affected if MySQL was started with the "--skip-networking" flag, or if it is firewalled.

If MySQL has been started with the "--skip-networking" option set (the Debian Linux package of MySQL server does this for example), you need to comment it out in the file /etc/mysql/my.cnf or /etc/my.cnf. Of course your my.cnf file might also exist in the `data` directory of your MySQL server, or anywhere else (depending on how MySQL was compiled for your system). Binaries created by MySQL AB always look in /etc/my.cnf and [datadir]/my.cnf. If your MySQL server has been firewalled, you will need to have the firewall configured to allow TCP/IP connections from the host where your Java code is running to the MySQL server on the port that MySQL is listening to (by default, 3306).

#### 19.3.5.3.4: I have a servlet/application that works fine for a day, and then stops working overnight

MySQL closes connections after 8 hours of inactivity. You either need to use a connection pool that handles stale connections or use the "autoReconnect" parameter (see [Section 18.3.4.1, "Driver/Datasource Class Names, URL Syntax and Configuration Properties for Connector/J"](#)).

Also, you should be catching SQLExceptions in your application and dealing with them, rather than propagating them all the way until your application exits, this is just good programming practice. MySQL Connector/J will set the SQLState (see `java.sql.SQLException.getSQLState()` in your APIDOCs) to "08S01" when it encounters network-connectivity issues during the processing of a query. Your application code should then attempt to re-connect to MySQL at this point.

The following (simplistic) example shows what code that can handle these exceptions might look like:

#### Example 18.12. Example of transaction with retry logic

```
public void doBusinessOp() throws SQLException {
    Connection conn = null;
    Statement stmt = null;
    ResultSet rs = null;

    //
    // How many times do you want to retry the transaction
    // (or at least _getting_ a connection)?
    //
    int retryCount = 5;

    boolean transactionCompleted = false;

    do {
        try {
            conn = getConnection(); // assume getting this from a
                                   // javax.sql.DataSource, or the
                                   // java.sql.DriverManager

            conn.setAutoCommit(false);
```

```
//
// Okay, at this point, the 'retry-ability' of the
// transaction really depends on your application logic,
// whether or not you're using autocommit (in this case
// not), and whether you're using transaccational storage
// engines
//
// For this example, we'll assume that it's not safe
// to retry the entire transaction, so we set retry count
// to 0 at this point
//
// If you were using exclusively transaction-safe tables,
// or your application could recover from a connection going
// bad in the middle of an operation, then you would not
// touch 'retryCount' here, and just let the loop repeat
// until retryCount == 0.
//
retryCount = 0;

stmt = conn.createStatement();

String query = "SELECT foo FROM bar ORDER BY baz";

rs = stmt.executeQuery(query);

while (rs.next()) {
}

rs.close();
rs = null;

stmt.close();
stmt = null;

conn.commit();
conn.close();
conn = null;

transactionCompleted = true;
} catch (SQLException sqlEx) {

//
// The two SQL states that are 'retry-able' are 08S01
// for a communications error, and 40001 for deadlock.
//
// Only retry if the error was due to a stale connection,
// communications problem or deadlock
//

String sqlState = sqlEx.getSQLState();

if ("08S01".equals(sqlState) || "40001".equals(sqlState)) {
    retryCount--;
} else {
    retryCount = 0;
}
} finally {
    if (rs != null) {
        try {
            rs.close();
        } catch (SQLException sqlEx) {
            // You'd probably want to log this . . .
        }
    }

    if (stmt != null) {
        try {
            stmt.close();
        } catch (SQLException sqlEx) {
            // You'd probably want to log this as well . . .
        }
    }

    if (conn != null) {
        try {
            //
            // If we got here, and conn is not null, the
            // transaction should be rolled back, as not
            // all work has been done

```



```

        try {
            conn.rollback();
        } finally {
            conn.close();
        }
    } catch (SQLException sqlEx) {
        //
        // If we got an exception here, something
        // pretty serious is going on, so we better
        // pass it up the stack, rather than just
        // logging it. . .
        throw sqlEx;
    }
}
} while (!transactionCompleted && (retryCount > 0));
}
}

```

**Note.** Use of the `autoReconnect` option is not recommended because there is no safe method of re-connecting to the MySQL server without risking some corruption of the connection state or database state information. Instead, you should use a connection pool which will enable your application to connect to the MySQL server using an available connection from the pool. The `autoReconnect` facility is deprecated, and may be removed in a future release.

#### 19.3.5.3.5: I'm trying to use JDBC-2.0 updatable result sets, and I get an exception saying my result set is not updatable.

Because MySQL does not have row identifiers, MySQL Connector/J can only update result sets that have come from queries on tables that have at least one primary key, the query must select every primary key and the query can only span one table (that is, no joins). This is outlined in the JDBC specification.

Note that this issue only occurs when using updatable result sets, and is caused because Connector/J is unable to guarantee that it can identify the correct rows within the result set to be updated without having a unique reference to each row. There is no requirement to have a unique field on a table if you are using `UPDATE` or `DELETE` statements on a table where you can individually specify the criteria to be matched using a `WHERE` clause.

## 18.3.6. Connector/J Support

### 18.3.6.1. Connector/J Community Support

MySQL AB provides assistance to the user community by means of its mailing lists. For Connector/J related issues, you can get help from experienced users by using the MySQL and Java mailing list. Archives and subscription information is available online at <http://lists.mysql.com/java>.

For information about subscribing to MySQL mailing lists or to browse list archives, visit <http://lists.mysql.com/>. See Section 1.7.1, “MySQL Mailing Lists”.

Community support from experienced users is also available through the [JDBC Forum](http://forums.mysql.com/list.php?39) [http://forums.mysql.com/list.php?39]. You may also find help from other users in the other MySQL Forums, located at <http://forums.mysql.com>. See Section 1.7.2, “MySQL Community Support at the MySQL Forums”.

### 18.3.6.2. How to Report Connector/J Bugs or Problems

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

If you have found a sensitive security bug in MySQL, you can send email to [security\\_at\\_mysql.com](mailto:security_at_mysql.com) [mailto:security\_at\_mysql.com].

Writing a good bug report takes patience, but doing it right the first time saves time both for us and for yourself. A good bug report, containing a full test case for the bug, makes it very likely that we will fix the bug in the next release.

This section will help you write your report correctly so that you don't waste your time doing things that may not help us much or at all.

If you have a repeatable bug report, please report it to the bugs database at <http://bugs.mysql.com/>. Any bug that we are able to repeat has a high chance of being fixed in the next MySQL release.

To report other problems, you can use one of the MySQL mailing lists.

Remember that it is possible for us to respond to a message containing too much information, but not to one containing too little. People often omit facts because they think they know the cause of a problem and assume that some details don't matter.

A good principle is this: If you are in doubt about stating something, state it. It is faster and less troublesome to write a couple more lines in your report than to wait longer for the answer if we must ask you to provide information that was missing from the initial report.

The most common errors made in bug reports are (a) not including the version number of Connector/J or MySQL used, and (b) not fully describing the platform on which Connector/J is installed (including the JVM version, and the platform type and version number that MySQL itself is installed on).

This is highly relevant information, and in 99 cases out of 100, the bug report is useless without it. Very often we get questions like, "Why doesn't this work for me?" Then we find that the feature requested wasn't implemented in that MySQL version, or that a bug described in a report has already been fixed in newer MySQL versions.

Sometimes the error is platform-dependent; in such cases, it is next to impossible for us to fix anything without knowing the operating system and the version number of the platform.

If at all possible, you should create a repeatable, standalone testcase that doesn't involve any third-party classes.

To streamline this process, we ship a base class for testcases with Connector/J, named `com.mysql.jdbc.util.BaseBugReport`. To create a testcase for Connector/J using this class, create your own class that inherits from `com.mysql.jdbc.util.BaseBugReport` and override the methods `setUp()`, `tearDown()` and `runTest()`.

In the `setUp()` method, create code that creates your tables, and populates them with any data needed to demonstrate the bug.

In the `runTest()` method, create code that demonstrates the bug using the tables and data you created in the `setUp` method.

In the `tearDown()` method, drop any tables you created in the `setUp()` method.

In any of the above three methods, you should use one of the variants of the `getConnection()` method to create a JDBC connection to MySQL:

- `getConnection()` - Provides a connection to the JDBC URL specified in `getUrl()`. If a connection already exists, that connection is returned, otherwise a new connection is created.
- `getNewConnection()` - Use this if you need to get a new connection for your bug report (i.e.

there's more than one connection involved).

- `getConnection(String url)` - Returns a connection using the given URL.
- `getConnection(String url, Properties props)` - Returns a connection using the given URL and properties.

If you need to use a JDBC URL that is different from 'jdbc:mysql://test', override the method `getUrl()` as well.

Use the `assertTrue(boolean expression)` and `assertTrue(String failureMessage, boolean expression)` methods to create conditions that must be met in your testcase demonstrating the behavior you are expecting (vs. the behavior you are observing, which is why you are most likely filing a bug report).

Finally, create a `main()` method that creates a new instance of your testcase, and calls the `run` method:

```
public static void main(String[] args) throws Exception {
    new MyBugReport().run();
}
```

Once you have finished your testcase, and have verified that it demonstrates the bug you are reporting, upload it with your bug report to <http://bugs.mysql.com/>.

### 18.3.6.3. Connector/J Change History

The Connector/J Change History (Changelog) is located with the main Changelog for MySQL. See [Section D.8, “MySQL Connector/J Change History”](#).

## 18.4. MySQL Connector/MXJ

MySQL Connector/MXJ is a solution for deploying the MySQL database engine (`mysqld`) intelligently from within a Java package.

### 18.4.1. Introduction to Connector/MXJ

MySQL Connector/MXJ is a Java Utility package for deploying and managing a MySQL database. Deploying and using MySQL can be as easy as adding an additional parameter to the JDBC connection url, which will result in the database being started when the first connection is made. This makes it easy for Java developers to deploy applications which require a database by reducing installation barriers for their end-users.

MySQL Connector/MXJ makes the MySQL database appear to be a java-based component. It does this by determining what platform the system is running on, selecting the appropriate binary, and launching the executable. It will also optionally deploy an initial database, with any specified parameters.

Included are instructions for use with a JDBC driver and deploying as a JMX MBean to JBoss.

You can download sources and binaries from: <http://dev.mysql.com/downloads/connector/mxj/>

This a beta release and feedback is welcome and encouraged.

Please send questions or comments to the [MySQL and Java mailing list](http://lists.mysql.com/java) [<http://lists.mysql.com/java>].

#### 18.4.1.1. Connector/MXJ Versions

## Connector/MX

- Connector/MXJ 5.x, currently in beta status, includes `mysqld` version 5.0.22 and includes binaries for Linux x86, Mac OS X PPC, Windows XP/NT/2000 x86 and Solaris SPARC. Connector/MXJ 5.x requires the Connector/J 5.x package.
- Connector/MXJ 1.x includes `mysqld` version 4.1.13 and includes binaries for Linux x86, Windows XP/NT/2000 x86 and Solaris SPARC. Connector/MXJ 1.x requires the Connector/J 3.x package.

This guide provides information on the Connector/MXJ 5.x release. For information on using the older releases, please see the documentation included with the appropriate distribution.

### 18.4.1.2. Connector/MXJ Overview

Connector/MXJ consists of a Java class, a copy of the `mysqld` binary for a specific list of platforms, and associated files and support utilities. The Java class controls the initialization of an instance of the embedded `mysqld` binary, and the ongoing management of the `mysqld` process. The entire sequence and management can be controlled entirely from within Java using the Connector/MXJ Java classes. You can see an overview of the contents of the Connector/MXJ package in the figure below.

Connector/MXJ Package		
Connector/MXJ Management Class	mysqld Binaries	MySQL Base Databases
	Linux x86	MySQL Stub Data
	Windows XP/NT/2000	
	Solaris SPARC	Support Tools
	MacOS X PowerPC	

It is important to note that Connector/MXJ is not an embedded version of MySQL, or a version of MySQL written as part of a Java class. Connector/MXJ works through the use of an embedded, compiled binary of `mysqld` as would normally be used when deploying a standard MySQL installation.

It is the Connector/MXJ wrapper, support classes and tools, that enable Connector/MXJ to appear as a MySQL instance.

When Connector/MXJ is initialized, the corresponding `mysqld` binary for the current platform is extracted, along with a pre-configured data directed. Both are contained within the Connector/MXJ JAR file. The `mysqld` instance is then started, with any additional options as specified during the initialization, and the MySQL database becomes accessible.

Because Connector/MXJ works in combination with Connector/J, you can access and integrate with the MySQL instance through a JDBC connection. When you have finished with the server, the instance is terminated, and, by default, any data created during the session is retained within the temporary directory created when the instance was started.

Connector/MXJ and the embedded `mysqld` instance can be deployed in a number of environments where relying on an existing database, or installing a MySQL instance would be impossible, including CD-ROM embedded database applications and temporary database requirements within a Java-based application environment.

## 18.4.2. Installing Connector/MXJ

Connector/MXJ does not have an installation application or process, but there are some steps you can follow to make the installation and deployment of Connector/MXJ easier.

Before you start, there are some baseline requirements for

- Java Runtime Environment (v1.4.0 or newer) if you are only going to deploy the package.
- Java Development Kit (v1.4.0 or newer) if you want to build Connector/MXJ from source.
- Connector/J 5.0 or newer.

Depending on your target installation/deployment environment you may also require:

- JBoss - 4.0rc1 or newer
- Apache Tomcat - 5.0 or newer
- Sun's JMX reference implementation version 1.2.1 (from <http://java.sun.com/products/JavaManagement/>)

### 18.4.2.1. Supported Platforms

Connector/MXJ is compatible with any platform supporting Java and MySQL. By default, Connector/MXJ incorporates the `mysqld` binary for a select number of platforms, as outlined below.

- Linux, i386
- Windows NT, Windows 2000, Windows XP, x86
- Solaris 8, SPARC 32 (compatible with Solaris 8, Solaris 9 and Solaris 10 on SPARC 32-bit and 64-bit platforms)
- Mac OS X, PowerPC

For more information on packaging your own Connector/MXJ with the platforms you require, see [Section 18.4.5.1, “Creating your own Connector/MXJ Package”](#)

### 18.4.2.2. Connector/MXJ Base Installation

Because there is no formal installation process, the method, installation directory, and access methods you use for Connector/MXJ are entirely up to your individual requirements.

To perform a basic installation, choose a target directory for the files included in the Connector/MXJ package. On Unix/Linux systems you may opt to use a directory such as `/usr/local/connector-mxj`; On Windows, you may want to install the files in the base directory, `C:\Connector-MXJ`, or within the `Program Files` directory.

To install the files:

1. Download the Connector/MXJ package, either in Tar/Gzip format (ideal for Unix/Linux systems)

or Zip format (Windows).

2. Extract the files from the package. This will create a directory `connector-mxj`. Copy and optionally rename this directory to your desired location.
3. For best results, you should update your global `CLASSPATH` variable with the location of the Connector/MXJ JAR (`connector-mxj.jar`). You will also need to add the AspectJ Runtime, located in `lib/aspectjrt.jar`.

Within Unix/Linux you can do this globally by editing the global shell profile, or on a user by user basis by editing their individual shell profile.

On Windows 2000, Windows NT and Windows XP, you can edit the global `CLASSPATH` by editing the `Environment Variables` configured through the `System` control panel.

### 18.4.2.3. Connector/MXJ Quick Start Guide

Once you have extracted the Connector/MXJ and Connector/J components you can run one of the sample applications that initiates a MySQL instance. You can test the installation by running the `ConnectorMXJUrlTestExample`:

```
java ConnectorMXJUrlTestExample
jdbc:mysql:mxj://localhost:3336/test?server.baseDir=/var/tmp/test-mxj
[MysqldResource] launching mysqld (getOptions)
[/var/tmp/test-mxj/bin/mysqld][--no-defaults][--pid-file=/var/tmp/test-mxj/data/MysqldResource.pid][--
[MysqldResource] launching mysqld (driver_launched_mysqld_1)
InnoDB: The first specified data file ./ibdata1 did not exist:
InnoDB: a new database to be created!
060726 15:40:42 InnoDB: Setting file ./ibdata1 size to 10 MB
InnoDB: Database physically writes the file full: wait...
060726 15:40:43 InnoDB: Log file ./ib_logfile0 did not exist: new to be created
InnoDB: Setting log file ./ib_logfile0 size to 5 MB
InnoDB: Database physically writes the file full: wait...
060726 15:40:43 InnoDB: Log file ./ib_logfile1 did not exist: new to be created
InnoDB: Setting log file ./ib_logfile1 size to 5 MB
InnoDB: Database physically writes the file full: wait...
InnoDB: Doublewrite buffer not found: creating new
InnoDB: Doublewrite buffer created
InnoDB: Creating foreign key constraint system tables
InnoDB: Foreign key constraint system tables created
060726 15:40:44 InnoDB: Started; log sequence number 0 0
060726 15:40:44 [Note] /var/tmp/test-mxj/bin/mysqld: ready for connections.
Version: '5.0.22-max' socket: 'mysql.sock' port: 3336 MySQL Community Edition - Experimental (GPL)
[MysqldResource] mysqld running as process: 1210
-----
5.0.22-max
-----
[MysqldResource] stopping mysqld (process: 1210)
060726 15:40:44 [Note] /var/tmp/test-mxj/bin/mysqld: Normal shutdown

060726 15:40:45 InnoDB: Starting shutdown...
060726 15:40:48 InnoDB: Shutdown completed; log sequence number 0 43655
060726 15:40:48 [Note] /var/tmp/test-mxj/bin/mysqld: Shutdown complete

[MysqldResource] clearing options
[MysqldResource] shutdown complete
```

The above output shows an instance of MySQL starting, the necessary files being created (log files, InnoDB data files) and the MySQL database entering the running state. The instance is then shutdown by Connector/MXJ before the example terminates.

### 18.4.2.4. Deploying Connector/MXJ using Driver Launch

Connector/MXJ and Connector/J work together to enable you to launch an instance of the `mysqld`

server through the use of a keyword in the JDBC connection string. Deploying Connector/MXJ within a Java application can be automated through this method, making the deployment of Connector/MXJ a simple process:

1. Download and unzip Connector/MXJ, add `connector-mxj.jar` to the `CLASSPATH`.
2. To the JDBC connection string, embed the `mxj` keyword, for example: `jdbc:mysql:mxj://localhost:PORT/DBNAME`.

For more details, see [Section 18.4.3, “Connector/MXJ Configuration”](#).

### 18.4.2.5. Deploying Connector/MXJ within JBoss

For deployment within a JBoss environment, you must configure the JBoss environment to use the Connector/MXJ component within the JDBC parameters:

1. Download Connector/MXJ copy the `connector-mxj.jar` file to the `$JBOSS_HOME/server/default/lib` directory.
2. Download Connector/J copy the `connector-mxj.jar` file to the `$JBOSS_HOME/server/default/lib` directory.
3. Create an MBean service xml file in the `$JBOSS_HOME/server/default/deploy` directory with any attributes set, for instance the `datadir` and `autostart`.
4. Set the JDBC parameters of your web application to use:

```
String driver = "com.mysql.jdbc.Driver";
String url = "jdbc:mysql:///test?propertiesTransform="+
            "com.mysql.management.jmx.ConnectorMXJPropertiesTransform";
String user = "root";
String password = "";
Class.forName(driver);
Connection conn = DriverManager.getConnection(url, user, password);
```

You may wish to create a separate users and database table spaces for each application, rather than using "root and test".

We highly suggest having a routine backup procedure for backing up the database files in the `datadir`.

### 18.4.2.6. Verifying Installation using JUnit

The best way to ensure that your platform is supported is to run the JUnit tests. These will test the Connector/MXJ classes and the associated components.

#### 18.4.2.6.1. JUnit Test Requirements

The first thing to do is make sure that the components will work on the platform. The `MySQLdRe-source` class is really a wrapper for a native version of MySQL, so not all platforms are supported. At the time of this writing, Linux on the i386 architecture has been tested and seems to work quite well, as does OS X v10.3. There has been limited testing on Windows and Solaris.

Requirements:

1. JDK-1.4 or newer (or the JRE if you aren't going to be compiling the source or JSPs).

2. MySQL Connector/J version 5.0 or newer (from <http://dev.mysql.com/downloads/connector/j/>) installed and available via your CLASSPATH.
3. The `javax.management` classes for JMX version 1.2.1, these are present in the following application servers:
  - JBoss - 4.0rc1 or newer.
  - Apache Tomcat - 5.0 or newer.
  - Sun's JMX reference implementation version 1.2.1 (from <http://java.sun.com/products/JavaManagement/>).
4. JUnit 3.8.1 (from <http://www.junit.org/>).

If building from source, All of the requirements from above, plus:

1. Ant version 1.5 or newer (download from <http://ant.apache.org/>).

#### 18.4.2.6.2. Running the JUnit Tests

1. The tests attempt to launch MySQL on the port 3336. If you have a MySQL running, it may conflict, but this isn't very likely because the default port for MySQL is 3306. However, You may set the "c-mxj\_test\_port" Java property to a port of your choosing. Alternatively, you may wish to start by shutting down any instances of MySQL you have running on the target machine.

The tests suppress output to the console by default. For verbose output, you may set the "c-mxj\_test\_silent" Java property to "false".

2. To run the JUnit test suite, the \$CLASSPATH must include the following:
  - JUnit
  - JMX
  - Connector/J
  - MySQL Connector/MXJ
3. If `connector-mxj.jar` is not present in your download, unzip MySQL Connector/MXJ source archive.

```
cd mysql djmx
ant dist
```

Then add `$TEMP/cmjxj/stage/connector-mxj/connector-mxj.jar` to the CLASSPATH.

4. if you have `junit`, execute the unit tests. From the command line, type:

```
java junit.textui.TestRunner com.mysql.management.AllTestsSuite
```

The output should look something like this:



```

.....
.....
Time: 259.438
OK (101 tests)

```

Note that the tests are a bit slow near the end, so please be patient.

## 18.4.3. Connector/MXJ Configuration

### 18.4.3.1. Running as part of the JDBC Driver

A feature of the MySQL Connector/J JDBC driver is the ability to specify a connection to an embedded Connector/MXJ instance through the use of the `mxj` keyword in the JDBC connection string.

In the following example, we have a program which creates a connection, executes a query, and prints the result to the `System.out`. The MySQL database will be deployed and started as part of the connection process, and shutdown as part of the finally block.

You can find this file in the Connector/MXJ package as `src/ConnectorMXJUrlTestExample.java`.

```

import java.io.File;
import java.sql.Connection;
import java.sql.DriverManager;
import java.sql.ResultSet;
import java.sql.Statement;

import com.mysql.management.driverlaunched.ServerLauncherSocketFactory;

public class ConnectorMXJUrlTestExample {
    public static String DRIVER = "com.mysql.jdbc.Driver";

    public static String JAVA_IO_TMPDIR = "java.io.tmpdir";

    public static void main(String[] args) throws Exception {
        File ourAppDir = new File(System.getProperty(JAVA_IO_TMPDIR));
        File databaseDir = new File(ourAppDir, "test-mxj");
        int port = 3336;

        String url = "jdbc:mysql:mxj://localhost:" + port + "/test" + "?"
            + "serverbasedir=" + databaseDir;

        System.out.println(url);

        String userName = "root";
        String password = "";

        Class.forName(DRIVER);
        Connection conn = null;
        try {
            conn = DriverManager.getConnection(url, userName, password);
            printQueryResults(conn, "SELECT VERSION()");
        } finally {
            try {
                if (conn != null)
                    conn.close();
            } catch (Exception e) {
                e.printStackTrace();
            }

            ServerLauncherSocketFactory.shutdown(databaseDir, null);
        }
    }

    public static void printQueryResults(Connection conn, String SQLquery)
        throws Exception {
        Statement stmt = conn.createStatement();

```

```

ResultSet rs = stmt.executeQuery(SQLquery);
int columns = rs.getMetaData().getColumnCount();
System.out.println("-----");
System.out.println();
while (rs.next()) {
    for (int i = 1; i <= columns; i++) {
        System.out.println(rs.getString(i));
    }
    System.out.println();
}
rs.close();
stmt.close();
System.out.println("-----");
System.out.flush();
Thread.sleep(100); // wait for System.out to finish flush
}
}

```

To run the above program, be sure to have connector-mxj.jar and Connector/J in the CLASSPATH. Then type:

```
java ConnectorMXJTestExample
```

### 18.4.3.2. Running within a Java Object

If you have a java application and wish to “embed” a MySQL database, make use of the `com.mysql.management.MysqldResource` class directly. This class may be instantiated with the default (no argument) constructor, or by passing in a `java.io.File` object representing the directory you wish the server to be “unzipped” into. It may also be instantiated with printstreams for “stdout” and “stderr” for logging.

Once instantiated, a `java.util.Map`, the object will be able to provide a `java.util.Map` of server options appropriate for the platform and version of MySQL which you will be using.

The `MysqldResource` enables you to “start” MySQL with a `java.util.Map` of server options which you provide, as well as “shutdown” the database. The following example shows a simplistic way to embed MySQL in an application using plain java objects.

You can find this file in the Connector/MXJ package as [src/ConnectorMXJObjectTestExample.java](#).

```

import java.io.File;
import java.sql.Connection;
import java.sql.DriverManager;
import java.sql.ResultSet;
import java.sql.Statement;
import java.util.HashMap;
import java.util.Map;

import com.mysql.management.MysqldResource;

public class ConnectorMXJObjectTestExample {
    public static String DRIVER = "com.mysql.jdbc.Driver";

    public static String JAVA_IO_TMPDIR = "java.io.tmpdir";

    public static void main(String[] args) throws Exception {
        File ourAppDir = new File(System.getProperty(JAVA_IO_TMPDIR));
        File databaseDir = new File(ourAppDir, "mysql-mxj");
        int port = 3336;

        MysqldResource mysqldResource = startDatabase(databaseDir, port);

        String userName = "root";
        String password = "";

        Class.forName(DRIVER);
        Connection conn = null;
        try {

```

```

        String url = "jdbc:mysql://localhost:" + port + "/test";
        conn = DriverManager.getConnection(url, userName, password);
        printQueryResults(conn, "SELECT VERSION()");
    } finally {
        try {
            if (conn != null) {
                conn.close();
            }
        } catch (Exception e) {
            e.printStackTrace();
        }
        try {
            mysqldResource.shutdown();
        } catch (Exception e) {
            e.printStackTrace();
        }
    }
}

public static MysqldResource startDatabase(File databaseDir, int port) {
    MysqldResource mysqldResource = new MysqldResource(databaseDir);

    Map database_options = new HashMap();
    database_options.put("port", Integer.toString(port));
    mysqldResource.start("test-mysqld-thread", database_options);

    if (!mysqldResource.isRunning()) {
        throw new RuntimeException("MySQL did not start.");
    }

    System.out.println("MySQL is running.");

    return mysqldResource;
}

public static void printQueryResults(Connection conn, String SQLquery)
    throws Exception {
    Statement stmt = conn.createStatement();
    ResultSet rs = stmt.executeQuery(SQLquery);
    int columns = rs.getMetaData().getColumnCount();
    System.out.println("-----");
    System.out.println();
    while (rs.next()) {
        for (int i = 1; i <= columns; i++) {
            System.out.println(rs.getString(i));
        }
        System.out.println();
    }
    rs.close();
    stmt.close();
    System.out.println("-----");
    System.out.flush();
    Thread.sleep(100); // wait for System.out to finish flush
}
}

```

### 18.4.3.3. Setting server options

Of course there are many options we may wish to set for a MySQL database. These options may be specified as part of the JDBC connection string simply by prefixing each server option with "server.". In the following example we set two driver parameters and two server parameters:

```

String url = "jdbc:mysql://" + hostColonPort + "/"
    + "?"
    + "cacheServerConfiguration=true"
    + "&"
    + "useLocalSessionState=true"
    + "&"
    + "server.basedir=/opt/myapp/db"
    + "&"
    + "server.datadir=/mnt/bigdisk/myapp/data";

```

## 18.4.4. Connector/MXJ Reference

### 18.4.4.1. MysqldResource API

#### 18.4.4.1.1. MysqldResource Constructors

The `MysqldResource` class supports three different constructor forms:

- `public MysqldResource(File baseDir, File dataDir, String mysqlVersionString, PrintStream out, PrintStream err, Utils util)`

The most detailed constructor, enables you to set the base directory, data directory, select a server by its version string, standard out and standard error and MySQL utilities class.

- `public MysqldResource(File baseDir, File dataDir, String mysqlVersionString, PrintStream out, PrintStream err)`

Enables you to set the base directory, data directory, select a server by its version string, standard out and standard error.

- `public MysqldResource(File baseDir, File dataDir, String mysqlVersionString)`

Enables you to set the base directory, data directory and select a server by its version string. Output for standard out and standard err are directed to `System.out` and `System.err`.

- `public MysqldResource(File baseDir, File dataDir)`

Enables you to set the base directory and data directory. The default MySQL version is selected, and output for standard out and standard err are directed to `System.out` and `System.err`.

- `public MysqldResource(File baseDir);`

Allows the setting of the "basedir" to deploy the MySQL files to. Output for standard out and standard err are directed to `System.out` and `System.err`.

- `public MysqldResource();`

The basedir is defaulted to a subdirectory of the `java.io.tmpdir`. Output for standard out and standard err are directed to `System.out` and `System.err`;

#### 18.4.4.1.2. MysqldResource Methods

`MysqldResource` API includes the following methods:

- `void start(String threadName, Map mysqlArgs);`

Deploys and starts MySQL. The "threadName" string is used to name the thread which actually performs the execution of the MySQL command line. The map is the set of arguments and their values to be passed to the command line.

- `void shutdown();`

Shuts down the MySQL instance managed by the `MysqldResource` object.

- `Map getServerOptions();`

Returns a map of all the options and their current (or default, if not running) options available for the MySQL database.

- `boolean isRunning();`

Returns true if the MySQL database is running.

- `boolean isReadyForConnections();`

Returns true once the database reports that is ready for connections.

- `void setKillDelay(int millis);`

The default “Kill Delay” is 30 seconds. This represents the amount of time to wait between the initial request to shutdown and issuing a “force kill” if the database has not shutdown by itself.

- `void addCompletionListener(Runnable listener);`

Allows for applications to be notified when the server process completes. Each "listener" will be fired off in its own thread.

- `String getVersion();`

Returns the version of MySQL.

- `void setVersion(int MajorVersion, int minorVersion, int patch-Level);`

The standard distribution comes with only one version of MySQL packaged. However, it is possible to package multiple versions, and specify which version to use.

## 18.4.5. Connector/MXJ Notes and Tips

This section contains notes and tips on using the Connector/MXJ component within your applications.

### 18.4.5.1. Creating your own Connector/MXJ Package

If you want to create a custom Connector/MXJ package that includes a specific `mysqld` version or platform then you must extract and rebuild the `connector-mxj.jar` file.

First, you should create a new directory into which you can extract the current `connector-mxj.jar`:

```
shell> mkdir custom-mxj
shell> cd custom-mxj
shell> jar -xf connector-mxj.jar
shell> ls
5-0-22/
ConnectorMXJObjectTestExample.class
ConnectorMXJUrlTestExample.class
META-INF/
TestDb.class
com/
kill.exe
```

The MySQL version directory, `5-0-22` in the above example, contains all of the files used to create an instance of MySQL when Connector/MXJ is executed. All of the files in this directory are required for each version of MySQL that you want to embed. Note as well the format of the version number, which uses hyphens instead of periods to separate the version number components.

Within the version specific directory are the platform specific directories, and archives of the [data](#) and [share](#) directory required by MySQL for the various platforms. For example, here is the listing for the default Connector/MXJ package:

```
shell>> ls
Linux-i386/
META-INF/
Mac_OS_X-ppc/
SunOS-sparc/
Win-x86/
com/
data_dir.jar
share_dir.jar
win_share_dir.jar
```

Platform specific directories are listed by their OS and platform - for example the [mysqld](#) for Mac OS X PowerPC is located within the [Mac\\_OS\\_X-ppc](#) directory. You can delete directories from this location that you do not require, and add new directories for additional platforms that you want to support.

To add a platform specific [mysqld](#), create a new directory with the corresponding name for your operating system/platform. For example, you could add a directory for Mac OS X/Intel using the directory [Mac\\_OS\\_X-i386](#).

On Unix systems, you can determine the platform using [uname](#):

```
shell> uname -p
i386
```

Now you need to download or compile [mysqld](#) for the MySQL version and platform you want to include in your custom [connector-mxj.jar](#) package into the new directory.

Create a file called [version.txt](#) in the OS/platform directory you have just created that contains the version string/path of the [mysqld](#) binary. For example:

```
mysql-5.0.22-osx10.3-i386/bin/mysqld
```

You can now recreate the [connector-mxj.jar](#) file with the added [mysqld](#):

```
shell> cd custom-mxj
shell> jar -cf ../connector-mxj.jar *
```

You should test this package using the steps outlined in [Section 18.4.2.3, “Connector/MXJ Quick Start Guide”](#).

## 18.4.5.2. Deploying Connector/MXJ with a pre-configured database

To include a pre-configured/populated database within your Connector/MXJ JAR file you must create a custom [data\\_dir.jar](#) file, as included within the main [connector-mxj.jar](#) file:

1. First extract the [connector-mxj.jar](#) file, as outlined in the previous section (see [Section 18.4.5.1, “Creating your own Connector/MXJ Package”](#)).
2. First, create your database and populate the database with the information you require in an existing instance of MySQL - including Connector/MXJ instances. Data file formats are compatible across platforms.
3. Shutdown the instance of MySQL.
4. Create a JAR file of the data directory and databases that you want to include your Connector/MXJ

package. You should include the `mysql` database, which includes user authentication information, in addition to the specific databases you want to include. For example, to create a JAR of the `mysql` and `mxjtest` databases:

```
shell> jar -cf ../data_dir.jar mysql mxjtest
```

5. Copy the `data_dir.jar` file into the extracted `connector-mxj.jar` directory, and then create an archive for `connector-mxj.jar`.

Note that if you are create databases using the InnoDB engine, you must include the `ibdata.*` and `ib_logfile*` files within the `data_dir.jar` archive.

### 18.4.5.3. Running within a JMX Agent (custom)

As a JMX MBean, MySQL Connector/MXJ requires a JMX v1.2 compliant MBean container, such as JBoss version 4. The MBean will uses the standard JMX management APIs to present (and allow the setting of) parameters which are appropriate for that platform.

If you are not using the SUN Reference implementation of the JMX libraries, you should skip this section. Or, if you are deploying to JBoss, you also may wish to skip to the next section.

We want to see the `MysqldDynamicMBean` in action inside of a JMX agent. In the `com.mysql.management.jmx.sunri` package is a custom JMX agent with two MBeans:

1. the `MysqldDynamicMBean`, and
2. a `com.sun.jdmk.comm.HtmlAdaptorServer`, which provides a web interface for manipulating the beans inside of a JMX agent.

When this very simple agent is started, it will allow a MySQL database to be started and stopped with a web browser.

1. Complete the testing of the platform as above.
  - current JDK, JUnit, Connector/J, MySQL Connector/MXJ
  - this section *requires* the SUN reference implementation of JMX
  - `PATH`, `JAVA_HOME`, `ANT_HOME`, `CLASSPATH`
2. If not building from source, skip to next step

rebuild with the "sunri.present"

```
ant -Dsunri.present=true dist
re-run tests:
java junit.textui.TestRunner com.mysql.management.AllTestsSuite
```

3. launch the test agent from the command line:

```
java com.mysql.management.jmx.sunri.MysqldTestAgentSunHtmlAdaptor &
```

4. from a browser:

```
http://localhost:9092/
```

5. under MysqlIdAgent,

```
select "name=mysqlId"
```

6. Observe the MBean View
7. scroll to the bottom of the screen press the `startMysqlId` button
8. click `Back to MBean View`
9. scroll to the bottom of the screen press `stopMysqlId` button
10. kill the java process running the Test Agent (jmx server)

#### 18.4.5.4. Deployment in a standard JMX Agent environment (JBoss)

Once there is confidence that the MBean will function on the platform, deploying the MBean inside of a standard JMX Agent is the next step. Included are instructions for deploying to JBoss.

1. Ensure a current version of java development kit (v1.4.x), see above.
  - Ensure `JAVA_HOME` is set (JBoss requires `JAVA_HOME`)
  - Ensure `JAVA_HOME/bin` is in the `PATH` (You will NOT need to set your `CLASSPATH`, nor will you need any of the jars used in the previous tests).
2. Ensure a current version of JBoss (v4.0RC1 or better)

```
http://www.jboss.org/index.html
select "Downloads"
select "jboss-4.0.zip"
pick a mirror
unzip ~/dload/jboss-4.0.zip
create a JBOSS_HOME environment variable set to the unzipped directory
unix only:
cd $JBOSS_HOME/bin
chmod +x *.sh
```

3. Deploy (copy) the `connector-mxj.jar` to `$JBOSS_HOME/server/default/lib`.
4. Deploy (copy) `mysql-connector-java-3.1.4-beta-bin.jar` to `$JBOSS_HOME/server/default/lib`.
5. Create a `mxjtest.war` directory in `$JBOSS_HOME/server/default/deploy`.
6. Deploy (copy) `index.jsp` to `$JBOSS_HOME/server/default/deploy/mxjtest.war`.
7. Create a `mysqlId-service.xml` file in `$JBOSS_HOME/server/default/deploy`.

```
<?xml version="1.0" encoding="UTF-8"?>
<server>
  <mbean code="com.mysql.management.jmx.jboss.JBossMysqlIdDynamicMBean"
    name="mysql:type=service,name=mysqlId">
    <attribute name="datadir">/tmp/xxx_data_xxx</attribute>
    <attribute name="autostart">true</attribute>
  </mbean>
```



```
</server>
```

8. Start jboss:
  - on unix: `$JBOSS_HOME/bin/run.sh`
  - on windows: `%JBOSS_HOME%\bin\run.bat`

Be ready: JBoss sends a lot of output to the screen.
9. When JBoss seems to have stopped sending output to the screen, open a web browser to: <http://localhost:8080/jmx-console>
10. Scroll down to the bottom of the page in the `mysql` section, select the bulleted `mysqld` link.
11. Observe the JMX MBean View page. MySQL should already be running.
12. (If "autostart=true" was set, you may skip this step.) Scroll to the bottom of the screen. You may press the `Invoke` button to stop (or start) MySQL observe `Operation completed successfully without a return value`. Click `Back to MBean View`
13. To confirm MySQL is running, open a web browser to <http://localhost:8080/mxjtest/> and you should see that

```
SELECT 1
```

returned with a result of

```
1
```

14. Guided by the `$JBOSS_HOME/server/default/deploy/mxjtest.war/index.jsp` you will be able to use MySQL in your Web Application. There is a `test` database and a `root` user (no password) ready to experiment with. Try creating a table, inserting some rows, and doing some selects.
15. Shut down MySQL. MySQL will be stopped automatically when JBoss is stopped, or: from the browser, scroll down to the bottom of the MBean View press the stop service `Invoke` button to halt the service. Observe `Operation completed successfully without a return value`. Using `ps` or `task manager` see that MySQL is no longer running

As of 1.0.6-beta version is the ability to have the MBean start the MySQL database upon start up. Also, we've taken advantage of the JBoss life-cycle extension methods so that the database will gracefully shut down when JBoss is shutdown.

## 18.4.6. Connector/MXJ Support

There are a wide variety of options available for obtaining support for using Connector/MXJ. You should contact the Connector/MXJ community for help before reporting a potential bug or problem. See [Section 18.4.6.1, "Connector/MXJ Community Support"](#).

### 18.4.6.1. Connector/MXJ Community Support

MySQL AB provides assistance to the user community by means of a number of mailing lists and web based forums.

You can find help and support through the [MySQL and Java](http://lists.mysql.com/java) [http://lists.mysql.com/java] mailing list.

For information about subscribing to MySQL mailing lists or to browse list archives, visit <http://lists.mysql.com/>. See [Section 1.7.1, “MySQL Mailing Lists”](#).

Community support from experienced users is also available through the [MyODBC Forum](http://forums.mysql.com/list.php?39) [http://forums.mysql.com/list.php?39]. You may also find help from other users in the other MySQL Forums, located at <http://forums.mysql.com>. See [Section 1.7.2, “MySQL Community Support at the MySQL Forums”](#).

### 18.4.6.2. How to Report Connector/MXJ Problems

If you encounter difficulties or problems with Connector/MXJ, contact the Connector/MXJ community [Section 18.4.6.1, “Connector/MXJ Community Support”](#).

If reporting a problem, you should ideally include the following information with the email:

- Operating system and version
- Connector/MXJ version
- MySQL server version
- Copies of error messages or other unexpected output
- Simple reproducible sample

Remember that the more information you can supply to us, the more likely it is that we can fix the problem.

If you believe the problem to be a bug, then you must report the bug through <http://bugs.mysql.com/>.

## 18.5. Connector/PHP

The PHP distribution and documentation are available from the PHP Web site. MySQL provides the `mysql` and `mysql_i` extensions for MySQL versions as of 4.1.16 for the Windows operating system on <http://dev.mysql.com/downloads/connector/php/>. You can find information why you should preferably use the extensions provided by MySQL on that page. For platforms other than Windows, you should use the `mysql` or `mysql_i` extensions shipped with the PHP sources. See [Section 17.3, “MySQL PHP API”](#).

---

# Chapter 19. Extending MySQL

## 19.1. MySQL Internals

This chapter describes a lot of things that you need to know when working on the MySQL code. If you plan to contribute to MySQL development, want to have access to the bleeding-edge versions of the code, or just want to keep track of development, follow the instructions in [Section 2.9.3, “Installing from the Development Source Tree”](#). If you are interested in MySQL internals, you should also subscribe to our `internals` mailing list. This list has relatively low traffic. For details on how to subscribe, please see [Section 1.7.1, “MySQL Mailing Lists”](#). All developers at MySQL AB are on the `internals` list and we help other people who are working on the MySQL code. Feel free to use this list both to ask questions about the code and to send patches that you would like to contribute to the MySQL project!

### 19.1.1. MySQL Threads

The MySQL server creates the following threads:

- One thread manages TCP/IP file connection requests and creates a new dedicated thread to handle the authentication and SQL statement processing for each connection. (On Unix, this thread also manages Unix socket file connection requests.) On Windows, a similar thread manages shared-memory connection requests, and on Windows NT-based systems, a thread manages named-pipe connection requests. Every client connection has its own thread, although the manager threads try to avoid creating threads by consulting the thread cache first to see whether a cached thread can be used for a new connection.
- On Windows NT, there is a named pipe handler thread that does the same work as the TCP/IP connection thread on named pipe connect requests.
- On a master replication server, slave server connections are like client connections: There is one thread per connected slave.
- On a slave replication server, an I/O thread is started to connect to the master server and read updates from it. An SQL thread is started to apply updates read from the master. These two threads run independently and can be started and stopped independently.
- The signal thread handles all signals. This thread also normally handles alarms and calls `process_alarm()` to force timeouts on connections that have been idle too long.
- If `mysqld` is compiled with `-DUSE_ALARM_THREAD`, a dedicated thread that handles alarms is created. This is only used on some systems where there are problems with `sigwait()` or if you want to use the `thr_alarm()` code in your application without a dedicated signal handling thread.
- If the server is started with the `--flush_time=val` option, a dedicated thread is created to flush all tables every `val` seconds.
- Each table for which `INSERT DELAYED` statements are issued gets its own thread.

`mysqladmin processlist` only shows the connection, `INSERT DELAYED`, and replication threads.

### 19.1.2. MySQL Test Suite

The test system that is included in Unix source and binary distributions makes it possible for users and developers to perform regression tests on the MySQL code. These tests can be run on Unix.

The current set of test cases doesn't test everything in MySQL, but it should catch most obvious bugs in the SQL processing code, operating system or library issues, and is quite thorough in testing replication. Our goal is to have the tests cover 100% of the code. We welcome contributions to our test suite. You may especially want to contribute tests that examine the functionality critical to your system because this ensures that all future MySQL releases work well with your applications.

The test system consists of a test language interpreter (`mysqltest`), a shell script to run all tests (`mysql-test-run`), the actual test cases written in a special test language, and their expected results. To run the test suite on your system after a build, type `make test` from the source root directory, or change location to the `mysql-test` directory and type `./mysql-test-run`. If you have installed a binary distribution, change location to the `mysql-test` directory under the installation root directory (for example, `/usr/local/mysql/mysql-test`), and run `./mysql-test-run`. All tests should succeed. If any do not, you should try to find out why and report the problem if it indicates a bug in MySQL. See [Section 1.8, “How to Report Bugs or Problems”](#).

If one test fails, you should run `mysql-test-run` with the `--force` option to check whether any other tests fail.

From MySQL 4.1 on, if you have a copy of `mysqld` running on the machine where you want to run the test suite, you do not have to stop it, as long as it is not using ports `9306` or `9307`. If either of those ports is taken, you should edit `mysql-test-run` and change the values of the master or slave port to one that is available.

Before MySQL 4.1, `mysql-test-run` does not try to run its own server by default but tries to use your currently running server. To override this and cause `mysql-test-run` to start its own server, run it with the `--local` option.

In the `mysql-test` directory, you can run an individual test case with `./mysql-test-run test_name`.

You can use the `mysqltest` language to write your own test cases. This is documented in the MySQL Test Framework manual, available at <http://dev.mysql.com/doc/>.

If you have a question about the test suite, or have a test case to contribute, send an email message to the MySQL [internals](#) mailing list. See [Section 1.7.1, “MySQL Mailing Lists”](#). This list does not accept attachments, so you should FTP all the relevant files to: <ftp://ftp.mysql.com/pub/mysql/upload/>

## 19.2. Adding New Functions to MySQL

There are two ways to add new functions to MySQL:

- You can add functions through the user-defined function (UDF) interface. User-defined functions are compiled as object files and then added to and removed from the server dynamically using the `CREATE FUNCTION` and `DROP FUNCTION` statements. See [Section 19.2.2, “CREATE FUNCTION Syntax”](#).
- You can add functions as native (built-in) MySQL functions. Native functions are compiled into the `mysqld` server and become available on a permanent basis.

Each method has advantages and disadvantages:

- If you write user-defined functions, you must install object files in addition to the server itself. If you compile your function into the server, you don't need to do that.
- Native functions require you to modify a source distribution. UDFs do not. You can add UDFs to a

binary MySQL distribution. No access to MySQL source is necessary.

- If you upgrade your MySQL distribution, you can continue to use your previously installed UDFs, unless you upgrade to a newer version for which the UDF interface changes. (An incompatible change occurred in MySQL 4.1.1 for aggregate functions. A function named `xxx_clear()` must be defined rather than `xxx_reset()`.) For native functions, you must repeat your modifications each time you upgrade.

Whichever method you use to add new functions, they can be invoked in SQL statements just like native functions such as `ABS()` or `SOUNDEX()`.

The following sections describe features of the UDF interface, provide instructions for writing UDFs, discuss security precautions that MySQL takes to prevent UDF misuse, and describe how to add native MySQL functions.

For example source code that illustrates how to write UDFs, take a look at the [sql/udf\\_example.cc](#) file that is provided in MySQL source distributions.

## 19.2.1. Features of the User-Defined Function Interface

The MySQL interface for user-defined functions provides the following features and capabilities:

- Functions can return string, integer, or real values.
- You can define simple functions that operate on a single row at a time, or aggregate functions that operate on groups of rows.
- Information is provided to functions that enables them to check the number and types of the arguments passed to them.
- You can tell MySQL to coerce arguments to a given type before passing them to a function.
- You can indicate that a function returns `NULL` or that an error occurred.

## 19.2.2. CREATE FUNCTION Syntax

```
CREATE [AGGREGATE] FUNCTION function_name RETURNS {STRING|INTEGER|REAL}
SONAME shared_library_name
```

A user-defined function (UDF) is a way to extend MySQL with a new function that works like a native (built-in) MySQL function such as `ABS()` or `CONCAT()`.

*function\_name* is the name that should be used in SQL statements to invoke the function. The `RETURNS` clause indicates the type of the function's return value. *shared\_library\_name* is the base-name of the shared object file that contains the code that implements the function. The file must be located in a directory that is searched by your system's dynamic linker.

To create a function, you must have the `INSERT` privilege for the `mysql` database. This is necessary because `CREATE FUNCTION` adds a row to the `mysql.func` system table that records the function's name, type, and shared library name. If you do not have this table, you should run the `mysql_fix_privilege_tables` script to create it. See [Section 5.5.1](#), “`mysql_fix_privilege_tables` — Upgrade MySQL System Tables”.

An active function is one that has been loaded with `CREATE FUNCTION` and not removed with `DROP FUNCTION`. All active functions are reloaded each time the server starts, unless you start `mysqld` with

the `--skip-grant-tables` option. In this case, UDF initialization is skipped and UDFs are unavailable.

For instructions on writing user-defined functions, see [Section 19.2.4, “Adding a New User-Defined Function”](#). For the UDF mechanism to work, functions must be written in C or C++ (or another language that can use C calling conventions), your operating system must support dynamic loading and you must have compiled `mysqld` dynamically (not statically).

`AGGREGATE` is a new option for MySQL 3.23. An `AGGREGATE` function works exactly like a native MySQL aggregate (summary) function such as `SUM` or `COUNT()`. For `AGGREGATE` to work, your `mysql.func` table must contain a `type` column. If your `mysql.func` table does not have this column, you should run the `mysql_fix_privilege_tables` script to create it.

### 19.2.3. DROP FUNCTION Syntax

```
DROP FUNCTION function_name
```

This statement drops the user-defined function (UDF) named *function\_name*.

To drop a function, you must have the `DELETE` privilege for the `mysql` database. This is because `DROP FUNCTION` removes a row from the `mysql.func` system table that records the function's name, type, and shared library name.

### 19.2.4. Adding a New User-Defined Function

For the UDF mechanism to work, functions must be written in C or C++ (or another language that can use C calling conventions), and your operating system must support dynamic loading. The MySQL source distribution includes a file `sql/udf_example.cc` that defines 5 new functions. Consult this file to see how UDF calling conventions work. UDF-related symbols and data structures are defined in the `include/mysql_com.h` header file. (You need not include this header file directly because it is included by `mysql.h`.)

A UDF contains code that becomes part of the running server, so when you write a UDF, you are bound by any and all constraints that otherwise apply to writing server code. For example, you may have problems if you attempt to use functions from the `libstdc++` library. Note that these constraints may change in future versions of the server, so it is possible that server upgrades will require revisions to UDFs that were originally written for older servers. For information about these constraints, see [Section 2.9.2, “Typical configure Options”](#), and [Section 2.9.4, “Dealing with Problems Compiling MySQL”](#).

To be able to use UDFs, you need to link `mysqld` dynamically. Don't configure MySQL using `--with-mysqld-ldflags=-all-static`. If you want to use a UDF that needs to access symbols from `mysqld` (for example, the `metaphone` function in `sql/udf_example.cc` that uses `default_charset_info`), you must link the program with `-rdynamic` (see `man dlopen`). If you plan to use UDFs, the rule of thumb is to configure MySQL with `--with-mysqld-ldflags=-rdynamic` unless you have a very good reason not to.

If you must use a precompiled distribution of MySQL, use MySQL-Max, which contains a dynamically linked server that supports dynamic loading.

For each function that you want to use in SQL statements, you should define corresponding C (or C++) functions. In the following discussion, the name “xxx” is used for an example function name. To distinguish between SQL and C/C++ usage, `XXX()` (uppercase) indicates an SQL function call, and `xxx()` (lowercase) indicates a C/C++ function call.

The C/C++ functions that you write to implement the interface for `XXX()` are:

- `xxx()` (required)

The main function. This is where the function result is computed. The correspondence between the SQL function data type and the return type of your C/C++ function is shown here:

SQL Type	C/C++ Type
STRING	char *
INTEGER	long long
REAL	double

- `xxx_init()` (optional)

The initialization function for `xxx()`. It can be used for the following purposes:

- To check the number of arguments to `xxx()`.
- To check that the arguments are of a required type or, alternatively, to tell MySQL to coerce arguments to the types you want when the main function is called.
- To allocate any memory required by the main function.
- To specify the maximum length of the result.
- To specify (for `REAL` functions) the maximum number of decimal places in the result.
- To specify whether the result can be `NULL`.

- `xxx_deinit()` (optional)

The deinitialization function for `xxx()`. It should deallocate any memory allocated by the initialization function.

When an SQL statement invokes `xxx()`, MySQL calls the initialization function `xxx_init()` to let it perform any required setup, such as argument checking or memory allocation. If `xxx_init()` returns an error, MySQL aborts the SQL statement with an error message and does not call the main or deinitialization functions. Otherwise, MySQL calls the main function `xxx()` once for each row. After all rows have been processed, MySQL calls the deinitialization function `xxx_deinit()` so that it can perform any required cleanup.

For aggregate functions that work like `SUM()`, you must also provide the following functions:

- `xxx_reset()` (required before 4.1.1)

Reset the current aggregate value and insert the argument as the initial aggregate value for a new group.

- `xxx_clear()` (required starting from 4.1.1)

Reset the current aggregate value but do not insert the argument as the initial aggregate value for a new group.

- `xxx_add()` (required)

Add the argument to the current aggregate value.

MySQL handles aggregate UDFs as follows:

1. Call `xxx_init()` to let the aggregate function allocate any memory it needs for storing results.
2. Sort the table according to the `GROUP BY` expression.
3. Call `xxx_clear()` for the first row in each new group.
4. Call `xxx_add()` for each new row that belongs in the same group.
5. Call `xxx()` to get the result for the aggregate when the group changes or after the last row has been processed.
6. Repeat 3-5 until all rows has been processed
7. Call `xxx_deinit()` to let the UDF free any memory it has allocated.

All functions must be thread-safe. This includes not just the main function, but the initialization and deinitialization functions as well, and also the additional functions required by aggregate functions. A consequence of this requirement is that you are not allowed to allocate any global or static variables that change! If you need memory, you should allocate it in `xxx_init()` and free it in `xxx_deinit()`.

### 19.2.4.1. UDF Calling Sequences for Simple Functions

This section describes the different functions that you need to define when you create a simple UDF. [Section 19.2.4, “Adding a New User-Defined Function”](#), describes the order in which MySQL calls these functions.

The main `xxx()` function should be declared as shown in this section. Note that the return type and parameters differ, depending on whether you declare the SQL function `xxx()` to return `STRING`, `INTEGER`, or `REAL` in the `CREATE FUNCTION` statement:

For `STRING` functions:

```
char *xxx(UDF_INIT *initid, UDF_ARGS *args,
         char *result, unsigned long *length,
         char *is_null, char *error);
```

For `INTEGER` functions:

```
long long xxx(UDF_INIT *initid, UDF_ARGS *args,
             char *is_null, char *error);
```

For `REAL` functions:

```
double xxx(UDF_INIT *initid, UDF_ARGS *args,
          char *is_null, char *error);
```

The initialization and deinitialization functions are declared like this:

```
my_bool xxx_init(UDF_INIT *initid, UDF_ARGS *args, char *message);
void xxx_deinit(UDF_INIT *initid);
```

The `initid` parameter is passed to all three functions. It points to a `UDF_INIT` structure that is used to communicate information between functions. The `UDF_INIT` structure members follow. The initialization function should fill in any members that it wishes to change. (To use the default for a member,



leave it unchanged.)

- `my_bool maybe_null`

`xxx_init()` should set `maybe_null` to 1 if `xxx()` can return `NULL`. The default value is 1 if any of the arguments are declared `maybe_null`.

- `unsigned int decimals`

The number of decimal digits to the right of the decimal point. The default value is the maximum number of decimal digits in the arguments passed to the main function. (For example, if the function is passed `1.34`, `1.345`, and `1.3`, the default would be 3, because `1.345` has 3 decimal digits.

- `unsigned int max_length`

The maximum length of the result. The default `max_length` value differs depending on the result type of the function. For string functions, the default is the length of the longest argument. For integer functions, the default is 21 digits. For real functions, the default is 13 plus the number of decimal digits indicated by `initid->decimals`. (For numeric functions, the length includes any sign or decimal point characters.)

If you want to return a blob value, you can set `max_length` to 65KB or 16MB. This memory is not allocated, but the value is used to decide which data type to use if there is a need to temporarily store the data.

- `char *ptr`

A pointer that the function can use for its own purposes. For example, functions can use `initid->ptr` to communicate allocated memory among themselves. `xxx_init()` should allocate the memory and assign it to this pointer:

```
initid->ptr = allocated_memory;
```

In `xxx()` and `xxx_deinit()`, refer to `initid->ptr` to use or deallocate the memory.

### 19.2.4.2. UDF Calling Sequences for Aggregate Functions

This section describes the different functions that you need to define when you create an aggregate UDF. [Section 19.2.4, “Adding a New User-Defined Function”](#), describes the order in which MySQL calls these functions.

- `xxx_reset()`

This function is called when MySQL finds the first row in a new group. It should reset any internal summary variables and then use the given `UDF_ARGS` argument as the first value in your internal summary value for the group. Declare `xxx_reset()` as follows:

```
char *xxx_reset(UDF_INIT *initid, UDF_ARGS *args,
               char *is_null, char *error);
```

`xxx_reset()` is needed only before MySQL 4.1.1. It is *not* needed or used as of MySQL 4.1.1, when the UDF interface changed to use `xxx_clear()` instead. However, you can define both `xxx_reset()` and `xxx_clear()` if you want to have your UDF work both before and after the interface change. (If you do include both functions, the `xxx_reset()` function in many cases can be implemented internally by calling `xxx_clear()` to reset all variables, and then calling `xxx_add()` to add the `UDF_ARGS` argument as the first value in the group.)

- `xxx_clear()`

This function is called when MySQL needs to reset the summary results. It is called at the beginning for each new group but can also be called to reset the values for a query where there were no matching rows. Declare `xxx_clear()` as follows:

```
char *xxx_clear(UDF_INIT *initid, char *is_null, char *error);
```

`is_null` is set to point to `CHAR(0)` before calling `xxx_clear()`.

If something went wrong, you can store a value in the variable to which the `error` argument points. `error` points to a single-byte variable, not to a string buffer.

`xxx_clear()` is required only by MySQL 4.1.1 and above. Before MySQL 4.1.1, use `xxx_reset()` instead.

- `xxx_add()`

This function is called for all rows that belong to the same group, except for the first row. You should use it to add the value in the `UDF_ARGS` argument to your internal summary variable.

```
char *xxx_add(UDF_INIT *initid, UDF_ARGS *args,
             char *is_null, char *error);
```

The `xxx()` function for an aggregate UDF should be declared the same way as for a non-aggregate UDF. See [Section 19.2.4.1, “UDF Calling Sequences for Simple Functions”](#).

For an aggregate UDF, MySQL calls the `xxx()` function after all rows in the group have been processed. You should normally never access its `UDF_ARGS` argument here but instead return a value based on your internal summary variables.

Return value handling in `xxx()` should be done the same way as for a non-aggregate UDF. See [Section 19.2.4.4, “UDF Return Values and Error Handling”](#).

The `xxx_reset()` and `xxx_add()` functions handle their `UDF_ARGS` argument the same way as functions for non-aggregate UDFs. See [Section 19.2.4.3, “UDF Argument Processing”](#).

The pointer arguments to `is_null` and `error` are the same for all calls to `xxx_reset()`, `xxx_clear()`, `xxx_add()` and `xxx()`. You can use this to remember that you got an error or whether the `xxx()` function should return `NULL`. You should not store a string into `*error`! `error` points to a single-byte variable, not to a string buffer.

`*is_null` is reset for each group (before calling `xxx_clear()`). `*error` is never reset.

If `*is_null` or `*error` are set when `xxx()` returns, MySQL returns `NULL` as the result for the group function.

### 19.2.4.3. UDF Argument Processing

The `args` parameter points to a `UDF_ARGS` structure that has the members listed here:

- `unsigned int arg_count`

The number of arguments. Check this value in the initialization function if you require your function to be called with a particular number of arguments. For example:

```
if (args->arg_count != 2)
```

```
{
    strcpy(message, "XXX() requires two arguments");
    return 1;
}
```

- `enum Item_result *arg_type`

A pointer to an array containing the types for each argument. The possible type values are `STRING_RESULT`, `INT_RESULT`, and `REAL_RESULT`.

To make sure that arguments are of a given type and return an error if they are not, check the `arg_type` array in the initialization function. For example:

```
if (args->arg_type[0] != STRING_RESULT ||
    args->arg_type[1] != INT_RESULT)
{
    strcpy(message, "XXX() requires a string and an integer");
    return 1;
}
```

As an alternative to requiring your function's arguments to be of particular types, you can use the initialization function to set the `arg_type` elements to the types you want. This causes MySQL to coerce arguments to those types for each call to `xxx()`. For example, to specify that the first two arguments should be coerced to string and integer, respectively, do this in `xxx_init()`:

```
args->arg_type[0] = STRING_RESULT;
args->arg_type[1] = INT_RESULT;
```

- `char **args`

`args->args` communicates information to the initialization function about the general nature of the arguments passed to your function. For a constant argument `i`, `args->args[i]` points to the argument value. (See below for instructions on how to access the value properly.) For a non-constant argument, `args->args[i]` is 0. A constant argument is an expression that uses only constants, such as `3` or `4*7-2` or `SIN(3.14)`. A non-constant argument is an expression that refers to values that may change from row to row, such as column names or functions that are called with non-constant arguments.

For each invocation of the main function, `args->args` contains the actual arguments that are passed for the row currently being processed.

If argument `i` represents `NULL`, `args->args[i]` is a null pointer (0). If the argument is not `NULL`, functions can refer to it as follows:

- An argument of type `STRING_RESULT` is given as a string pointer plus a length, to allow handling of binary data or data of arbitrary length. The string contents are available as `args->args[i]` and the string length is `args->lengths[i]`. You should not assume that strings are null-terminated.
- For an argument of type `INT_RESULT`, you must cast `args->args[i]` to a `long long` value:

```
long long int_val;
int_val = *((long long*) args->args[i]);
```

- For an argument of type `REAL_RESULT`, you must cast `args->args[i]` to a `double` value:

```
double real_val;
real_val = *((double*) args->args[i]);
```

- `unsigned long *lengths`

For the initialization function, the `lengths` array indicates the maximum string length for each argument. You should not change these. For each invocation of the main function, `lengths` contains the actual lengths of any string arguments that are passed for the row currently being processed. For arguments of types `INT_RESULT` or `REAL_RESULT`, `lengths` still contains the maximum length of the argument (as for the initialization function).

#### 19.2.4.4. UDF Return Values and Error Handling

The initialization function should return `0` if no error occurred and `1` otherwise. If an error occurs, `xxx_init()` should store a null-terminated error message in the `message` parameter. The message is returned to the client. The message buffer is `MYSQL_ERRMSG_SIZE` characters long, but you should try to keep the message to less than 80 characters so that it fits the width of a standard terminal screen.

The return value of the main function `xxx()` is the function value, for `long long` and `double` functions. A string function should return a pointer to the result and set `*result` and `*length` to the contents and length of the return value. For example:

```
memcpy(result, "result string", 13);
*length = 13;
```

The `result` buffer that is passed to the `xxx()` function is 255 bytes long. If your result fits in this, you don't have to worry about memory allocation for results.

If your string function needs to return a string longer than 255 bytes, you must allocate the space for it with `malloc()` in your `xxx_init()` function or your `xxx()` function and free it in your `xxx_deinit()` function. You can store the allocated memory in the `ptr` slot in the `UDF_INIT` structure for reuse by future `xxx()` calls. See Section 19.2.4.1, “UDF Calling Sequences for Simple Functions”.

To indicate a return value of `NULL` in the main function, set `*is_null` to `1`:

```
*is_null = 1;
```

To indicate an error return in the main function, set `*error` to `1`:

```
*error = 1;
```

If `xxx()` sets `*error` to `1` for any row, the function value is `NULL` for the current row and for any subsequent rows processed by the statement in which `xxx()` was invoked. (`xxx()` is not even called for subsequent rows.) **Note:** Before MySQL 3.22.10, you should set both `*error` and `*is_null`:

```
*error = 1;
*is_null = 1;
```

#### 19.2.4.5. Compiling and Installing User-Defined Functions

Files implementing UDFs must be compiled and installed on the host where the server runs. This process is described below for the example UDF file `sql/udf_example.cc` that is included in the MySQL source distribution.

The immediately following instructions are for Unix. Instructions for Windows are given later in this section.

The `udf_example.cc` file contains the following functions:

- `metaphon()` returns a metaphon string of the string argument. This is something like a soundex string, but it's more tuned for English.
- `myfunc_double()` returns the sum of the ASCII values of the characters in its arguments, divided by the sum of the length of its arguments.
- `myfunc_int()` returns the sum of the length of its arguments.
- `sequence([const int])` returns a sequence starting from the given number or 1 if no number has been given.
- `lookup()` returns the IP number for a hostname.
- `reverse_lookup()` returns the hostname for an IP number. The function may be called either with a single string argument of the form '`xxx.xxx.xxx.xxx`' or with four numbers.

A dynamically loadable file should be compiled as a sharable object file, using a command something like this:

```
shell> gcc -shared -o udf_example.so udf_example.cc
```

If you are using `gcc`, you should be able to create `udf_example.so` with a simpler command:

```
shell> make udf_example.so
```

You can easily determine the correct compiler options for your system by running this command in the `sql` directory of your MySQL source tree:

```
shell> make udf_example.o
```

You should run a compile command similar to the one that `make` displays, except that you should remove the `-c` option near the end of the line and add `-o udf_example.so` to the end of the line. (On some systems, you may need to leave the `-c` on the command.)

After you compile a shared object containing UDFs, you must install it and tell MySQL about it. Compiling a shared object from `udf_example.cc` produces a file named something like `udf_example.so` (the exact name may vary from platform to platform). Copy this file to some directory such as `/usr/lib` that searched by your system's dynamic (runtime) linker, or add the directory in which you placed the shared object to the linker configuration file (for example, `/etc/ld.so.conf`).

The dynamic linker name is system-specific (for example, `ld-elf.so.1` on FreeBSD, `ld.so` on Linux, or `dyld` on Mac OS X). Consult your system documentation for information about the linker name and how to configure it.

On many systems, you can also set the `LD_LIBRARY` or `LD_LIBRARY_PATH` environment variable to point at the directory where you have the files for your UDF. The `dlopen` manual page tells you which variable you should use on your system. You should set this in `mysql.server` or `mysqld_safe` startup scripts and restart `mysqld`.

On some systems, the `ldconfig` program that configures the dynamic linker does not recognize a shared object unless its name begins with `lib`. In this case you should rename a file such as `udf_example.so` to `libudf_example.so`.

On Windows, you can compile user-defined functions by using the following procedure:

1. You need to obtain the BitKeeper source repository for MySQL 4.0 or higher. See [Section 2.9.3, “Installing from the Development Source Tree”](#).
2. In the source repository, look in the `VC++Files/examples/udf_example` directory. There are files named `udf_example.def`, `udf_example.dsp`, and `udf_example.dsw` there.
3. In the source repository, look in the `sql` directory. Copy the `udf_example.cc` from this directory to the `VC++Files/examples/udf_example` directory and rename the file to `udf_example.cpp`.
4. Open the `udf_example.dsw` file with Visual Studio VC++ and use it to compile the UDFs as a normal project.

After the shared object file has been installed, notify `mysqld` about the new functions with these statements:

```
mysql> CREATE FUNCTION metaphon RETURNS STRING SONAME 'udf_example.so';
mysql> CREATE FUNCTION myfunc_double RETURNS REAL SONAME 'udf_example.so';
mysql> CREATE FUNCTION myfunc_int RETURNS INTEGER SONAME 'udf_example.so';
mysql> CREATE FUNCTION lookup RETURNS STRING SONAME 'udf_example.so';
mysql> CREATE FUNCTION reverse_lookup
-> RETURNS STRING SONAME 'udf_example.so';
mysql> CREATE AGGREGATE FUNCTION avgcost
-> RETURNS REAL SONAME 'udf_example.so';
```

Functions can be deleted using `DROP FUNCTION`:

```
mysql> DROP FUNCTION metaphon;
mysql> DROP FUNCTION myfunc_double;
mysql> DROP FUNCTION myfunc_int;
mysql> DROP FUNCTION lookup;
mysql> DROP FUNCTION reverse_lookup;
mysql> DROP FUNCTION avgcost;
```

The `CREATE FUNCTION` and `DROP FUNCTION` statements update the `func` system table in the `mysql` database. The function's name, type and shared library name are saved in the table. You must have the `INSERT` and `DELETE` privileges for the `mysql` database to create and drop functions.

You should not use `CREATE FUNCTION` to add a function that has previously been created. If you need to reinstall a function, you should remove it with `DROP FUNCTION` and then reinstall it with `CREATE FUNCTION`. You would need to do this, for example, if you recompile a new version of your function, so that `mysqld` gets the new version. Otherwise, the server continues to use the old version.

An active function is one that has been loaded with `CREATE FUNCTION` and not removed with `DROP FUNCTION`. All active functions are reloaded each time the server starts, unless you start `mysqld` with the `--skip-grant-tables` option. In this case, UDF initialization is skipped and UDFs are unavailable.

### 19.2.4.6. User-Defined Function Security Precautions

MySQL takes the following measures to prevent misuse of user-defined functions.

You must have the `INSERT` privilege to be able to use `CREATE FUNCTION` and the `DELETE` privilege to be able to use `DROP FUNCTION`. This is necessary because these statements add and delete rows from the `mysql.func` table.

UDFs should have at least one symbol defined in addition to the `xxx` symbol that corresponds to the main `xxx()` function. These auxiliary symbols correspond to the `xxx_init()`, `xxx_deinit()`, `xxx_reset()`, `xxx_clear()`, and `xxx_add()` functions. As of MySQL 4.0.24 and 4.1.10a, `mysqld` supports an `--allow-suspicious-udfs` option that controls whether UDFs that have

only an `xxx` symbol can be loaded. By default, the option is off, to prevent attempts at loading functions from shared object files other than those containing legitimate UDFs. If you have older UDFs that contain only the `xxx` symbol and that cannot be recompiled to include an auxiliary symbol, it may be necessary to specify the `--allow-suspicious-udfs` option. Otherwise, you should avoid enabling this capability.

UDF object files cannot be placed in arbitrary directories. They must be located in some system directory that the dynamic linker is configured to search. To enforce this restriction and prevent attempts at specifying pathnames outside of directories searched by the dynamic linker, MySQL checks the shared object file name specified in `CREATE FUNCTION` statements for pathname delimiter characters. As of MySQL 4.0.24 and 4.1.10a, MySQL also checks for pathname delimiters in filenames stored in the `mysql.func` table when it loads functions. This prevents attempts at specifying illegitimate pathnames through direct manipulation of the `mysql.func` table. For information about UDFs and the runtime linker, see [Section 19.2.4.5, “Compiling and Installing User-Defined Functions”](#).

## 19.2.5. Adding a New Native Function

The procedure for adding a new native function is described here. Note that you cannot add native functions to a binary distribution because the procedure involves modifying MySQL source code. You must compile MySQL yourself from a source distribution. Also note that if you migrate to another version of MySQL (for example, when a new version is released), you need to repeat the procedure with the new version.

To add a new native MySQL function, follow these steps:

1. Add one line to `lex.h` that defines the function name in the `sql_functions[]` array.
2. If the function prototype is simple (just takes zero, one, two or three arguments), you should in `lex.h` specify `SYM(FUNC_ARGN)` (where `N` is the number of arguments) as the second argument in the `sql_functions[]` array and add a function that creates a function object in `item_create.cc`. Take a look at `"ABS"` and `create_funcs_abs()` for an example of this.

If the function prototype is complicated (for example, if it takes a variable number of arguments), you should add two lines to `sql_yacc.yy`. One indicates the preprocessor symbol that `yacc` should define (this should be added at the beginning of the file). Then define the function parameters and add an “item” with these parameters to the `simple_expr` parsing rule. For an example, check all occurrences of `ATAN` in `sql_yacc.yy` to see how this is done.

3. In `item_func.h`, declare a class inheriting from `Item_num_func` or `Item_str_func`, depending on whether your function returns a number or a string.
4. In `item_func.cc`, add one of the following declarations, depending on whether you are defining a numeric or string function:

```
double Item_func_newname::val()
longlong Item_func_newname::val_int()
String *Item_func_newname::Str(String *str)
```

If you inherit your object from any of the standard items (like `Item_num_func`), you probably only have to define one of these functions and let the parent object take care of the other functions. For example, the `Item_str_func` class defines a `val()` function that executes `atof()` on the value returned by `::str()`.

5. You should probably also define the following object function:

```
void Item_func_newname::fix_length_and_dec()
```

This function should at least calculate `max_length` based on the given arguments.

`max_length` is the maximum number of characters the function may return. This function should also set `maybe_null = 0` if the main function can't return a `NULL` value. The function can check whether any of the function arguments can return `NULL` by checking the arguments' `maybe_null` variable. You can take a look at `Item_func_mod::fix_length_and_dec` for a typical example of how to do this.

All functions must be thread-safe. In other words, don't use any global or static variables in the functions without protecting them with mutexes)

If you want to return `NULL`, from `::val()`, `::val_int()` or `::str()` you should set `null_value` to 1 and return 0.

For `::str()` object functions, there are some additional considerations to be aware of:

- The `String *str` argument provides a string buffer that may be used to hold the result. (For more information about the `String` type, take a look at the `sql_string.h` file.)
- The `::str()` function should return the string that holds the result or `(char*) 0` if the result is `NULL`.
- All current string functions try to avoid allocating any memory unless absolutely necessary!

## 19.3. Adding New Procedures to MySQL

In MySQL, you can define a procedure in C++ that can access and modify the data in a query before it is sent to the client. The modification can be done on a row-by-row or `GROUP BY` level.

We have created an example procedure in MySQL 3.23 to show you what can be done.

Additionally, we recommend that you take a look at `mysqlua`. With this you can use the LUA language to load a procedure at runtime into `mysqld`.

### 19.3.1. Procedure Analyse

```
analyse([max_elements],[max_memory])
```

This procedure is defined in the `sql/sql_analyse.cc`. This examines the result from your query and returns an analysis of the results:

- `max_elements` (default 256) is the maximum number of distinct values `analyse` does notice per column. This is used by `analyse` to check whether the optimal data type should be of type `ENUM`.
- `max_memory` (default 8192) is the maximum amount of memory that `analyse` should allocate per column while trying to find all distinct values.

```
SELECT ... FROM ... WHERE ... PROCEDURE ANALYSE([max_elements],[max_memory])
```

### 19.3.2. Writing a Procedure

For the moment, the only documentation for this is the source.

You can find all information about procedures by examining the following files:



- `sql/sql_analyse.cc`
- `sql/procedure.h`
- `sql/procedure.cc`
- `sql/sql_select.cc`

---

# Appendix A. Problems and Common Errors

This appendix lists some common problems and error messages that you may encounter. It describes how to determine the causes of the problems and what to do to solve them.

## A.1. How to Determine What Is Causing a Problem

When you run into a problem, the first thing you should do is to find out which program or piece of equipment is causing it:

- If you have one of the following symptoms, then it is probably a hardware problems (such as memory, motherboard, CPU, or hard disk) or kernel problem:
  - The keyboard doesn't work. This can normally be checked by pressing the Caps Lock key. If the Caps Lock light doesn't change, you have to replace your keyboard. (Before doing this, you should try to restart your computer and check all cables to the keyboard.)
  - The mouse pointer doesn't move.
  - The machine doesn't answer to a remote machine's pings.
  - Other programs that are not related to MySQL don't behave correctly.
  - Your system restarted unexpectedly. (A faulty user-level program should never be able to take down your system.)

In this case, you should start by checking all your cables and run some diagnostic tool to check your hardware! You should also check whether there are any patches, updates, or service packs for your operating system that could likely solve your problem. Check also that all your libraries (such as [glibc](#)) are up to date.

It's always good to use a machine with ECC memory to discover memory problems early.

- If your keyboard is locked up, you may be able to recover by logging in to your machine from another machine and executing `kbd_mode -a`.
- Please examine your system log file (`/var/log/messages` or similar) for reasons for your problem. If you think the problem is in MySQL, you should also examine MySQL's log files. See [Section 5.11, "MySQL Server Logs"](#).
- If you don't think you have hardware problems, you should try to find out which program is causing problems. Try using `top`, `ps`, Task Manager, or some similar program, to check which program is taking all CPU or is locking the machine.
- Use `top`, `df`, or a similar program to check whether you are out of memory, disk space, file descriptors, or some other critical resource.
- If the problem is some runaway process, you can always try to kill it. If it doesn't want to die, there is probably a bug in the operating system.

If after you have examined all other possibilities and you have concluded that the MySQL server or a MySQL client is causing the problem, it's time to create a bug report for our mailing list or our support team. In the bug report, try to give a very detailed description of how the system is behaving and what you think is happening. You should also state why you think that MySQL is causing the problem. Take into consideration all the situations in this chapter. State any problems exactly how they appear when

you examine your system. Use the “copy and paste” method for any output and error messages from programs and log files.

Try to describe in detail which program is not working and all symptoms you see. We have in the past received many bug reports that state only “the system doesn't work.” This doesn't provide us with any information about what could be the problem.

If a program fails, it's always useful to know the following information:

- Has the program in question made a segmentation fault (did it dump core)?
- Is the program taking up all available CPU time? Check with `top`. Let the program run for a while, it may simply be evaluating something computationally intensive.
- If the `mysqld` server is causing problems, can you get any response from it with `mysqladmin -u root ping` or `mysqladmin -u root processlist`?
- What does a client program say when you try to connect to the MySQL server? (Try with `mysql`, for example.) Does the client jam? Do you get any output from the program?

When sending a bug report, you should follow the outline described in [Section 1.8, “How to Report Bugs or Problems”](#).

## A.2. Common Errors When Using MySQL Programs

This section lists some errors that users frequently encounter when running MySQL programs. Although the problems show up when you try to run client programs, the solutions to many of the problems involves changing the configuration of the MySQL server.

### A.2.1. Access denied

An `Access denied` error can have many causes. Often the problem is related to the MySQL accounts that the server allows client programs to use when connecting. See [Section 5.7.8, “Causes of Access denied Errors”](#), and [Section 5.7.2, “How the Privilege System Works”](#).

### A.2.2. Can't connect to [local] MySQL server

A MySQL client on Unix can connect to the `mysqld` server in two different ways: By using a Unix socket file to connect through a file in the filesystem (default `/tmp/mysql.sock`), or by using TCP/IP, which connects through a port number. A Unix socket file connection is faster than TCP/IP, but can be used only when connecting to a server on the same computer. A Unix socket file is used if you don't specify a hostname or if you specify the special hostname `localhost`.

If the MySQL server is running on Windows 9x or Me, you can connect only via TCP/IP. If the server is running on Windows NT, 2000, XP, or 2003 and is started with the `--enable-named-pipe` option, you can also connect with named pipes if you run the client on the host where the server is running. The name of the named pipe is `MySQL` by default. If you don't give a hostname when connecting to `mysqld`, a MySQL client first tries to connect to the named pipe. If that doesn't work, it connects to the TCP/IP port. You can force the use of named pipes on Windows by using `.` as the hostname.

The error (2002) `Can't connect to ...` normally means that there is no MySQL server running on the system or that you are using an incorrect Unix socket filename or TCP/IP port number when trying to connect to the server.

The error (2003) `Can't connect to MySQL server on 'server' (10061)` indicates that

the network connection has been refused. You should check that there is a MySQL server running, that it has network connections enabled, the network port you specified is the one configured on the server, and that the TCP/IP port you are using has not been blocked by a firewall or port blocking service.

Start by checking whether there is a process named `mysqld` running on your server host. (Use `ps xa | grep mysqld` on Unix or the Task Manager on Windows.) If there is no such process, you should start the server. See [Section 2.10.2.3](#), “Starting and Troubleshooting the MySQL Server”.

If a `mysqld` process is running, you can check it by trying the following commands. The port number or Unix socket filename might be different in your setup. `host_ip` represents the IP number of the machine where the server is running.

```
shell> mysqladmin version
shell> mysqladmin variables
shell> mysqladmin -h `hostname` version variables
shell> mysqladmin -h `hostname` --port=3306 version
shell> mysqladmin -h host_ip version
shell> mysqladmin --protocol=socket --socket=/tmp/mysql.sock version
```

Note the use of backticks rather than forward quotes with the `hostname` command; these cause the output of `hostname` (that is, the current hostname) to be substituted into the `mysqladmin` command. If you have no `hostname` command or are running on Windows, you can manually type the hostname of your machine (without backticks) following the `-h` option. You can also try `-h 127.0.0.1` to connect with TCP/IP to the local host.

Here are some reasons the `Can't connect to local MySQL server` error might occur:

- `mysqld` is not running. Check your operating system's process list to ensure the `mysqld` process is present.
- You're running a MySQL server on Windows with many TCP/IP connections to it. If you're experiencing that quite often your clients get that error, you can find a workaround here: [Section A.2.2.1](#), “Connection to MySQL Server Failing on Windows”.
- You are running on a system that uses MIT-pthreads. If you are running on a system that doesn't have native threads, `mysqld` uses the MIT-pthreads package. See [Section 2.1.1](#), “Operating Systems Supported by MySQL”. However, not all MIT-pthreads versions support Unix socket files. On a system without socket file support, you must always specify the hostname explicitly when connecting to the server. Try using this command to check the connection to the server:

```
shell> mysqladmin -h `hostname` version
```

- Someone has removed the Unix socket file that `mysqld` uses (`/tmp/mysql.sock` by default). For example, you might have a `cron` job that removes old files from the `/tmp` directory. You can always run `mysqladmin version` to check whether the Unix socket file that `mysqladmin` is trying to use really exists. The fix in this case is to change the `cron` job to not remove `mysql.sock` or to place the socket file somewhere else. See [Section A.4.5](#), “How to Protect or Change the MySQL Unix Socket File”.
- You have started the `mysqld` server with the `--socket=/path/to/socket` option, but forgotten to tell client programs the new name of the socket file. If you change the socket pathname for the server, you must also notify the MySQL clients. You can do this by providing the same `--socket` option when you run client programs. You also need to ensure that clients have permission to access the `mysql.sock` file. To find out where the socket file is, you can do:

```
shell> netstat -ln | grep mysql
```

See [Section A.4.5](#), “How to Protect or Change the MySQL Unix Socket File”.

- You are using Linux and one server thread has died (dumped core). In this case, you must kill the other `mysqld` threads (for example, with `kill` or with the `mysql_zap` script) before you can restart the MySQL server. See [Section A.4.2, “What to Do If MySQL Keeps Crashing”](#).
- The server or client program might not have the proper access privileges for the directory that holds the Unix socket file or the socket file itself. In this case, you must either change the access privileges for the directory or socket file so that the server and clients can access them, or restart `mysqld` with a `--socket` option that specifies a socket filename in a directory where the server can create it and where client programs can access it.

If you get the error message `Can't connect to MySQL server on some_host`, you can try the following things to find out what the problem is:

- Check whether the server is running on that host by executing `telnet some_host 3306` and pressing the Enter key a couple of times. (3306 is the default MySQL port number. Change the value if your server is listening to a different port.) If there is a MySQL server running and listening to the port, you should get a response that includes the server's version number. If you get an error such as `telnet: Unable to connect to remote host: Connection refused`, then there is no server running on the given port.
- If the server is running on the local host, try using `mysqladmin -h localhost variables` to connect using the Unix socket file. Verify the TCP/IP port number that the server is configured to listen to (it is the value of the `port` variable.)
- Make sure that your `mysqld` server was not started with the `--skip-networking` option. If it was, you cannot connect to it using TCP/IP.
- Check to make sure that there is no firewall blocking access to MySQL. Applications such as ZoneAlarm and the Windows XP personal firewall may need to be configured to allow external access to a MySQL server.

### A.2.2.1. Connection to MySQL Server Failing on Windows

When you're running a MySQL server on Windows with many TCP/IP connections to it, and you're experiencing that quite often your clients get a `Can't connect to MySQL server` error, the reason might be that Windows doesn't allow for enough ephemeral (short-lived) ports to serve those connections.

By default, Windows allows 5000 ephemeral (short-lived) TCP ports to the user. After any port is closed it will remain in a `TIME_WAIT` status for 120 seconds. This status allows the connection to be reused at a much lower cost than reinitializing a brand new connection. However, the port will not be available again until this time expires.

With a small stack of available TCP ports (5000) and a high number of TCP ports being open and closed over a short period of time along with the `TIME_WAIT` status you have a good chance for running out of ports. There are two ways to address this problem:

- Reduce the number of TCP ports consumed quickly by investigating connection pooling or persistent connections where possible
- Tune some settings in the Windows registry (see below)

**IMPORTANT: The following procedure involves modifying the Windows registry. Before you modify the registry, make sure to back it up and make sure that you understand how to restore**

the registry if a problem occurs. For information about how to back up, restore, and edit the registry, view the following article in the Microsoft Knowledge Base: <http://support.microsoft.com/kb/256986/EN-US/>.

1. Start Registry Editor ([Regedt32.exe](#)).
2. Locate the following key in the registry:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters
```

3. On the **Edit** menu, click **Add Value**, and then add the following registry value:

```
Value Name: MaxUserPort  
Data Type: REG_DWORD  
Value: 65534
```

This sets the number of ephemeral ports available to any user. The valid range is between 5000 and 65534 (decimal). The default value is 0x1388 (5000 decimal).

4. On the **Edit** menu, click **Add Value**, and then add the following registry value:

```
Value Name: TcpTimedWaitDelay  
Data Type: REG_DWORD  
Value: 30
```

This sets the number of seconds to hold a TCP port connection in **TIME\_WAIT** state before closing. The valid range is between 0 (zero) and 300 (decimal). The default value is 0x78 (120 decimal).

5. Quit Registry Editor.
6. Reboot the machine.

Note: Undoing the above should be as simple as deleting the registry entries you've created.

### A.2.3. Client does not support authentication protocol

MySQL 4.1 and up uses an authentication protocol based on a password hashing algorithm that is incompatible with that used by older clients. If you upgrade the server from 4.0, attempts to connect to it with an older client may fail with the following message:

```
shell> mysql  
Client does not support authentication protocol requested  
by server; consider upgrading MySQL client
```

To solve this problem, you should use one of the following approaches:

- Upgrade all client programs to use a 4.1.1 or newer client library.
- When connecting to the server with a pre-4.1 client program, use an account that still has a pre-4.1-style password.
- Reset the password to pre-4.1 style for each user that needs to use a pre-4.1 client program. This can be done using the **SET PASSWORD** statement and the **OLD\_PASSWORD()** function:

```
mysql> SET PASSWORD FOR
-> 'some_user'@'some_host' = OLD_PASSWORD('newpwd');
```

Alternatively, use `UPDATE` and `FLUSH PRIVILEGES`:

```
mysql> UPDATE mysql.user SET Password = OLD_PASSWORD('newpwd')
-> WHERE Host = 'some_host' AND User = 'some_user';
mysql> FLUSH PRIVILEGES;
```

Substitute the password you want to use for “*newpwd*” in the preceding examples. MySQL cannot tell you what the original password was, so you'll need to pick a new one.

- Tell the server to use the older password hashing algorithm:
  1. Start `mysqld` with the `--old-passwords` option.
  2. Assign an old-format password to each account that has had its password updated to the longer 4.1 format. You can identify these accounts with the following query:

```
mysql> SELECT Host, User, Password FROM mysql.user
-> WHERE LENGTH>Password) > 16;
```

For each account record displayed by the query, use the `Host` and `User` values and assign a password using the `OLD_PASSWORD()` function and either `SET PASSWORD` or `UPDATE`, as described earlier.

**Note:** In PHP, the standard `mysql` extension does not support the new authentication protocol in MySQL 4.1.1 and higher. This is true regardless of the PHP version being used. If you wish to use the `mysql` extension with MySQL 4.1 or newer, you will need to follow one of the options discussed above for configuring MySQL to work with old clients. The `mysqli` extension (stands for “MySQL, Improved”; new in PHP 5) is compatible with the improved password hashing employed in MySQL 4.1 and higher, and no special configuration of MySQL need be done to use this newer MySQL client library for PHP. For more information about the `mysqli` extension, see <http://php.net/mysqli>.

It may also be possible to compile the older `mysql` extension against the new MySQL client library. This is beyond the scope of this Manual; consult the PHP documentation for more information. You also be able to obtain assistance with these issues in our [MySQL with PHP forum](http://forums.mysql.com/list.php?52) [http://forums.mysql.com/list.php?52].

For additional background on password hashing and authentication, see [Section 5.7.9, “Password Hashing as of MySQL 4.1”](#).

## A.2.4. Password Fails When Entered Interactively

MySQL client programs prompt for a password when invoked with a `--password` or `-p` option that has no following password value:

```
shell> mysql -u user_name -p
Enter password:
```

On some systems, you may find that your password works when specified in an option file or on the command line, but not when you enter it interactively at the `Enter password:` prompt. This occurs when the library provided by the system to read passwords limits password values to a small number of characters (typically eight). That is a problem with the system library, not with MySQL. To work around it, change your MySQL password to a value that is eight or fewer characters long, or put your password in an option file.

## A.2.5. Host '*host\_name*' is blocked

If you get the following error, it means that `mysqld` has received many connect requests from the host '*host\_name*' that have been interrupted in the middle:

```
Host 'host_name' is blocked because of many connection errors.  
Unblock with 'mysqladmin flush-hosts'
```

The number of interrupted connect requests allowed is determined by the value of the `max_connect_errors` system variable. After `max_connect_errors` failed requests, `mysqld` assumes that something is wrong (for example, that someone is trying to break in), and blocks the host from further connections until you execute a `mysqladmin flush-hosts` command or issue a `FLUSH HOSTS` statement. See [Section 5.2.2, “System Variables”](#).

By default, `mysqld` blocks a host after 10 connection errors. You can adjust the value by starting the server like this:

```
shell> mysql_safe --max_connect_errors=10000 &
```

If you get this error message for a given host, you should first verify that there isn't anything wrong with TCP/IP connections from that host. If you are having network problems, it does you no good to increase the value of the `max_connect_errors` variable.

## A.2.6. Too many connections

If you get a `Too many connections` error when you try to connect to the `mysqld` server, this means that all available connections are in use by other clients.

The number of connections allowed is controlled by the `max_connections` system variable. Its default value is 100. If you need to support more connections, you should restart `mysqld` with a larger value for this variable.

`mysqld` actually allows `max_connections+1` clients to connect. The extra connection is reserved for use by accounts that have the `SUPER` privilege. By granting the `SUPER` privilege to administrators and not to normal users (who should not need it), an administrator can connect to the server and use `SHOW PROCESSLIST` to diagnose problems even if the maximum number of unprivileged clients are connected. See [Section 13.5.4.16, “SHOW PROCESSLIST Syntax”](#).

The maximum number of connections MySQL can support depends on the quality of the thread library on a given platform. Linux or Solaris should be able to support 500-1000 simultaneous connections, depending on how much RAM you have and what your clients are doing. Static Linux binaries provided by MySQL AB can support up to 4000 connections.

## A.2.7. Out of memory

If you issue a query using the `mysql` client program and receive an error like the following one, it means that `mysql` does not have enough memory to store the entire query result:

```
mysql: Out of memory at line 42, 'malloc.c'  
mysql: needed 8136 byte (8k), memory in use: 12481367 bytes (12189k)  
ERROR 2008: MySQL client ran out of memory
```

To remedy the problem, first check whether your query is correct. Is it reasonable that it should return so many rows? If not, correct the query and try again. Otherwise, you can invoke `mysql` with the `-quick` option. This causes it to use the `mysql_use_result()` C API function to retrieve the result set, which places less of a load on the client (but more on the server).



## A.2.8. MySQL server has gone away

This section also covers the related `Lost connection to server during query` error.

The most common reason for the `MySQL server has gone away` error is that the server timed out and closed the connection. In this case, you normally get one of the following error codes (which one you get is operating system-dependent):

Error Code	Description
<code>CR_SERVER_GONE_ERROR</code>	The client couldn't send a question to the server.
<code>CR_SERVER_LOST</code>	The client didn't get an error when writing to the server, but it didn't get a full answer (or any answer) to the question.

By default, the server closes the connection after eight hours if nothing has happened. You can change the time limit by setting the `wait_timeout` variable when you start `mysqld`. See [Section 5.2.2, “System Variables”](#).

If you have a script, you just have to issue the query again for the client to do an automatic reconnection. This assumes that you have automatic reconnection in the client enabled (which is the default for the `mysql` command-line client).

Some other common reasons for the `MySQL server has gone away` error are:

- You (or the db administrator) has killed the running thread with a `KILL` statement or a `mysqladmin kill` command.
- You tried to run a query after closing the connection to the server. This indicates a logic error in the application that should be corrected.
- A client application running on a different host does not have the necessary privileges to connect to the MySQL server from that host.
- You got a timeout from the TCP/IP connection on the client side. This may happen if you have been using the commands: `mysql_options(..., MYSQL_OPT_READ_TIMEOUT,...)` or `mysql_options(..., MYSQL_OPT_WRITE_TIMEOUT,...)`. In this case increasing the timeout may help solve the problem.
- You have encountered a timeout on the server side and the automatic reconnection in the client is disabled (the `reconnect` flag in the `MYSQL` structure is equal to 0).
- You are using a Windows client and the server had dropped the connection (probably because `wait_timeout` expired) before the command was issued.

The problem on Windows is that in some cases MySQL doesn't get an error from the OS when writing to the TCP/IP connection to the server, but instead gets the error when trying to read the answer from the connection.

In this case, even if the `reconnect` flag in the `MYSQL` structure is equal to 1, MySQL does not automatically reconnect and re-issue the query as it doesn't know if the server did get the original query or not.

The solution to this is to either do a `mysql_ping` on the connection if there has been a long time since the last query (this is what `MyODBC` does) or set `wait_timeout` on the `mysqld` server so high that it in practice never times out.

- You can also get these errors if you send a query to the server that is incorrect or too large. If

`mysqld` receives a packet that is too large or out of order, it assumes that something has gone wrong with the client and closes the connection. If you need big queries (for example, if you are working with big `BLOB` columns), you can increase the query limit by setting the server's `max_allowed_packet` variable, which has a default value of 1MB. You may also need to increase the maximum packet size on the client end. More information on setting the packet size is given in [Section A.2.9, “Packet too large”](#).

An `INSERT` or `REPLACE` statement that inserts a great many rows can also cause these sorts of errors. Either one of these statements sends a single request to the server irrespective of the number of rows to be inserted; thus, you can often avoid the error by reducing the number of rows sent per `INSERT` or `REPLACE`.

- You also get a lost connection if you are sending a packet 16MB or larger if your client is older than 4.0.8 and your server is 4.0.8 and above, or the other way around.
- It is also possible to see this error if hostname lookups fail (for example, if the DNS server on which your server or network relies goes down). This is because MySQL is dependent on the host system for name resolution, but has no way of knowing whether it is working — from MySQL's point of view the problem is indistinguishable from any other network timeout.

You may also see the `MySQL server has gone away` error if MySQL is started with the `--skip-networking` option.

Another networking issue that can cause this error occurs if the MySQL port (default 3306) is blocked by your firewall, thus preventing any connections at all to the MySQL server.

- You can also encounter this error with applications that fork child processes, all of which try to use the same connection to the MySQL server. This can be avoided by using a separate connection for each child process.
- You have encountered a bug where the server died while executing the query.

You can check whether the MySQL server died and restarted by executing `mysqladmin version` and examining the server's uptime. If the client connection was broken because `mysqld` crashed and restarted, you should concentrate on finding the reason for the crash. Start by checking whether issuing the query again kills the server again. See [Section A.4.2, “What to Do If MySQL Keeps Crashing”](#).

You can get more information about the lost connections by starting `mysqld` with the `--log-warnings=2` option. This logs some of the disconnected errors in the `hostname.err` file. See [Section 5.11.1, “The Error Log”](#).

If you want to create a bug report regarding this problem, be sure that you include the following information:

- Indicate whether the MySQL server died. You can find information about this in the server error log. See [Section A.4.2, “What to Do If MySQL Keeps Crashing”](#).
- If a specific query kills `mysqld` and the tables involved were checked with `CHECK TABLE` before you ran the query, can you provide a reproducible test case? See [Section E.1.6, “Making a Test Case If You Experience Table Corruption”](#).
- What is the value of the `wait_timeout` system variable in the MySQL server? (`mysqladmin variables` gives you the value of this variable.)
- Have you tried to run `mysqld` with the `--log` option to determine whether the problem query appears in the log?

See also [Section A.2.10, “Communication Errors and Aborted Connections”](#), and [Section 1.8, “How to Report Bugs or Problems”](#).

## A.2.9. Packet too large

A communication packet is a single SQL statement sent to the MySQL server or a single row that is sent to the client.

In MySQL 3.23, the largest possible packet is 16MB, due to limits in the client/server protocol. In MySQL 4.0.1 and up, the limit is 1GB.

When a MySQL client or the `mysqld` server receives a packet bigger than `max_allowed_packet` bytes, it issues a `Packet too large` error and closes the connection. With some clients, you may also get a `Lost connection to MySQL server during query` error if the communication packet is too large.

Both the client and the server have their own `max_allowed_packet` variable, so if you want to handle big packets, you must increase this variable both in the client and in the server.

If you are using the `mysql` client program, its default `max_allowed_packet` variable is 16MB. That is also the maximum value before MySQL 4.0. To set a larger value from 4.0 on, start `mysql` like this:

```
shell> mysql --max_allowed_packet=32M
```

That sets the packet size to 32MB.

The server's default `max_allowed_packet` value is 1MB. You can increase this if the server needs to handle big queries (for example, if you are working with big `BLOB` columns). For example, to set the variable to 16MB, start the server like this:

```
shell> mysqld --max_allowed_packet=16M
```

Before MySQL 4.0, use this syntax instead:

```
shell> mysqld --set-variable=max_allowed_packet=16M
```

You can also use an option file to set `max_allowed_packet`. For example, to set the size for the server to 16MB, add the following lines in an option file:

```
[mysqld]
max_allowed_packet=16M
```

Before MySQL 4.0, use this syntax instead:

```
[mysqld]
set-variable = max_allowed_packet=16M
```

It's safe to increase the value of this variable because the extra memory is allocated only when needed. For example, `mysqld` allocates more memory only when you issue a long query or when `mysqld` must return a large result row. The small default value of the variable is a precaution to catch incorrect packets between the client and server and also to ensure that you don't run out of memory by using large packets accidentally.

You can also get strange problems with large packets if you are using large `BLOB` values but have not given `mysqld` access to enough memory to handle the query. If you suspect this is the case, try adding `ulimit -d 256000` to the beginning of the `mysqld_safe` script and restarting `mysqld`.

## A.2.10. Communication Errors and Aborted Connections

The server error log can be a useful source of information about connection problems. See [Section 5.11.1, “The Error Log”](#). Starting with MySQL 3.23.40, if you start the server with the `--warnings` option (or `--log-warnings` from MySQL 4.0.3 on), you might find messages like this in your error log:

```
010301 14:38:23 Aborted connection 854 to db: 'users' user: 'josh'
```

If `Aborted connections` messages appear in the error log, the cause can be any of the following:

- The client program did not call `mysql_close()` before exiting.
- The client had been sleeping more than `wait_timeout` or `interactive_timeout` seconds without issuing any requests to the server. See [Section 5.2.2, “System Variables”](#).
- The client program ended abruptly in the middle of a data transfer.

When any of these things happen, the server increments the `Aborted_clients` status variable.

The server increments the `Aborted_connects` status variable when the following things happen:

- A client doesn't have privileges to connect to a database.
- A client uses an incorrect password.
- A connection packet doesn't contain the right information.
- It takes more than `connect_timeout` seconds to get a connect packet. See [Section 5.2.2, “System Variables”](#).

If these kinds of things happen, it might indicate that someone is trying to break into your server!

Other reasons for problems with aborted clients or aborted connections:

- Use of Ethernet protocol with Linux, both half and full duplex. Many Linux Ethernet drivers have this bug. You should test for this bug by transferring a huge file via FTP between the client and server machines. If a transfer goes in burst-pause-burst-pause mode, you are experiencing a Linux duplex syndrome. The only solution is switching the duplex mode for both your network card and hub/switch to either full duplex or to half duplex and testing the results to determine the best setting.
- Some problem with the thread library that causes interrupts on reads.
- Badly configured TCP/IP.
- Faulty Ethernets, hubs, switches, cables, and so forth. This can be diagnosed properly only by replacing hardware.
- The `max_allowed_packet` variable value is too small or queries require more memory than you have allocated for `mysqld`. See [Section A.2.9, “Packet too large”](#).

See also [Section A.2.8, “MySQL server has gone away”](#).

## A.2.11. The table is full

There are several ways a full-table error can occur:

- You are using a MySQL server older than 3.23 and an in-memory temporary table becomes larger than `tmp_table_size` bytes. To avoid this problem, you can use the `--tmp_table_size=val` option to make `mysqld` increase the temporary table size or use the SQL option `SQL_BIG_TABLES` before you issue the problematic query. See [Section 13.5.3, “SET Syntax”](#).

You can also start `mysqld` with the `--big-tables` option. This is exactly the same as using `SQL_BIG_TABLES` for all queries.

As of MySQL 3.23, this problem should not occur. If an in-memory temporary table becomes larger than `tmp_table_size`, the server automatically converts it to a disk-based `MyISAM` table.

- You are using `InnoDB` tables and run out of room in the `InnoDB` tablespace. In this case, the solution is to extend the `InnoDB` tablespace. See [Section 14.2.8, “Adding and Removing InnoDB Data and Log Files”](#).
- You are using `ISAM` or `MyISAM` tables on an operating system that supports files only up to 2GB in size and you have hit this limit for the data file or index file.
- You are using a `MyISAM` table and the space required for the table exceeds what is allowed by the internal pointer size. If you don't specify the `MAX_ROWS` table option when you create a table, MySQL uses the `myisam_data_pointer_size` system variable. The default value is 4 bytes, which is enough to allow only 4GB of data. See [Section 5.2.2, “System Variables”](#).

You can check the maximum data/index sizes by using this statement:

```
SHOW TABLE STATUS FROM database LIKE 'tbl_name';
```

You also can use `myisamchk -dv /path/to/table-index-file`.

If the pointer size is too small, you can fix the problem by using `ALTER TABLE`:

```
ALTER TABLE tbl_name MAX_ROWS=1000000000 AVG_ROW_LENGTH=nnn;
```

You have to specify `AVG_ROW_LENGTH` only for tables with `BLOB` or `TEXT` columns; in this case, MySQL can't optimize the space required based only on the number of rows.

## A.2.12. Can't create/write to file

If you get an error of the following type for some queries, it means that MySQL cannot create a temporary file for the result set in the temporary directory:

```
Can't create/write to file '\\sqla3fe_0.ism'.
```

The preceding error is a typical message for Windows; the Unix message is similar.

One fix is to start `mysqld` with the `--tmpdir` option or to add the option to the `[mysqld]` section of your option file. For example, to specify a directory of `C:\temp`, use these lines:

```
[mysqld]
tmpdir=C:/temp
```

The `C:\temp` directory must exist and have sufficient space for the MySQL server to write to. See [Section 4.3.2, “Using Option Files”](#).

Another cause of this error can be permissions issues. Make sure that the MySQL server can write to the `tmpdir` directory.

Check also the error code that you get with `pererror`. One reason the server cannot write to a table is that the filesystem is full:

```
shell> pererror 28
Error code 28: No space left on device
```

## A.2.13. Commands out of sync

If you get `Commands out of sync; you can't run this command now` in your client code, you are calling client functions in the wrong order.

This can happen, for example, if you are using `mysql_use_result()` and try to execute a new query before you have called `mysql_free_result()`. It can also happen if you try to execute two queries that return data without calling `mysql_use_result()` or `mysql_store_result()` in between.

## A.2.14. Ignoring user

If you get the following error, it means that when `mysqld` was started or when it reloaded the grant tables, it found an account in the `user` table that had an invalid password.

```
Found wrong password for user 'some_user'@'some_host'; ignoring user
```

As a result, the account is simply ignored by the permission system.

The following list indicates possible causes of and fixes for this problem:

- You may be running a new version of `mysqld` with an old `user` table. You can check this by executing `mysqlshow mysql user` to see whether the `Password` column is shorter than 16 characters. If so, you can correct this condition by running the `scripts/add_long_password` script.
- The account has an old password (eight characters long) and you didn't start `mysqld` with the `--old-protocol` option. Update the account in the `user` table to have a new password or restart `mysqld` with the `--old-protocol` option.
- You have specified a password in the `user` table without using the `PASSWORD()` function. Use `mysql` to update the account in the `user` table with a new password, making sure to use the `PASSWORD()` function:

```
mysql> UPDATE user SET Password=PASSWORD('newpwd')
-> WHERE User='some_user' AND Host='some_host';
```

## A.2.15. Table '`tbl_name`' doesn't exist

If you get either of the following errors, it usually means that no table exists in the default database with the given name:

```
Table 'tbl_name' doesn't exist
Can't find file: 'tbl_name' (errno: 2)
```

In some cases, it may be that the table does exist but that you are referring to it incorrectly:

- Because MySQL uses directories and files to store databases and tables, database and table names are case sensitive if they are located on a filesystem that has case-sensitive filenames.
- Even for filesystems that are not case sensitive, such as on Windows, all references to a given table within a query must use the same lettercase.

You can check which tables are in the default database with `SHOW TABLES`. See [Section 13.5.4](#), “`SHOW Syntax`”.

## A.2.16. Can't initialize character set

You might see an error like this if you have character set problems:

```
MySQL Connection Failed: Can't initialize character set charset_name
```

This error can have any of the following causes:

- The character set is a multi-byte character set and you have no support for the character set in the client. In this case, you need to recompile the client by running `configure` with the `-with-charset=charset_name` or `--with-extra-charsets=charset_name` option. See [Section 2.9.2](#), “`Typical configure Options`”.

All standard MySQL binaries are compiled with `-with-extra-character-sets=complex`, which enables support for all multi-byte character sets. See [Section 5.10.1](#), “`The Character Set Used for Data and Sorting`”.

- The character set is a simple character set that is not compiled into `mysqld`, and the character set definition files are not in the place where the client expects to find them.

In this case, you need to use one of the following methods to solve the problem:

- Recompile the client with support for the character set. See [Section 2.9.2](#), “`Typical configure Options`”.
- Specify to the client the directory where the character set definition files are located. For many clients, you can do this with the `--character-sets-dir` option.
- Copy the character definition files to the path where the client expects them to be.

## A.2.17. File Not Found

If you get `ERROR '...' not found (errno: 23), Can't open file: ... (errno: 24)`, or any other error with `errno 23` or `errno 24` from MySQL, it means that you haven't allocated enough file descriptors for the MySQL server. You can use the `perror` utility to get a description of what the error number means:

```
shell> perror 23
Error code 23: File table overflow
shell> perror 24
Error code 24: Too many open files
shell> perror 11
Error code 11: Resource temporarily unavailable
```

The problem here is that `mysqld` is trying to keep open too many files simultaneously. You can either tell `mysqld` not to open so many files at once or increase the number of file descriptors available to `mysqld`.

To tell `mysqld` to keep open fewer files at a time, you can make the table cache smaller by reducing the value of the `table_cache` system variable (the default value is 64). Reducing the value of `max_connections` also reduces the number of open files (the default value is 100).

To change the number of file descriptors available to `mysqld`, you can use the `--open-files-limit` option to `mysqld_safe` or (as of MySQL 3.23.30) set the `open_files_limit` system variable. See [Section 5.2.2, “System Variables”](#). The easiest way to set these values is to add an option to your option file. See [Section 4.3.2, “Using Option Files”](#). If you have an old version of `mysqld` that doesn't support setting the open files limit, you can edit the `mysqld_safe` script. There is a commented-out line `ulimit -n 256` in the script. You can remove the `#` character to uncomment this line, and change the number `256` to set the number of file descriptors to be made available to `mysqld`.

`--open-files-limit` and `ulimit` can increase the number of file descriptors, but only up to the limit imposed by the operating system. There is also a “hard” limit that can be overridden only if you start `mysqld_safe` or `mysqld` as `root` (just remember that you also need to start the server with the `--user` option in this case so that it does not continue to run as `root` after it starts up). If you need to increase the operating system limit on the number of file descriptors available to each process, consult the documentation for your system.

**Note:** If you run the `tcsh` shell, `ulimit` does not work! `tcsh` also reports incorrect values when you ask for the current limits. In this case, you should start `mysqld_safe` using `sh`.

## A.3. Installation-Related Issues

### A.3.1. Problems Linking to the MySQL Client Library

When you are linking an application program to use the MySQL client library, you might get undefined reference errors for symbols that start with `mysql_`, such as those shown here:

```
/tmp/ccFKsdPa.o: In function `main':  
/tmp/ccFKsdPa.o(.text+0xb): undefined reference to `mysql_init'  
/tmp/ccFKsdPa.o(.text+0x31): undefined reference to `mysql_real_connect'  
/tmp/ccFKsdPa.o(.text+0x57): undefined reference to `mysql_real_connect'  
/tmp/ccFKsdPa.o(.text+0x69): undefined reference to `mysql_error'  
/tmp/ccFKsdPa.o(.text+0x9a): undefined reference to `mysql_close'
```

You should be able to solve this problem by adding `-Ldir_path -lmysqlclient` at the end of your link command, where `dir_path` represents the pathname of the directory where the client library is located. To determine the correct directory, try this command:

```
shell> mysql_config --libs
```

The output from `mysql_config` might indicate other libraries that should be specified on the link command as well.

If you get `undefined reference` errors for the `uncompress` or `compress` function, add `-lz` to the end of your link command and try again.

If you get `undefined reference` errors for a function that should exist on your system, such as `connect`, check the manual page for the function in question to determine which libraries you should add to the link command.

You might get `undefined reference` errors such as the following for functions that don't exist on



your system:

```
mf_format.o(.text+0x201): undefined reference to `__lxstat'
```

This usually means that your MySQL client library was compiled on a system that is not 100% compatible with yours. In this case, you should download the latest MySQL source distribution and compile MySQL yourself. See [Section 2.9, “MySQL Installation Using a Source Distribution”](#).

You might get undefined reference errors at runtime when you try to execute a MySQL program. If these errors specify symbols that start with `mysql_` or indicate that the `mysqlclient` library can't be found, it means that your system can't find the shared `libmysqlclient.so` library. The fix for this is to tell your system to search for shared libraries where the library is located. Use whichever of the following methods is appropriate for your system:

- Add the path to the directory where `libmysqlclient.so` is located to the `LD_LIBRARY_PATH` environment variable.
- Add the path to the directory where `libmysqlclient.so` is located to the `LD_LIBRARY` environment variable.
- Copy `libmysqlclient.so` to some directory that is searched by your system, such as `/lib`, and update the shared library information by executing `ldconfig`.

Another way to solve this problem is by linking your program statically with the `-static` option, or by removing the dynamic MySQL libraries before linking your code. Before trying the second method, you should be sure that no other programs are using the dynamic libraries.

## A.3.2. Problems with File Permissions

If you have problems with file permissions, the `UMASK` environment variable might be set incorrectly when `mysqld` starts. For example, MySQL might issue the following error message when you create a table:

```
ERROR: Can't find file: 'path/with/filename.frm' (Errcode: 13)
```

The default `UMASK` value is `0660`. You can change this behavior by starting `mysqld_safe` as follows:

```
shell> UMASK=384 # = 600 in octal
shell> export UMASK
shell> mysqld_safe &
```

By default, MySQL creates database and `RAID` directories with an access permission value of `0700`. You can modify this behavior by setting the `UMASK_DIR` variable. If you set its value, new directories are created with the combined `UMASK` and `UMASK_DIR` values. For example, if you want to give group access to all new directories, you can do this:

```
shell> UMASK_DIR=504 # = 770 in octal
shell> export UMASK_DIR
shell> mysqld_safe &
```

In MySQL 3.23.25 and above, MySQL assumes that the value for `UMASK` and `UMASK_DIR` is in octal if it starts with a zero.

See [Appendix F, \*Environment Variables\*](#).

## A.4. Administration-Related Issues

### A.4.1. How to Reset the Root Password

If you have never set a `root` password for MySQL, the server does not require a password at all for connecting as `root`. However, it is recommended to set a password for each account. See [Section 5.6.1, “General Security Guidelines”](#).

If you set a `root` password previously, but have forgotten what it was, you can set a new password. The following procedure is for Windows systems. The procedure for Unix systems is given later in this section.

The procedure under Windows:

1. Log on to your system as Administrator.
2. Stop the MySQL server if it is running. For a server that is running as a Windows service, go to the Services manager:

```
Start Menu -> Control Panel -> Administrative Tools -> Services
```

Then find the MySQL service in the list, and stop it.

If your server is not running as a service, you may need to use the Task Manager to force it to stop.

3. Create a text file and place the following command within it on a single line:

```
SET PASSWORD FOR 'root'@'localhost' = PASSWORD('MyNewPassword');
```

Save the file with any name. For this example the file will be `C:\mysql-init.txt`.

4. Open a console window to get to the DOS command prompt:

```
Start Menu -> Run -> cmd
```

5. We are assuming that you installed MySQL to `C:\mysql`. If you installed MySQL to another location, adjust the following commands accordingly.

At the DOS command prompt, execute this command:

```
C:\> C:\mysql\bin\mysqld-nt --init-file=C:\mysql-init.txt
```

The contents of the file named by the `--init-file` option are executed at server startup, changing the `root` password. After the server has started successfully, you should delete `C:\mysql-init.txt`.

Users of MySQL 4.1 and higher who install MySQL using the MySQL Installation Wizard may need to specify a `--defaults-file` option:

```
C:\> "C:\Program Files\MySQL\MySQL Server 4.1\bin\mysqld-nt.exe"  
--defaults-file="C:\Program Files\MySQL\MySQL Server 4.1\my.ini"  
--init-file=C:\mysql-init.txt
```

The appropriate `--defaults-file` setting can be found using the Services Manager:

```
Start Menu -> Control Panel -> Administrative Tools -> Services
```

Find the MySQL service in the list, right-click on it, and choose the [Properties](#) option. The [Path to executable](#) field contains the `--defaults-file` setting.

6. Stop the MySQL server, then restart it in normal mode again. If you run the server as a service, start it from the Windows Services window. If you start the server manually, use whatever command you normally use.
7. You should be able to connect using the new password.

In a Unix environment, the procedure for resetting the `root` password is as follows:

1. Log on to your system as either the Unix `root` user or as the same user that the `mysqld` server runs as.
2. Locate the `.pid` file that contains the server's process ID. The exact location and name of this file depend on your distribution, hostname, and configuration. Common locations are `/var/lib/mysql/`, `/var/run/mysqld/`, and `/usr/local/mysql/data/`. Generally, the filename has the extension of `.pid` and begins with either `mysqld` or your system's hostname.

You can stop the MySQL server by sending a normal `kill` (not `kill -9`) to the `mysqld` process, using the pathname of the `.pid` file in the following command:

```
shell> kill `cat /mysql-data-directory/host_name.pid`
```

Note the use of backticks rather than forward quotes with the `cat` command; these cause the output of `cat` to be substituted into the `kill` command.

3. Create a text file and place the following command within it on a single line:

```
SET PASSWORD FOR 'root'@'localhost' = PASSWORD('MyNewPassword');
```

Save the file with any name. For this example the file will be `~/mysql-init`.

4. Restart the MySQL server with the special `--init-file=~/mysql-init` option:

```
shell> mysqld_safe --init-file=~/mysql-init &
```

The contents of the init-file are executed at server startup, changing the root password. After the server has started successfully you should delete `~/mysql-init`.

5. You should be able to connect using the new password.

Alternatively, on any platform, you can set the new password using the `mysql` client (but this approach is less secure):

1. Stop `mysqld` and restart it with the `--skip-grant-tables --user=root` options (Windows users omit the `--user=root` portion).
2. Connect to the `mysqld` server with this command:

```
shell> mysql -u root
```

3. Issue the following statements in the `mysql` client:

```
mysql> UPDATE mysql.user SET Password=PASSWORD('newpwd')
-> WHERE User='root';
mysql> FLUSH PRIVILEGES;
```

Replace “*newpwd*” with the actual `root` password that you want to use.

4. You should be able to connect using the new password.

## A.4.2. What to Do If MySQL Keeps Crashing

Each MySQL version is tested on many platforms before it is released. This doesn't mean that there are no bugs in MySQL, but if there are bugs, they should be very few and can be hard to find. If you have a problem, it always helps if you try to find out exactly what crashes your system, because you have a much better chance of getting the problem fixed quickly.

First, you should try to find out whether the problem is that the `mysqld` server dies or whether your problem has to do with your client. You can check how long your `mysqld` server has been up by executing `mysqladmin version`. If `mysqld` has died and restarted, you may find the reason by looking in the server's error log. See [Section 5.11.1, “The Error Log”](#).

On some systems, you can find in the error log a stack trace of where `mysqld` died that you can resolve with the `resolve_stack_dump` program. See [Section E.1.4, “Using a Stack Trace”](#). Note that the variable values written in the error log may not always be 100% correct.

Many server crashes are caused by corrupted data files or index files. MySQL updates the files on disk with the `write()` system call after every SQL statement and before the client is notified about the result. (This is not true if you are running with `--delay-key-write`, in which case data files are written but not index files.) This means that data file contents are safe even if `mysqld` crashes, because the operating system ensures that the unflushed data is written to disk. You can force MySQL to flush everything to disk after every SQL statement by starting `mysqld` with the `--flush` option.

The preceding means that normally you should not get corrupted tables unless one of the following happens:

- The MySQL server or the server host was killed in the middle of an update.
- You have found a bug in `mysqld` that caused it to die in the middle of an update.
- Some external program is manipulating data files or index files at the same time as `mysqld` without locking the table properly.
- You are running many `mysqld` servers using the same data directory on a system that doesn't support good filesystem locks (normally handled by the `lockd` lock manager), or you are running multiple servers with external locking disabled.
- You have a crashed data file or index file that contains very corrupt data that confused `mysqld`.
- You have found a bug in the data storage code. This isn't likely, but it's at least possible. In this case, you can try to change the storage engine to another engine by using `ALTER TABLE` on a repaired copy of the table.

Because it is very difficult to know why something is crashing, first try to check whether things that work for others crash for you. Please try the following things:

- Stop the `mysqld` server with `mysqladmin shutdown`, run `myisamchk --silent --force */* .MYI` from the data directory to check all MyISAM tables, and restart `mysqld`. This ensures that you are running from a clean state. See [Chapter 5, Database Administration](#).
- Start `mysqld` with the `--log` option and try to determine from the information written to the log whether some specific query kills the server. About 95% of all bugs are related to a particular query. Normally, this is one of the last queries in the log file just before the server restarts. See [Section 5.11.2, “The General Query Log”](#). If you can repeatedly kill MySQL with a specific query, even when you have checked all tables just before issuing it, then you have been able to locate the bug and should submit a bug report for it. See [Section 1.8, “How to Report Bugs or Problems”](#).
- Try to make a test case that we can use to repeat the problem. See [Section E.1.6, “Making a Test Case If You Experience Table Corruption”](#).
- Try running the tests in the `mysql-test` directory and the MySQL benchmarks. See [Section 19.1.2, “MySQL Test Suite”](#). They should test MySQL rather well. You can also add code to the benchmarks that simulates your application. The benchmarks can be found in the `sql-bench` directory in a source distribution or, for a binary distribution, in the `sql-bench` directory under your MySQL installation directory.
- Try the `fork_big.pl` script. (It is located in the `tests` directory of source distributions.)
- If you configure MySQL for debugging, it is much easier to gather information about possible errors if something goes wrong. Configuring MySQL for debugging causes a safe memory allocator to be included that can find some errors. It also provides a lot of output about what is happening. Reconfigure MySQL with the `--with-debug` or `--with-debug=full` option to `configure` and then recompile. See [Section E.1, “Debugging a MySQL Server”](#).
- Make sure that you have applied the latest patches for your operating system.
- Use the `--skip-external-locking` option to `mysqld`. On some systems, the `lockd` lock manager does not work properly; the `--skip-external-locking` option tells `mysqld` not to use external locking. (This means that you cannot run two `mysqld` servers on the same data directory and that you must be careful if you use `myisamchk`. Nevertheless, it may be instructive to try the option as a test.)
- Have you tried `mysqladmin -u root processlist` when `mysqld` appears to be running but not responding? Sometimes `mysqld` is not comatose even though you might think so. The problem may be that all connections are in use, or there may be some internal lock problem. `mysqladmin -u root processlist` usually is able to make a connection even in these cases, and can provide useful information about the current number of connections and their status.
- Run the command `mysqladmin -i 5 status` or `mysqladmin -i 5 -r status` in a separate window to produce statistics while you run your other queries.
- Try the following:
  1. Start `mysqld` from `gdb` (or another debugger). See [Section E.1.3, “Debugging mysqld under gdb”](#).
  2. Run your test scripts.
  3. Print the backtrace and the local variables at the three lowest levels. In `gdb`, you can do this with the following commands when `mysqld` has crashed inside `gdb`:

```
backtrace
info local
up
info local
up
```

```
info local
```

With `gdb`, you can also examine which threads exist with `info threads` and switch to a specific thread with `thread N`, where `N` is the thread ID.

- Try to simulate your application with a Perl script to force MySQL to crash or misbehave.
- Send a normal bug report. See [Section 1.8, “How to Report Bugs or Problems”](#). Be even more detailed than usual. Because MySQL works for many people, it may be that the crash results from something that exists only on your computer (for example, an error that is related to your particular system libraries).
- If you have a problem with tables containing dynamic-length rows and you are using only `VARCHAR` columns (not `BLOB` or `TEXT` columns), you can try to change all `VARCHAR` to `CHAR` with `ALTER TABLE`. This forces MySQL to use fixed-size rows. Fixed-size rows take a little extra space, but are much more tolerant to corruption.

The current dynamic row code has been in use at MySQL AB for several years with very few problems, but dynamic-length rows are by nature more prone to errors, so it may be a good idea to try this strategy to see whether it helps.

- Do not rule out your server hardware when diagnosing problems. Defective hardware can be the cause of data corruption. Particular attention should be paid to both RAMS and hard-drives when troubleshooting hardware.

### A.4.3. How MySQL Handles a Full Disk

This section describes how MySQL responds to disk-full errors (such as “no space left on device”), and, as of MySQL 4.0.22, to quota-exceeded errors (such as “write failed” or “user block limit reached”).

This section is relevant for writes to `MyISAM` tables. As of MySQL 4.1.9, it also applies for writes to binary log files and binary log index file, except that references to “row” and “record” should be understood to mean “event.”

When a disk-full condition occurs, MySQL does the following:

- It checks once every minute to see whether there is enough space to write the current row. If there is enough space, it continues as if nothing had happened.
- Every 10 minutes it writes an entry to the log file, warning about the disk-full condition.

To alleviate the problem, you can take the following actions:

- To continue, you only have to free enough disk space to insert all records.
- To abort the thread, you must use `mysqladmin kill`. The thread is aborted the next time it checks the disk (in one minute).
- Other threads might be waiting for the table that caused the disk-full condition. If you have several “locked” threads, killing the one thread that is waiting on the disk-full condition allows the other threads to continue.

Exceptions to the preceding behavior are when you use `REPAIR TABLE` or `OPTIMIZE TABLE` or when the indexes are created in a batch after `LOAD DATA INFILE` or after an `ALTER TABLE` statement. All of these statements may create large temporary files that, if left to themselves, would cause big problems for the rest of the system. If the disk becomes full while MySQL is doing any of these operations, it removes the big temporary files and mark the table as crashed. The exception is that for `ALTER TABLE`, the old table is left unchanged.

## A.4.4. Where MySQL Stores Temporary Files

MySQL uses the value of the `TMPDIR` environment variable as the pathname of the directory in which to store temporary files. If you don't have `TMPDIR` set, MySQL uses the system default, which is normally `/tmp`, `/var/tmp`, or `/usr/tmp`. If the filesystem containing your temporary file directory is too small, you can use the `--tmpdir` option to `mysqld` to specify a directory in a filesystem where you have enough space.

Starting from MySQL 4.1, the `--tmpdir` option can be set to a list of several paths that are used in round-robin fashion. Paths should be separated by colon characters (':') on Unix and semicolon characters (;) on Windows, NetWare, and OS/2. **Note:** To spread the load effectively, these paths should be located on different *physical* disks, not different partitions of the same disk.

If the MySQL server is acting as a replication slave, you should not set `--tmpdir` to point to a directory on a memory-based filesystem or to a directory that is cleared when the server host restarts. A replication slave needs some of its temporary files to survive a machine restart so that it can replicate temporary tables or `LOAD DATA INFILE` operations. If files in the temporary file directory are lost when the server restarts, replication fails.

MySQL creates all temporary files as hidden files. This ensures that the temporary files are removed if `mysqld` is terminated. The disadvantage of using hidden files is that you do not see a big temporary file that fills up the filesystem in which the temporary file directory is located.

When sorting (`ORDER BY` or `GROUP BY`), MySQL normally uses one or two temporary files. The maximum disk space required is determined by the following expression:

```
(length of what is sorted + sizeof(row pointer))
* number of matched rows
* 2
```

The row pointer size is usually four bytes, but may grow in the future for really big tables.

For some `SELECT` queries, MySQL also creates temporary SQL tables. These are not hidden and have names of the form `SQL_*`.

`ALTER TABLE` creates a temporary table in the same directory as the original table.

## A.4.5. How to Protect or Change the MySQL Unix Socket File

The default location for the Unix socket file that the server uses for communication with local clients is `/tmp/mysql.sock`. (For some distribution formats, the directory might be different, such as `/var/lib/mysql` for RPMs.)

On some versions of Unix, anyone can delete files in the `/tmp` directory or other similar directories used for temporary files. If the socket file is located in such a directory on your system, this might cause problems.

On most versions of Unix, you can protect your `/tmp` directory so that files can be deleted only by their owners or the superuser (`root`). To do this, set the `sticky` bit on the `/tmp` directory by logging in as `root` and using the following command:

```
shell> chmod +t /tmp
```

You can check whether the `sticky` bit is set by executing `ls -ld /tmp`. If the last permission character is `t`, the bit is set.

Another approach is to change the place where the server creates the Unix socket file. If you do this, you should also let client programs know the new location of the file. You can specify the file location in several ways:

- Specify the path in a global or local option file. For example, put the following lines in `/etc/my.cnf`:

```
[mysqld]
socket=/path/to/socket

[client]
socket=/path/to/socket
```

See [Section 4.3.2, “Using Option Files”](#).

- Specify a `--socket` option on the command line to `mysqld_safe` and when you run client programs.
- Set the `MYSQL_UNIX_PORT` environment variable to the path of the Unix socket file.
- Recompile MySQL from source to use a different default Unix socket file location. Define the path to the file with the `--with-unix-socket-path` option when you run `configure`. See [Section 2.9.2, “Typical configure Options”](#).

You can test whether the new socket location works by attempting to connect to the server with this command:

```
shell> mysqladmin --socket=/path/to/socket version
```

## A.4.6. Time Zone Problems

If you have a problem with `SELECT NOW()` returning values in UTC and not your local time, you have to tell the server your current time zone. The same applies if `UNIX_TIMESTAMP()` returns the wrong value. This should be done for the environment in which the server runs; for example, in `mysqld_safe` or `mysql.server`. See [Appendix F, \*Environment Variables\*](#).

You can set the time zone for the server with the `--timezone=timezone_name` option to `mysqld_safe`. You can also set it by setting the `TZ` environment variable before you start `mysqld`.

The allowable values for `--timezone` or `TZ` are system-dependent. Consult your operating system documentation to see what values are acceptable.

## A.5. Query-Related Issues

### A.5.1. Case Sensitivity in Searches

By default, MySQL searches are not case sensitive (although there are some character sets that are never case insensitive, such as `czech`). This means that if you search with `col_name LIKE 'a%'`, you get all column values that start with `A` or `a`. If you want to make this search case sensitive, make sure that one of the operands has a case sensitive or binary collation. For example, if you are comparing a



column and a string that both have the `latin1` character set, you can use the `COLLATE` operator to cause either operand to have the `latin1_general_cs` or `latin1_bin` collation. For example:

```
col_name COLLATE latin1_general_cs LIKE 'a%'
col_name LIKE 'a%' COLLATE latin1_general_cs
col_name COLLATE latin1_bin LIKE 'a%'
col_name LIKE 'a%' COLLATE latin1_bin
```

If you want a column always to be treated in case-sensitive fashion, declare it with a case sensitive or binary collation. See [Section 13.1.5, “CREATE TABLE Syntax”](#).

Before MySQL 4.1, `COLLATE` is unavailable. Use the `BINARY` operator in expressions to treat a string as a binary string: `BINARY col_name LIKE 'a%'` or `col_name LIKE BINARY 'a%'`. In column declarations, use the `BINARY` attribute.

Simple comparison operations (`>=`, `>`, `=`, `<`, `<=`, sorting, and grouping) are based on each character’s “sort value.” Characters with the same sort value (such as ‘E’, ‘e’, and ‘Ã#Â©’) are treated as the same character.

## A.5.2. Problems Using `DATE` Columns

The format of a `DATE` value is ‘YYYY-MM-DD’. According to standard SQL, no other format is allowed. You should use this format in `UPDATE` expressions and in the `WHERE` clause of `SELECT` statements. For example:

```
mysql> SELECT * FROM tbl_name WHERE date >= '2003-05-05';
```

As a convenience, MySQL automatically converts a date to a number if the date is used in a numeric context (and vice versa). It is also smart enough to allow a “relaxed” string form when updating and in a `WHERE` clause that compares a date to a `TIMESTAMP`, `DATE`, or `DATETIME` column. (“Relaxed form” means that any punctuation character may be used as the separator between parts. For example, ‘2004-08-15’ and ‘2004#08#15’ are equivalent.) MySQL can also convert a string containing no separators (such as ‘20040815’), provided it makes sense as a date.

When you compare a `DATE`, `TIME`, `DATETIME`, or `TIMESTAMP` to a constant string with the `<`, `<=`, `=`, `>=`, `>`, or `BETWEEN` operators, MySQL normally converts the string to an internal long integer for faster comparison (and also for a bit more “relaxed” string checking). However, this conversion is subject to the following exceptions:

- When you compare two columns
- When you compare a `DATE`, `TIME`, `DATETIME`, or `TIMESTAMP` column to an expression
- When you use any other comparison method than those just listed, such as `IN` or `STRCMP()`.

For these exceptional cases, the comparison is done by converting the objects to strings and performing a string comparison.

To keep things safe, assume that strings are compared as strings and use the appropriate string functions if you want to compare a temporal value to a string.

The special date ‘0000-00-00’ can be stored and retrieved as ‘0000-00-00’. When using a ‘0000-00-00’ date through MyODBC, it is automatically converted to `NULL` in MyODBC 2.50.12 and above, because ODBC can’t handle this kind of date.

Because MySQL performs the conversions described above, the following statements work:

```
mysql> INSERT INTO tbl_name (idate) VALUES (19970505);
```

```
mysql> INSERT INTO tbl_name (idate) VALUES ('19970505');
mysql> INSERT INTO tbl_name (idate) VALUES ('97-05-05');
mysql> INSERT INTO tbl_name (idate) VALUES ('1997.05.05');
mysql> INSERT INTO tbl_name (idate) VALUES ('1997 05 05');
mysql> INSERT INTO tbl_name (idate) VALUES ('0000-00-00');

mysql> SELECT idate FROM tbl_name WHERE idate >= '1997-05-05';
mysql> SELECT idate FROM tbl_name WHERE idate >= 19970505;
mysql> SELECT MOD(idate,100) FROM tbl_name WHERE idate >= 19970505;
mysql> SELECT idate FROM tbl_name WHERE idate >= '19970505';
```

However, the following does not work:

```
mysql> SELECT idate FROM tbl_name WHERE STRCMP(idate,'20030505')=0;
```

`STRCMP()` is a string function, so it converts `idate` to a string in 'YYYY-MM-DD' format and performs a string comparison. It does not convert '20030505' to the date '2003-05-05' and perform a date comparison.

If the date cannot be converted to any reasonable value, a 0 is stored in the `DATE` column, which is retrieved as '0000-00-00'. This is both a speed and a convenience issue. We believe that the database server's responsibility is to retrieve the same date you stored (even if the data was not logically correct in all cases). We think it is up to the application and not the server to check the dates.

### A.5.3. Problems with `NULL` Values

The concept of the `NULL` value is a common source of confusion for newcomers to SQL, who often think that `NULL` is the same thing as an empty string ''. This is not the case. For example, the following statements are completely different:

```
mysql> INSERT INTO my_table (phone) VALUES (NULL);
mysql> INSERT INTO my_table (phone) VALUES ('');
```

Both statements insert a value into the `phone` column, but the first inserts a `NULL` value and the second inserts an empty string. The meaning of the first can be regarded as “phone number is not known” and the meaning of the second can be regarded as “the person is known to have no phone, and thus no phone number.”

To help with `NULL` handling, you can use the `IS NULL` and `IS NOT NULL` operators and the `IFNULL()` function.

In SQL, the `NULL` value is never true in comparison to any other value, even `NULL`. An expression that contains `NULL` always produces a `NULL` value unless otherwise indicated in the documentation for the operators and functions involved in the expression. All columns in the following example return `NULL`:

```
mysql> SELECT NULL, 1+NULL, CONCAT('Invisible',NULL);
```

If you want to search for column values that are `NULL`, you cannot use an `expr = NULL` test. The following statement returns no rows, because `expr = NULL` is never true for any expression:

```
mysql> SELECT * FROM my_table WHERE phone = NULL;
```

To look for `NULL` values, you must use the `IS NULL` test. The following statements show how to find the `NULL` phone number and the empty phone number:

```
mysql> SELECT * FROM my_table WHERE phone IS NULL;
mysql> SELECT * FROM my_table WHERE phone = '';
```

See [Section 3.3.4.6, “Working with `NULL` Values”](#), for additional information and examples.

You can add an index on a column that can have `NULL` values if you are using MySQL 3.23.2 or newer and are using the `MyISAM`, `InnoDB`, or `BDB` storage engine. As of MySQL 4.0.2, the `MEMORY` storage engine also supports `NULL` values in indexes. Otherwise, you must declare an indexed column `NOT NULL` and you cannot insert `NULL` into the column.

When reading data with `LOAD DATA INFILE`, empty or missing columns are updated with `' '`. If you want a `NULL` value in a column, you should use `\N` in the data file. The literal word `"NULL"` may also be used under some circumstances. See [Section 13.2.5, "LOAD DATA INFILE Syntax"](#).

When using `DISTINCT`, `GROUP BY`, or `ORDER BY`, all `NULL` values are regarded as equal.

When using `ORDER BY`, `NULL` values are presented first, or last if you specify `DESC` to sort in descending order. Exception: In MySQL 4.0.2 through 4.0.10, `NULL` values sort first regardless of sort order.

Aggregate (summary) functions such as `COUNT()`, `MIN()`, and `SUM()` ignore `NULL` values. The exception to this is `COUNT(*)`, which counts rows and not individual column values. For example, the following statement produces two counts. The first is a count of the number of rows in the table, and the second is a count of the number of non-`NULL` values in the `age` column:

```
mysql> SELECT COUNT(*), COUNT(age) FROM person;
```

For some data types, MySQL handles `NULL` values specially. If you insert `NULL` into a `TIMESTAMP` column, the current date and time is inserted. If you insert `NULL` into an integer column that has the `AUTO_INCREMENT` attribute, the next number in the sequence is inserted.

## A.5.4. Problems with Column Aliases

You can use an alias to refer to a column in `GROUP BY`, `ORDER BY`, or `HAVING` clauses. Aliases can also be used to give columns better names:

```
SELECT SQRT(a*b) AS root FROM tbl_name GROUP BY root HAVING root > 0;
SELECT id, COUNT(*) AS cnt FROM tbl_name GROUP BY id HAVING cnt > 0;
SELECT id AS 'Customer identity' FROM tbl_name;
```

Standard SQL doesn't allow you to refer to a column alias in a `WHERE` clause. This restriction is imposed because when the `WHERE` code is executed, the column value may not yet be determined. For example, the following query is illegal:

```
SELECT id, COUNT(*) AS cnt FROM tbl_name WHERE cnt > 0 GROUP BY id;
```

The `WHERE` statement is executed to determine which rows should be included in the `GROUP BY` part, whereas `HAVING` is used to decide which rows from the result set should be used.

## A.5.5. Rollback Failure for Non-Transactional Tables

If you receive the following message when trying to perform a `ROLLBACK`, it means that one or more of the tables you used in the transaction do not support transactions:

```
Warning: Some non-transactional changed tables couldn't be rolled back
```

These non-transactional tables are not affected by the `ROLLBACK` statement.

If you were not deliberately mixing transactional and non-transactional tables within the transaction, the most likely cause for this message is that a table you thought was transactional actually is not. This can happen if you try to create a table using a transactional storage engine that is not supported by your `mysqld` server (or that was disabled with a startup option). If `mysqld` doesn't support a storage engine,

it instead creates the table as a `MyISAM` table, which is non-transactional.

You can check the storage engine for a table by using either of these statements:

```
SHOW TABLE STATUS LIKE 'tbl_name';
SHOW CREATE TABLE tbl_name;
```

See [Section 13.5.4.18, “SHOW TABLE STATUS Syntax”](#), and [Section 13.5.4.5, “SHOW CREATE TABLE Syntax”](#).

You can check which storage engines your `mysqld` server supports by using this statement:

```
SHOW ENGINES;
```

Before MySQL 4.1.2, `SHOW ENGINES` is unavailable. Use the following statement instead and check the value of the variable that is associated with the storage engine in which you are interested:

```
SHOW VARIABLES LIKE 'have_%';
```

For example, to determine whether the `InnoDB` storage engine is available, check the value of the `have_innodb` variable.

See [Section 13.5.4.8, “SHOW ENGINES Syntax”](#), and [Section 13.5.4.20, “SHOW VARIABLES Syntax”](#).

## A.5.6. Deleting Rows from Related Tables

MySQL does not support subqueries prior to version 4.1, or the use of more than one table in the `DELETE` statement prior to version 4.0. If your version of MySQL does not support subqueries or multiple-table `DELETE` statements, you can use the following approach to delete rows from two related tables:

1. `SELECT` the rows based on some `WHERE` condition in the main table.
2. `DELETE` the rows in the main table based on the same condition.
3. `DELETE FROM related_table WHERE related_column IN (selected_rows)`.

If the total length of the `DELETE` statement for `related_table` is more than 1MB (the default value of the `max_allowed_packet` system variable), you should split it into smaller parts and execute multiple `DELETE` statements. You probably get the fastest `DELETE` by specifying only 100 to 1,000 `related_column` values per statement if the `related_column` is indexed. If the `related_column` isn't indexed, the speed is independent of the number of arguments in the `IN` clause.

## A.5.7. Solving Problems with No Matching Rows

If you have a complicated query that uses many tables but that doesn't return any rows, you should use the following procedure to find out what is wrong:

1. Test the query with `EXPLAIN` to check whether you can find something that is obviously wrong. See [Section 7.2.1, “Optimizing Queries with EXPLAIN”](#).
2. Select only those columns that are used in the `WHERE` clause.
3. Remove one table at a time from the query until it returns some rows. If the tables are large, it's a good idea to use `LIMIT 10` with the query.

- Issue a `SELECT` for the column that should have matched a row against the table that was last removed from the query.
- If you are comparing `FLOAT` or `DOUBLE` columns with numbers that have decimals, you can't use equality (`=`) comparisons. This problem is common in most computer languages because not all floating-point values can be stored with exact precision. In some cases, changing the `FLOAT` to a `DOUBLE` fixes this. See [Section A.5.8, "Problems with Floating-Point Comparisons"](#).

Similar problems may be encountered when comparing `DECIMAL` values.

- If you still can't figure out what's wrong, create a minimal test that can be run with `mysql test < query.sql` that shows your problems. You can create a test file by dumping the tables with `mysqldump --quick db_name tbl_name_1 ... tbl_name_n > query.sql`. Open the file in an editor, remove some insert lines (if there are more than needed to demonstrate the problem), and add your `SELECT` statement at the end of the file.

Verify that the test file demonstrates the problem by executing these commands:

```
shell> mysqladmin create test2
shell> mysql test2 < query.sql
```

Attach the test file to a bug report, which you can file using the instructions in [Section 1.8, "How to Report Bugs or Problems"](#).

## A.5.8. Problems with Floating-Point Comparisons

Floating-point numbers sometimes cause confusion because they are approximate. That is, they are not stored as exact values inside computer architecture. What you can see on the screen usually is not the exact value of the number. The `FLOAT` and `DOUBLE` data types are such, and `DECIMAL` operations before MySQL 5.0.3 are approximate as well.

Prior to MySQL 5.0.3, `DECIMAL` columns store values with exact precision because they are represented as strings, but calculations on `DECIMAL` values are done using floating-point operations. As of 5.0.3, MySQL performs `DECIMAL` operations with a precision of 64 decimal digits, which should solve most common inaccuracy problems when it comes to `DECIMAL` columns. (If your server is from MySQL 5.0.3 or higher, but you have `DECIMAL` columns in tables that were created before 5.0.3, the old behavior still applies to those columns. To convert the tables to the newer `DECIMAL` format, dump them with `mysqldump` and reload them.)

The following example (for versions of MySQL older than 5.0.3) demonstrates the problem. It shows that even for older `DECIMAL` columns, calculations that are done using floating-point operations are subject to floating-point error. (Were you to replace the `DECIMAL` columns with `FLOAT`, similar problems would occur for all versions of MySQL.)

```
mysql> CREATE TABLE t1 (i INT, d1 DECIMAL(9,2), d2 DECIMAL(9,2));
mysql> INSERT INTO t1 VALUES (1, 101.40, 21.40), (1, -80.00, 0.00),
-> (2, 0.00, 0.00), (2, -13.20, 0.00), (2, 59.60, 46.40),
-> (2, 30.40, 30.40), (3, 37.00, 7.40), (3, -29.60, 0.00),
-> (4, 60.00, 15.40), (4, -10.60, 0.00), (4, -34.00, 0.00),
-> (5, 33.00, 0.00), (5, -25.80, 0.00), (5, 0.00, 7.20),
-> (6, 0.00, 0.00), (6, -51.40, 0.00);

mysql> SELECT i, SUM(d1) AS a, SUM(d2) AS b
-> FROM t1 GROUP BY i HAVING a <> b;
```

i	a	b
1	21.40	21.40
2	76.80	76.80
3	7.40	7.40
4	15.40	15.40
5	7.20	7.20

```
| 6 | -51.40 | 0.00 |
+-----+-----+
```

The result is correct. Although the first five records look like they should not satisfy the comparison (the values of `a` and `b` do not appear to be different), they may do so because the difference between the numbers shows up around the tenth decimal or so, depending on factors such as computer architecture or the compiler version or optimization level. For example, different CPUs may evaluate floating-point numbers differently.

As of MySQL 5.0.3, you will get only the last row in the above result.

The problem cannot be solved by using `ROUND()` or similar functions, because the result is still a floating-point number:

```
mysql> SELECT i, ROUND(SUM(d1), 2) AS a, ROUND(SUM(d2), 2) AS b
-> FROM t1 GROUP BY i HAVING a <> b;
+-----+-----+
| i | a | b |
+-----+-----+
| 1 | 21.40 | 21.40 |
| 2 | 76.80 | 76.80 |
| 3 | 7.40 | 7.40 |
| 4 | 15.40 | 15.40 |
| 5 | 7.20 | 7.20 |
| 6 | -51.40 | 0.00 |
+-----+-----+
```

This is what the numbers in column `a` look like when displayed with more decimal places:

```
mysql> SELECT i, ROUND(SUM(d1), 2)*1.0000000000000000 AS a,
-> ROUND(SUM(d2), 2) AS b FROM t1 GROUP BY i HAVING a <> b;
+-----+-----+-----+
| i | a | b |
+-----+-----+-----+
| 1 | 21.3999999999999986 | 21.40 |
| 2 | 76.7999999999999972 | 76.80 |
| 3 | 7.4000000000000004 | 7.40 |
| 4 | 15.4000000000000004 | 15.40 |
| 5 | 7.2000000000000002 | 7.20 |
| 6 | -51.399999999999986 | 0.00 |
+-----+-----+-----+
```

Depending on your computer architecture, you may or may not see similar results. For example, on some machines you may get the “correct” results by multiplying both arguments by 1, as the following example shows.

**Warning:** Never use this method in your applications. It is not an example of a trustworthy method!

```
mysql> SELECT i, ROUND(SUM(d1), 2)*1 AS a, ROUND(SUM(d2), 2)*1 AS b
-> FROM t1 GROUP BY i HAVING a <> b;
+-----+-----+
| i | a | b |
+-----+-----+
| 6 | -51.40 | 0.00 |
+-----+-----+
```

The reason that the preceding example seems to work is that on the particular machine where the test was done, CPU floating-point arithmetic happens to round the numbers to the same value. However, there is no rule that any CPU should do so, so this method cannot be trusted.

The correct way to do floating-point number comparison is to first decide on an acceptable tolerance for differences between the numbers and then do the comparison against the tolerance value. For example, if we agree that floating-point numbers should be regarded the same if they are same within a precision of one in ten thousand (0.0001), the comparison should be written to find differences larger than the tolerance value:

```
mysql> SELECT i, SUM(d1) AS a, SUM(d2) AS b FROM t1
-> GROUP BY i HAVING ABS(a - b) > 0.0001;
+-----+-----+-----+
| i | a | b |
+-----+-----+-----+
| 6 | -51.40 | 0.00 |
+-----+-----+-----+
1 row in set (0.00 sec)
```

Conversely, to get rows where the numbers are the same, the test should find differences within the tolerance value:

```
mysql> SELECT i, SUM(d1) AS a, SUM(d2) AS b FROM t1
-> GROUP BY i HAVING ABS(a - b) <= 0.0001;
+-----+-----+-----+
| i | a | b |
+-----+-----+-----+
| 1 | 21.40 | 21.40 |
| 2 | 76.80 | 76.80 |
| 3 | 7.40 | 7.40 |
| 4 | 15.40 | 15.40 |
| 5 | 7.20 | 7.20 |
+-----+-----+-----+
```

## A.6. Optimizer-Related Issues

MySQL uses a cost-based optimizer to determine the best way to resolve a query. In many cases, MySQL can calculate the best possible query plan, but sometimes MySQL doesn't have enough information about the data at hand and has to make "educated" guesses about the data.

For the cases when MySQL does not do the "right" thing, tools that you have available to help MySQL are:

- Use the `EXPLAIN` statement to get information about how MySQL processes a query. To use it, just add the keyword `EXPLAIN` to the front of your `SELECT` statement:

```
mysql> EXPLAIN SELECT * FROM t1, t2 WHERE t1.i = t2.i;
```

`EXPLAIN` is discussed in more detail in [Section 7.2.1, "Optimizing Queries with EXPLAIN"](#).

- Use `ANALYZE TABLE tbl_name` to update the key distributions for the scanned table. See [Section 13.5.2.1, "ANALYZE TABLE Syntax"](#).
- Use `FORCE INDEX` for the scanned table to tell MySQL that table scans are very expensive compared to using the given index. See [Section 13.2.7, "SELECT Syntax"](#).

```
SELECT * FROM t1, t2 FORCE INDEX (index_for_column)
WHERE t1.col_name=t2.col_name;
```

`USE INDEX` and `IGNORE INDEX` may also be useful.

- Global and table-level `STRAIGHT_JOIN`. See [Section 13.2.7, "SELECT Syntax"](#).
- You can tune global or thread-specific system variables. For example, Start `mysqld` with the `--max-seeks-for-key=1000` option or use `SET max_seeks_for_key=1000` to tell the optimizer to assume that no key scan causes more than 1,000 key seeks. See [Section 5.2.2, "System Variables"](#).

## A.7. Table Definition-Related Issues

## A.7.1. Problems with `ALTER TABLE`

`ALTER TABLE` changes a table to the current character set. If you get a duplicate-key error during `ALTER TABLE`, the cause is either that the new character sets maps two keys to the same value or that the table is corrupted. In the latter case, you should run `REPAIR TABLE` on the table.

If `ALTER TABLE` dies with the following error, the problem may be that MySQL crashed during an earlier `ALTER TABLE` operation and there is an old table named `A-xxx` or `B-xxx` lying around:

```
Error on rename of './database/name.frm'
to './database/B-xxx.frm' (Errcode: 17)
```

In this case, go to the MySQL data directory and delete all files that have names starting with `A-` or `B-`. (You may want to move them elsewhere instead of deleting them.)

`ALTER TABLE` works in the following way:

- Create a new table named `A-xxx` with the requested structural changes.
- Copy all rows from the original table to `A-xxx`.
- Rename the original table to `B-xxx`.
- Rename `A-xxx` to your original table name.
- Delete `B-xxx`.

If something goes wrong with the renaming operation, MySQL tries to undo the changes. If something goes seriously wrong (although this shouldn't happen), MySQL may leave the old table as `B-xxx`. A simple rename of the table files at the system level should get your data back.

If you use `ALTER TABLE` on a transactional table or if you are using Windows or OS/2, `ALTER TABLE` unlocks the table if you had done a `LOCK TABLE` on it. This is done because `InnoDB` and these operating systems cannot drop a table that is in use.

## A.7.2. How to Change the Order of Columns in a Table

First, consider whether you really need to change the column order in a table. The whole point of SQL is to abstract the application from the data storage format. You should always specify the order in which you wish to retrieve your data. The first of the following statements returns columns in the order `col_name1, col_name2, col_name3`, whereas the second returns them in the order `col_name1, col_name3, col_name2`:

```
mysql> SELECT col_name1, col_name2, col_name3 FROM tbl_name;
mysql> SELECT col_name1, col_name3, col_name2 FROM tbl_name;
```

If you decide to change the order of table columns anyway, you can do so as follows:

1. Create a new table with the columns in the new order.
2. Execute this statement:

```
mysql> INSERT INTO new_table
-> SELECT columns-in-new-order FROM old_table;
```

3. Drop or rename `old_table`.



4. Rename the new table to the original name:

```
mysql> ALTER TABLE new_table RENAME old_table;
```

`SELECT *` is quite suitable for testing queries. However, in an application, you should *never* rely on using `SELECT *` and retrieving the columns based on their position. The order and position in which columns are returned does not remain the same if you add, move, or delete columns. A simple change to your table structure could cause your application to fail.

### A.7.3. TEMPORARY TABLE Problems

The following list indicates limitations on the use of `TEMPORARY` tables:

- A `TEMPORARY` table can only be of type `HEAP`, `ISAM`, `MyISAM`, `MERGE`, or `InnoDB`.

Temporary tables are not supported for MySQL Cluster.

- You cannot refer to a `TEMPORARY` table more than once in the same query. For example, the following does not work:

```
mysql> SELECT * FROM temp_table, temp_table AS t2;
ERROR 1137: Can't reopen table: 'temp_table'
```

- The `SHOW TABLES` statement does not list `TEMPORARY` tables.
- You cannot use `RENAME` to rename a `TEMPORARY` table. However, you can use `ALTER TABLE` instead:

```
mysql> ALTER TABLE orig_name RENAME new_name;
```

- There are known issues in using temporary tables with replication. See [Section 6.7, “Replication Features and Known Problems”](#), for more information.

## A.8. Known Issues in MySQL

This section is a list of the known issues in recent versions of MySQL.

For information about platform-specific issues, see the installation and porting instructions in [Section 2.13, “Operating System-Specific Notes”](#), and [Appendix E, \*Porting to Other Systems\*](#).

### A.8.1. Issues in MySQL 3.23 Fixed in a Later MySQL Version

The following known issues have not been fixed in MySQL 3.23 for various reasons, and are not classified as critical.

- Fixed in MySQL 4.0: Avoid using spaces at the end of column names because this can cause unexpected behavior. ([Bug#4196](#) [<http://bugs.mysql.com/4196>])
- Fixed in MySQL 4.0.12: You can get a deadlock (hung thread) if you use `LOCK TABLE` to lock multiple tables and then in the same connection use `DROP TABLE` to drop one of them while another thread is trying to lock it. (To break the deadlock, you can use `KILL` to terminate any of the threads involved.)

- Fixed in MySQL 4.0.11: `SELECT MAX(key_column) FROM t1,t2,t3...` where one of the tables are empty doesn't return `NULL` but instead returns the maximum value for the column.
- `DELETE FROM heap_table` without a `WHERE` clause doesn't work on a locked `HEAP` table.

## A.8.2. Issues in MySQL 4.0 Fixed in a Later Version

The following known issues have not been fixed in MySQL 4.0 for various reasons, and are not classified as critical.

- Fixed in MySQL 4.1.10: Using `HAVING`, you can get a crash or wrong result if you use an alias to a `RAND()` function. This will not be fixed in 4.0 because the fix may break compatibility with some applications.
- Fixed in MySQL 4.1.1: In a `UNION`, the first `SELECT` determines the type, `max_length`, and `NULL` properties for the resulting columns.
- Fixed in MySQL 4.1: In `DELETE` with many tables, you can't refer to tables to be deleted through an alias.
- Fixed in MySQL 4.1.2: You cannot mix `UNION ALL` and `UNION DISTINCT` in the same query. If you use `ALL` for one `UNION`, it is used for all of them.
- `FLUSH TABLES WITH READ LOCK` does not block `CREATE TABLE`, which may cause a problem with the binary log position when doing a full backup of tables and the binary log.
- Fixed in MySQL 4.1.8: `mysqldump --single-transaction --master-data` behaved like `mysqldump --master-data`, so the dump was a blocking one.
- When using the `RPAD()` function (or any function adding spaces to the right) in a query that had to be resolved by using a temporary table, all resulting strings had rightmost spaces removed (i.e. `RPAD()` did not work).

## A.8.3. Issues in MySQL 4.1 Fixed in a Later Version

The following known issues have not been fixed in MySQL 4.1 for various reasons, and are not classified as critical.

- Fixed in 5.0.3: `VARCHAR` and `VARBINARY` did not remember end space.

## A.8.4. Open Issues in MySQL

The following problems are known and fixing them is a high priority:

- If you compare a `NULL` value to a subquery using `ALL/ANY/SOME` and the subquery returns an empty result, the comparison might evaluate to the non-standard result of `NULL` rather than to `TRUE` or `FALSE`. This will be fixed in MySQL 5.1.
- Subquery optimization for `IN` is not as effective as for `=`.
- Even if you use `lower_case_table_names=2` (which enables MySQL to remember the case used for databases and table names), MySQL does not remember the case used for database names

for the function `DATABASE()` or within the various logs (on case-insensitive systems).

- Dropping a `FOREIGN KEY` constraint doesn't work in replication because the constraint may have another name on the slave.
- `REPLACE` (and `LOAD DATA` with the `REPLACE` option) does not trigger `ON DELETE CASCADE`.
- `DISTINCT` with `ORDER BY` doesn't work inside `GROUP_CONCAT()` if you don't use all and only those columns that are in the `DISTINCT` list.
- If one user has a long-running transaction and another user drops a table that is updated in the transaction, there is small chance that the binary log may contain the `DROP TABLE` command before the table is used in the transaction itself. We plan to fix this by having the `DROP TABLE` command wait until the table is not being used in any transaction.
- When inserting a big integer value (between  $2^{63}$  and  $2^{64}-1$ ) into a decimal or string column, it is inserted as a negative value because the number is evaluated in a signed integer context.
- `FLUSH TABLES WITH READ LOCK` does not block `COMMIT` if the server is running without binary logging, which may cause a problem (of consistency between tables) when doing a full backup.
- `ANALYZE TABLE` on a `BDB` table may in some cases make the table unusable until you restart `mysqld`. If this happens, look for errors of the following form in the MySQL error file:

```
001207 22:07:56 bdb: log_flush: LSN past current end-of-log
```

- Don't execute `ALTER TABLE` on a `BDB` table on which you are running multiple-statement transactions until all those transactions complete. (The transaction might be ignored.)
- `ANALYZE TABLE`, `OPTIMIZE TABLE`, and `REPAIR TABLE` may cause problems on tables for which you are using `INSERT DELAYED`.
- Performing `LOCK TABLE ...` and `FLUSH TABLES ...` doesn't guarantee that there isn't a half-finished transaction in progress on the table.
- `BDB` tables are relatively slow to open. If you have many `BDB` tables in a database, it takes a long time to use the `mysql` client on the database if you are not using the `-A` option or if you are using `rehash`. This is especially noticeable when you have a large table cache.
- Replication uses query-level logging: The master writes the executed queries to the binary log. This is a very fast, compact, and efficient logging method that works perfectly in most cases.

It is possible for the data on the master and slave to become different if a query is designed in such a way that the data modification is non-deterministic (generally not a recommended practice, even outside of replication).

For example:

- `CREATE ... SELECT` or `INSERT ... SELECT` statements that insert zero or `NULL` values into an `AUTO_INCREMENT` column.
- `DELETE` if you are deleting rows from a table that has foreign keys with `ON DELETE CASCADE` properties.
- `REPLACE ... SELECT`, `INSERT IGNORE ... SELECT` if you have duplicate key values in the inserted data.

**If and only if the preceding queries have no `ORDER BY` clause guaranteeing a deterministic or-**

der.

For example, for `INSERT ... SELECT` with no `ORDER BY`, the `SELECT` may return rows in a different order (which results in a row having different ranks, hence getting a different number in the `AUTO_INCREMENT` column), depending on the choices made by the optimizers on the master and slave.

A query is optimized differently on the master and slave only if:

- The files used by the two queries are not exactly the same; for example, `OPTIMIZE TABLE` was run on the master tables and not on the slave tables. (To fix this, `OPTIMIZE TABLE`, `ANALYZE TABLE`, and `REPAIR TABLE` are written to the binary log as of MySQL 4.1.1).
- The table is stored using a different storage engine on the master than on the slave. (It is possible to use different storage engines on the master and slave. For example, you can use `InnoDB` on the master, but `MyISAM` on the slave if the slave has less available disk space.)
- MySQL buffer sizes (`key_buffer_size`, and so on) are different on the master and slave.
- The master and slave run different MySQL versions, and the optimizer code differs between these versions.

This problem may also affect database restoration using `mysqlbinlog|mysql`.

The easiest way to avoid this problem is to add an `ORDER BY` clause to the aforementioned non-deterministic queries to ensure that the rows are always stored or modified in the same order.

In future MySQL versions, we will automatically add an `ORDER BY` clause when needed.

The following issues are known and will be fixed in due time:

- Log filenames are based on the server hostname (if you don't specify a filename with the startup option). You have to use options such as `--log-bin=old_host_name-bin` if you change your hostname to something else. Another option is to rename the old files to reflect your hostname change (if these are binary logs, you need to edit the binary log index file and fix the binlog names there as well). See [Section 5.2.1, "Command Options"](#).
- `mysqlbinlog` does not delete temporary files left after a `LOAD DATA INFILE` command. See [Section 8.10, "mysqlbinlog — Utility for Processing Binary Log Files"](#).
- `RENAME` doesn't work with `TEMPORARY` tables or tables used in a `MERGE` table.
- Due to the way table format (`.frm`) files are stored, you cannot use character 255 (`CHAR(255)`) in table names, column names, or enumerations. This is scheduled to be fixed in version 5.1 when we implement new table definition format files.
- When using `SET CHARACTER SET`, you can't use translated characters in database, table, and column names.
- You can't use `'_'` or `'%'` with `ESCAPE` in `LIKE ... ESCAPE`.
- If you have a `DECIMAL` column in which the same number is stored in different formats (for example, `+01.00`, `1.00`, `01.00`), `GROUP BY` may regard each value as a different value.
- You cannot build the server in another directory when using MIT-pthreads. Because this requires changes to MIT-pthreads, we are not likely to fix this. See [Section 2.9.5, "MIT-pthreads Notes"](#).

- `BLOB` and `TEXT` values can't “reliably” be used in `GROUP BY`, `ORDER BY` or `DISTINCT`. Only the first `max_sort_length` bytes are used when comparing `BLOB` values in these cases. The default value of `max_sort_length` value is 1024 and can be changed at server startup time. As of MySQL 4.0.3, it can be changed at runtime. For older versions, a workaround is to use a substring. For example:

```
SELECT DISTINCT LEFT(blob_col,2048) FROM tbl_name;
```

- Numeric calculations are done with `BIGINT` or `DOUBLE` (both are normally 64 bits long). Which precision you get depends on the function. The general rule is that bit functions are performed with `BIGINT` precision, `IF` and `ELT()` with `BIGINT` or `DOUBLE` precision, and the rest with `DOUBLE` precision. You should try to avoid using unsigned long long values if they resolve to be larger than 63 bits (9223372036854775807) for anything other than bit fields. MySQL Server 4.0 has better `BIGINT` handling than 3.23.
- You can have up to 255 `ENUM` and `SET` columns in one table.
- In `MIN()`, `MAX()`, and other aggregate functions, MySQL currently compares `ENUM` and `SET` columns by their string value rather than by the string's relative position in the set.
- `mysqld_safe` redirects all messages from `mysqld` to the `mysqld` log. One problem with this is that if you execute `mysqladmin refresh` to close and reopen the log, `stdout` and `stderr` are still redirected to the old log. If you use `--log` extensively, you should edit `mysqld_safe` to log to `host_name.err` instead of `host_name.log` so that you can easily reclaim the space for the old log by deleting it and executing `mysqladmin refresh`.
- In an `UPDATE` statement, columns are updated from left to right. If you refer to an updated column, you get the updated value instead of the original value. For example, the following statement increments `KEY` by 2, not 1:

```
mysql> UPDATE tbl_name SET KEY=KEY+1,KEY=KEY+1;
```

- You can refer to multiple temporary tables in the same query, but you cannot refer to any given temporary table more than once. For example, the following doesn't work:

```
mysql> SELECT * FROM temp_table, temp_table AS t2;
ERROR 1137: Can't reopen table: 'temp_table'
```

- The optimizer may handle `DISTINCT` differently when you are using “hidden” columns in a join than when you are not. In a join, hidden columns are counted as part of the result (even if they are not shown), whereas in normal queries, hidden columns don't participate in the `DISTINCT` comparison. We will probably change this in the future to never compare the hidden columns when executing `DISTINCT`.

An example of this is:

```
SELECT DISTINCT mp3id FROM band_downloads
WHERE userid = 9 ORDER BY id DESC;
```

and

```
SELECT DISTINCT band_downloads.mp3id
FROM band_downloads,band_mp3
WHERE band_downloads.userid = 9
AND band_mp3.id = band_downloads.mp3id
ORDER BY band_downloads.id DESC;
```

In the second case, using MySQL Server 3.23.x, you may get two identical rows in the result set

(because the values in the hidden `id` column may differ).

Note that this happens only for queries where that do not have the `ORDER BY` columns in the result.

- If you execute a `PROCEDURE` on a query that returns an empty set, in some cases the `PROCEDURE` does not transform the columns.
- Creation of a table of type `MERGE` doesn't check whether the underlying tables are compatible types.
- If you use `ALTER TABLE` to add a `UNIQUE` index to a table used in a `MERGE` table and then add a normal index on the `MERGE` table, the key order is different for the tables if there was an old, non-`UNIQUE` key in the table. This is because `ALTER TABLE` puts `UNIQUE` indexes before normal indexes to be able to detect duplicate keys as early as possible.

---

# Appendix B. Error Codes and Messages

This appendix lists the errors that may appear when you call MySQL from any host language. The first list displays server error messages. The second list displays client program messages.

## B.1. Server Error Codes and Messages

MySQL programs have access to several types of error information when the server returns an error. For example, the `mysql` client program displays errors using the following format:

```
shell> SELECT * FROM no_such_table;
ERROR 1146 (42S02): Table 'test.no_such_table' doesn't exist
```

The message displayed contains three types of information:

- A numeric error value (`1146`). This number is MySQL-specific and is not portable to other database systems.
- A five-character SQLSTATE value (`'42S02'`). The values are specified by ANSI SQL and ODBC and are more standardized. Not all MySQL error numbers are mapped to SQLSTATE error codes. The value `'HY000'` (general error) is used for unmapped errors.
- A string that provides a textual description of the error.

Server error information comes from the following source files. For details about the way that error information is defined, see the MySQL Internals manual, available at <http://dev.mysql.com/doc/>.

- The Error values and the symbols in parentheses correspond to definitions in the `include/mysql_d_error.h` MySQL source file.
- The SQLSTATE values correspond to definitions in the `include/sql_state.h` MySQL source file.

SQLSTATE error codes are displayed only if you use MySQL version 4.1 and up. SQLSTATE codes were added for compatibility with X/Open, ANSI, and ODBC behavior.

- The Message values correspond to the error messages that are listed in the `sql/share/errmsg.txt` file. `%d` and `%s` represent numbers and strings, respectively, that are substituted into the messages when they are displayed.

Because updates are frequent, it is possible that those files will contain additional error information not listed here.

- Error: `1000` SQLSTATE: `HY000` (`ER_HASHCHK`)  
Message: hashchk
- Error: `1001` SQLSTATE: `HY000` (`ER_NISAMCHK`)  
Message: isamchk
- Error: `1002` SQLSTATE: `HY000` (`ER_NO`)

Message: NO

- Error: 1003 SQLSTATE: HY000 (ER\_YES)

Message: YES

- Error: 1004 SQLSTATE: HY000 (ER\_CANT\_CREATE\_FILE)

Message: Can't create file '%s' (errno: %d)

- Error: 1005 SQLSTATE: HY000 (ER\_CANT\_CREATE\_TABLE)

Message: Can't create table '%s' (errno: %d)

- Error: 1006 SQLSTATE: HY000 (ER\_CANT\_CREATE\_DB)

Message: Can't create database '%s' (errno: %d)

- Error: 1007 SQLSTATE: HY000 (ER\_DB\_CREATE\_EXISTS)

Message: Can't create database '%s'; database exists

- Error: 1008 SQLSTATE: HY000 (ER\_DB\_DROP\_EXISTS)

Message: Can't drop database '%s'; database doesn't exist

- Error: 1009 SQLSTATE: HY000 (ER\_DB\_DROP\_DELETE)

Message: Error dropping database (can't delete '%s', errno: %d)

- Error: 1010 SQLSTATE: HY000 (ER\_DB\_DROP\_RMDIR)

Message: Error dropping database (can't rmdir '%s', errno: %d)

- Error: 1011 SQLSTATE: HY000 (ER\_CANT\_DELETE\_FILE)

Message: Error on delete of '%s' (errno: %d)

- Error: 1012 SQLSTATE: HY000 (ER\_CANT\_FIND\_SYSTEM\_REC)

Message: Can't read record in system table

- Error: 1013 SQLSTATE: HY000 (ER\_CANT\_GET\_STAT)

Message: Can't get status of '%s' (errno: %d)

- Error: 1014 SQLSTATE: HY000 (ER\_CANT\_GET\_WD)

Message: Can't get working directory (errno: %d)

- Error: 1015 SQLSTATE: HY000 (ER\_CANT\_LOCK)

Message: Can't lock file (errno: %d)

- Error: 1016 SQLSTATE: HY000 (ER\_CANT\_OPEN\_FILE)

Message: Can't open file: '%s' (errno: %d)

- Error: 1017 SQLSTATE: HY000 (ER\_FILE\_NOT\_FOUND)



- Message: Can't find file: '%s' (errno: %d)

  - Error: 1018 SQLSTATE: HY000 (ER\_CANT\_READ\_DIR)  
Message: Can't read dir of '%s' (errno: %d)
  - Error: 1019 SQLSTATE: HY000 (ER\_CANT\_SET\_WD)  
Message: Can't change dir to '%s' (errno: %d)
  - Error: 1020 SQLSTATE: HY000 (ER\_CHECKREAD)  
Message: Record has changed since last read in table '%s'
  - Error: 1021 SQLSTATE: HY000 (ER\_DISK\_FULL)  
Message: Disk full (%s); waiting for someone to free some space...
  - Error: 1022 SQLSTATE: 23000 (ER\_DUP\_KEY)  
Message: Can't write; duplicate key in table '%s'
  - Error: 1023 SQLSTATE: HY000 (ER\_ERROR\_ON\_CLOSE)  
Message: Error on close of '%s' (errno: %d)
  - Error: 1024 SQLSTATE: HY000 (ER\_ERROR\_ON\_READ)  
Message: Error reading file '%s' (errno: %d)
  - Error: 1025 SQLSTATE: HY000 (ER\_ERROR\_ON\_RENAME)  
Message: Error on rename of '%s' to '%s' (errno: %d)
  - Error: 1026 SQLSTATE: HY000 (ER\_ERROR\_ON\_WRITE)  
Message: Error writing file '%s' (errno: %d)
  - Error: 1027 SQLSTATE: HY000 (ER\_FILE\_USED)  
Message: '%s' is locked against change
  - Error: 1028 SQLSTATE: HY000 (ER\_FILSORT\_ABORT)  
Message: Sort aborted
  - Error: 1029 SQLSTATE: HY000 (ER\_FORM\_NOT\_FOUND)  
Message: View '%s' doesn't exist for '%s'
  - Error: 1030 SQLSTATE: HY000 (ER\_GET\_ERRNO)  
Message: Got error %d from storage engine
  - Error: 1031 SQLSTATE: HY000 (ER\_ILLEGAL\_HA)  
Message: Table storage engine for '%s' doesn't have this option
  - Error: 1032 SQLSTATE: HY000 (ER\_KEY\_NOT\_FOUND)

Message: Can't find record in '%s'

- Error: 1033 SQLSTATE: HY000 (ER\_NOT\_FORM\_FILE)

Message: Incorrect information in file: '%s'

- Error: 1034 SQLSTATE: HY000 (ER\_NOT\_KEYFILE)

Message: Incorrect key file for table '%s'; try to repair it

- Error: 1035 SQLSTATE: HY000 (ER\_OLD\_KEYFILE)

Message: Old key file for table '%s'; repair it!

- Error: 1036 SQLSTATE: HY000 (ER\_OPEN\_AS\_READONLY)

Message: Table '%s' is read only

- Error: 1037 SQLSTATE: HY001 (ER\_OUTOFMEMORY)

Message: Out of memory; restart server and try again (needed %d bytes)

- Error: 1038 SQLSTATE: HY001 (ER\_OUT\_OF\_SORTMEMORY)

Message: Out of sort memory; increase server sort buffer size

- Error: 1039 SQLSTATE: HY000 (ER\_UNEXPECTED\_EOF)

Message: Unexpected EOF found when reading file '%s' (errno: %d)

- Error: 1040 SQLSTATE: 08004 (ER\_CON\_COUNT\_ERROR)

Message: Too many connections

- Error: 1041 SQLSTATE: HY000 (ER\_OUT\_OF\_RESOURCES)

Message: Out of memory; check if mysqld or some other process uses all available memory; if not, you may have to use 'ulimit' to allow mysqld to use more memory or you can add more swap space

- Error: 1042 SQLSTATE: 08S01 (ER\_BAD\_HOST\_ERROR)

Message: Can't get hostname for your address

- Error: 1043 SQLSTATE: 08S01 (ER\_HANDSHAKE\_ERROR)

Message: Bad handshake

- Error: 1044 SQLSTATE: 42000 (ER\_DBACCESS\_DENIED\_ERROR)

Message: Access denied for user '%s'@'%s' to database '%s'

- Error: 1045 SQLSTATE: 28000 (ER\_ACCESS\_DENIED\_ERROR)

Message: Access denied for user '%s'@'%s' (using password: %s)

- Error: 1046 SQLSTATE: 3D000 (ER\_NO\_DB\_ERROR)

Message: No database selected

- Error: 1047 SQLSTATE: 08S01 (ER\_UNKNOWN\_COM\_ERROR)  
Message: Unknown command
- Error: 1048 SQLSTATE: 23000 (ER\_BAD\_NULL\_ERROR)  
Message: Column '%s' cannot be null
- Error: 1049 SQLSTATE: 42000 (ER\_BAD\_DB\_ERROR)  
Message: Unknown database '%s'
- Error: 1050 SQLSTATE: 42S01 (ER\_TABLE\_EXISTS\_ERROR)  
Message: Table '%s' already exists
- Error: 1051 SQLSTATE: 42S02 (ER\_BAD\_TABLE\_ERROR)  
Message: Unknown table '%s'
- Error: 1052 SQLSTATE: 23000 (ER\_NON\_UNIQ\_ERROR)  
Message: Column '%s' in %s is ambiguous
- Error: 1053 SQLSTATE: 08S01 (ER\_SERVER\_SHUTDOWN)  
Message: Server shutdown in progress
- Error: 1054 SQLSTATE: 42S22 (ER\_BAD\_FIELD\_ERROR)  
Message: Unknown column '%s' in '%s'
- Error: 1055 SQLSTATE: 42000 (ER\_WRONG\_FIELD\_WITH\_GROUP)  
Message: '%s' isn't in GROUP BY
- Error: 1056 SQLSTATE: 42000 (ER\_WRONG\_GROUP\_FIELD)  
Message: Can't group on '%s'
- Error: 1057 SQLSTATE: 42000 (ER\_WRONG\_SUM\_SELECT)  
Message: Statement has sum functions and columns in same statement
- Error: 1058 SQLSTATE: 21S01 (ER\_WRONG\_VALUE\_COUNT)  
Message: Column count doesn't match value count
- Error: 1059 SQLSTATE: 42000 (ER\_TOO\_LONG\_IDENT)  
Message: Identifier name '%s' is too long
- Error: 1060 SQLSTATE: 42S21 (ER\_DUP\_FIELDNAME)  
Message: Duplicate column name '%s'
- Error: 1061 SQLSTATE: 42000 (ER\_DUP\_KEYNAME)  
Message: Duplicate key name '%s'

- Error: 1062 SQLSTATE: 23000 (ER\_DUP\_ENTRY)  
Message: Duplicate entry '%s' for key %d
- Error: 1063 SQLSTATE: 42000 (ER\_WRONG\_FIELD\_SPEC)  
Message: Incorrect column specifier for column '%s'
- Error: 1064 SQLSTATE: 42000 (ER\_PARSE\_ERROR)  
Message: %s near '%s' at line %d
- Error: 1065 SQLSTATE: HY000 (ER\_EMPTY\_QUERY)  
Message: Query was empty
- Error: 1066 SQLSTATE: 42000 (ER\_NONUNIQ\_TABLE)  
Message: Not unique table/alias: '%s'
- Error: 1067 SQLSTATE: 42000 (ER\_INVALID\_DEFAULT)  
Message: Invalid default value for '%s'
- Error: 1068 SQLSTATE: 42000 (ER\_MULTIPLE\_PRI\_KEY)  
Message: Multiple primary key defined
- Error: 1069 SQLSTATE: 42000 (ER\_TOO\_MANY\_KEYS)  
Message: Too many keys specified; max %d keys allowed
- Error: 1070 SQLSTATE: 42000 (ER\_TOO\_MANY\_KEY\_PARTS)  
Message: Too many key parts specified; max %d parts allowed
- Error: 1071 SQLSTATE: 42000 (ER\_TOO\_LONG\_KEY)  
Message: Specified key was too long; max key length is %d bytes
- Error: 1072 SQLSTATE: 42000 (ER\_KEY\_COLUMN\_DOES\_NOT\_EXISTS)  
Message: Key column '%s' doesn't exist in table
- Error: 1073 SQLSTATE: 42000 (ER\_BLOB\_USED\_AS\_KEY)  
Message: BLOB column '%s' can't be used in key specification with the used table type
- Error: 1074 SQLSTATE: 42000 (ER\_TOO\_BIG\_FIELDLENGTH)  
Message: Column length too big for column '%s' (max = %d); use BLOB or TEXT instead
- Error: 1075 SQLSTATE: 42000 (ER\_WRONG\_AUTO\_KEY)  
Message: Incorrect table definition; there can be only one auto column and it must be defined as a key
- Error: 1076 SQLSTATE: HY000 (ER\_READY)  
Message: %s: ready for connections. Version: '%s' socket: '%s' port: %d

- Error: 1077 SQLSTATE: HY000 (ER\_NORMAL\_SHUTDOWN)  
Message: %s: Normal shutdown
- Error: 1078 SQLSTATE: HY000 (ER\_GOT\_SIGNAL)  
Message: %s: Got signal %d. Aborting!
- Error: 1079 SQLSTATE: HY000 (ER\_SHUTDOWN\_COMPLETE)  
Message: %s: Shutdown complete
- Error: 1080 SQLSTATE: 08S01 (ER\_FORCING\_CLOSE)  
Message: %s: Forcing close of thread %ld user: '%s'
- Error: 1081 SQLSTATE: 08S01 (ER\_IPSOCKET\_ERROR)  
Message: Can't create IP socket
- Error: 1082 SQLSTATE: 42S12 (ER\_NO\_SUCH\_INDEX)  
Message: Table '%s' has no index like the one used in CREATE INDEX; recreate the table
- Error: 1083 SQLSTATE: 42000 (ER\_WRONG\_FIELD\_TERMINATORS)  
Message: Field separator argument is not what is expected; check the manual
- Error: 1084 SQLSTATE: 42000 (ER\_BLOBS\_AND\_NO\_TERMINATED)  
Message: You can't use fixed rowlength with BLOBs; please use 'fields terminated by'
- Error: 1085 SQLSTATE: HY000 (ER\_TEXTFILE\_NOT\_READABLE)  
Message: The file '%s' must be in the database directory or be readable by all
- Error: 1086 SQLSTATE: HY000 (ER\_FILE\_EXISTS\_ERROR)  
Message: File '%s' already exists
- Error: 1087 SQLSTATE: HY000 (ER\_LOAD\_INFO)  
Message: Records: %ld Deleted: %ld Skipped: %ld Warnings: %ld
- Error: 1088 SQLSTATE: HY000 (ER\_ALTER\_INFO)  
Message: Records: %ld Duplicates: %ld
- Error: 1089 SQLSTATE: HY000 (ER\_WRONG\_SUB\_KEY)  
Message: Incorrect sub part key; the used key part isn't a string, the used length is longer than the key part, or the storage engine doesn't support unique sub keys
- Error: 1090 SQLSTATE: 42000 (ER\_CANT\_REMOVE\_ALL\_FIELDS)  
Message: You can't delete all columns with ALTER TABLE; use DROP TABLE instead
- Error: 1091 SQLSTATE: 42000 (ER\_CANT\_DROP\_FIELD\_OR\_KEY)  
Message: Can't DROP '%s'; check that column/key exists

- Error: 1092 SQLSTATE: HY000 (ER\_INSERT\_INFO)  
Message: Records: %ld Duplicates: %ld Warnings: %ld
- Error: 1093 SQLSTATE: HY000 (ER\_UPDATE\_TABLE\_USED)  
Message: You can't specify target table '%s' for update in FROM clause
- Error: 1094 SQLSTATE: HY000 (ER\_NO\_SUCH\_THREAD)  
Message: Unknown thread id: %lu
- Error: 1095 SQLSTATE: HY000 (ER\_KILL\_DENIED\_ERROR)  
Message: You are not owner of thread %lu
- Error: 1096 SQLSTATE: HY000 (ER\_NO\_TABLES\_USED)  
Message: No tables used
- Error: 1097 SQLSTATE: HY000 (ER\_TOO\_BIG\_SET)  
Message: Too many strings for column %s and SET
- Error: 1098 SQLSTATE: HY000 (ER\_NO\_UNIQUE\_LOGFILE)  
Message: Can't generate a unique log-filename %s.(1-999)
- Error: 1099 SQLSTATE: HY000 (ER\_TABLE\_NOT\_LOCKED\_FOR\_WRITE)  
Message: Table '%s' was locked with a READ lock and can't be updated
- Error: 1100 SQLSTATE: HY000 (ER\_TABLE\_NOT\_LOCKED)  
Message: Table '%s' was not locked with LOCK TABLES
- Error: 1101 SQLSTATE: 42000 (ER\_BLOB\_CANT\_HAVE\_DEFAULT)  
Message: BLOB/TEXT column '%s' can't have a default value
- Error: 1102 SQLSTATE: 42000 (ER\_WRONG\_DB\_NAME)  
Message: Incorrect database name '%s'
- Error: 1103 SQLSTATE: 42000 (ER\_WRONG\_TABLE\_NAME)  
Message: Incorrect table name '%s'
- Error: 1104 SQLSTATE: 42000 (ER\_TOO\_BIG\_SELECT)  
Message: The SELECT would examine more than MAX\_JOIN\_SIZE rows; check your WHERE and use SET SQL\_BIG\_SELECTS=1 or SET SQL\_MAX\_JOIN\_SIZE=# if the SELECT is okay
- Error: 1105 SQLSTATE: HY000 (ER\_UNKNOWN\_ERROR)  
Message: Unknown error
- Error: 1106 SQLSTATE: 42000 (ER\_UNKNOWN\_PROCEDURE)  
Message: Unknown procedure '%s'

- Error: [1107 SQLSTATE: 42000 \(ER\\_WRONG\\_PARAMCOUNT\\_TO\\_PROCEDURE\)](#)  
Message: Incorrect parameter count to procedure '%s'
- Error: [1108 SQLSTATE: HY000 \(ER\\_WRONG\\_PARAMETERS\\_TO\\_PROCEDURE\)](#)  
Message: Incorrect parameters to procedure '%s'
- Error: [1109 SQLSTATE: 42S02 \(ER\\_UNKNOWN\\_TABLE\)](#)  
Message: Unknown table '%s' in %s
- Error: [1110 SQLSTATE: 42000 \(ER\\_FIELD\\_SPECIFIED\\_TWICE\)](#)  
Message: Column '%s' specified twice
- Error: [1111 SQLSTATE: HY000 \(ER\\_INVALID\\_GROUP\\_FUNC\\_USE\)](#)  
Message: Invalid use of group function
- Error: [1112 SQLSTATE: 42000 \(ER\\_UNSUPPORTED\\_EXTENSION\)](#)  
Message: Table '%s' uses an extension that doesn't exist in this MySQL version
- Error: [1113 SQLSTATE: 42000 \(ER\\_TABLE\\_MUST\\_HAVE\\_COLUMNS\)](#)  
Message: A table must have at least 1 column
- Error: [1114 SQLSTATE: HY000 \(ER\\_RECORD\\_FILE\\_FULL\)](#)  
Message: The table '%s' is full
- Error: [1115 SQLSTATE: 42000 \(ER\\_UNKNOWN\\_CHARACTER\\_SET\)](#)  
Message: Unknown character set: '%s'
- Error: [1116 SQLSTATE: HY000 \(ER\\_TOO\\_MANY\\_TABLES\)](#)  
Message: Too many tables; MySQL can only use %d tables in a join
- Error: [1117 SQLSTATE: HY000 \(ER\\_TOO\\_MANY\\_FIELDS\)](#)  
Message: Too many columns
- Error: [1118 SQLSTATE: 42000 \(ER\\_TOO\\_BIG\\_ROWSIZE\)](#)  
Message: Row size too large. The maximum row size for the used table type, not counting BLOBs, is %ld. You have to change some columns to TEXT or BLOBs
- Error: [1119 SQLSTATE: HY000 \(ER\\_STACK\\_OVERRUN\)](#)  
Message: Thread stack overrun: Used: %ld of a %ld stack. Use 'mysqld -O thread\_stack=#' to specify a bigger stack if needed
- Error: [1120 SQLSTATE: 42000 \(ER\\_WRONG\\_OUTER\\_JOIN\)](#)  
Message: Cross dependency found in OUTER JOIN; examine your ON conditions
- Error: [1121 SQLSTATE: 42000 \(ER\\_NULL\\_COLUMN\\_IN\\_INDEX\)](#)

Message: Column '%s' is used with UNIQUE or INDEX but is not defined as NOT NULL

- Error: 1122 SQLSTATE: HY000 (ER\_CANT\_FIND\_UDF)

Message: Can't load function '%s'

- Error: 1123 SQLSTATE: HY000 (ER\_CANT\_INITIALIZE\_UDF)

Message: Can't initialize function '%s'; %s

- Error: 1124 SQLSTATE: HY000 (ER\_UDF\_NO\_PATHS)

Message: No paths allowed for shared library

- Error: 1125 SQLSTATE: HY000 (ER\_UDF\_EXISTS)

Message: Function '%s' already exists

- Error: 1126 SQLSTATE: HY000 (ER\_CANT\_OPEN\_LIBRARY)

Message: Can't open shared library '%s' (errno: %d %s)

- Error: 1127 SQLSTATE: HY000 (ER\_CANT\_FIND\_DL\_ENTRY)

Message: Can't find function '%s' in library

- Error: 1128 SQLSTATE: HY000 (ER\_FUNCTION\_NOT\_DEFINED)

Message: Function '%s' is not defined

- Error: 1129 SQLSTATE: HY000 (ER\_HOST\_IS\_BLOCKED)

Message: Host '%s' is blocked because of many connection errors; unblock with 'mysqladmin flush-hosts'

- Error: 1130 SQLSTATE: HY000 (ER\_HOST\_NOT\_PRIVILEGED)

Message: Host '%s' is not allowed to connect to this MySQL server

- Error: 1131 SQLSTATE: 42000 (ER\_PASSWORD\_ANONYMOUS\_USER)

Message: You are using MySQL as an anonymous user and anonymous users are not allowed to change passwords

- Error: 1132 SQLSTATE: 42000 (ER\_PASSWORD\_NOT\_ALLOWED)

Message: You must have privileges to update tables in the mysql database to be able to change passwords for others

- Error: 1133 SQLSTATE: 42000 (ER\_PASSWORD\_NO\_MATCH)

Message: Can't find any matching row in the user table

- Error: 1134 SQLSTATE: HY000 (ER\_UPDATE\_INFO)

Message: Rows matched: %ld Changed: %ld Warnings: %ld

- Error: 1135 SQLSTATE: HY000 (ER\_CANT\_CREATE\_THREAD)



Message: Can't create a new thread (errno %d); if you are not out of available memory, you can consult the manual for a possible OS-dependent bug

- Error: 1136 SQLSTATE: 21S01 (ER\_WRONG\_VALUE\_COUNT\_ON\_ROW)

Message: Column count doesn't match value count at row %ld

- Error: 1137 SQLSTATE: HY000 (ER\_CANT\_REOPEN\_TABLE)

Message: Can't reopen table: '%s'

- Error: 1138 SQLSTATE: 42000 (ER\_INVALID\_USE\_OF\_NULL)

Message: Invalid use of NULL value

- Error: 1139 SQLSTATE: 42000 (ER\_REGEXP\_ERROR)

Message: Got error '%s' from regexp

- Error: 1140 SQLSTATE: 42000 (ER\_MIX\_OF\_GROUP\_FUNC\_AND\_FIELDS)

Message: Mixing of GROUP columns (MIN(),MAX(),COUNT(),...) with no GROUP columns is illegal if there is no GROUP BY clause

- Error: 1141 SQLSTATE: 42000 (ER\_NONEXISTING\_GRANT)

Message: There is no such grant defined for user '%s' on host '%s'

- Error: 1142 SQLSTATE: 42000 (ER\_TABLEACCESS\_DENIED\_ERROR)

Message: %s command denied to user '%s'@'%s' for table '%s'

- Error: 1143 SQLSTATE: 42000 (ER\_COLUMNACCESS\_DENIED\_ERROR)

Message: %s command denied to user '%s'@'%s' for column '%s' in table '%s'

- Error: 1144 SQLSTATE: 42000 (ER\_ILLEGAL\_GRANT\_FOR\_TABLE)

Message: Illegal GRANT/REVOKE command; please consult the manual to see which privileges can be used

- Error: 1145 SQLSTATE: 42000 (ER\_GRANT\_WRONG\_HOST\_OR\_USER)

Message: The host or user argument to GRANT is too long

- Error: 1146 SQLSTATE: 42S02 (ER\_NO\_SUCH\_TABLE)

Message: Table '%s.%s' doesn't exist

- Error: 1147 SQLSTATE: 42000 (ER\_NONEXISTING\_TABLE\_GRANT)

Message: There is no such grant defined for user '%s' on host '%s' on table '%s'

- Error: 1148 SQLSTATE: 42000 (ER\_NOT\_ALLOWED\_COMMAND)

Message: The used command is not allowed with this MySQL version

- Error: 1149 SQLSTATE: 42000 (ER\_SYNTAX\_ERROR)

Message: You have an error in your SQL syntax; check the manual that corresponds to your MySQL server version for the right syntax to use

- Error: 1150 SQLSTATE: HY000 (ER\_DELAYED\_CANT\_CHANGE\_LOCK)

Message: Delayed insert thread couldn't get requested lock for table %s

- Error: 1151 SQLSTATE: HY000 (ER\_TOO\_MANY\_DELAYED\_THREADS)

Message: Too many delayed threads in use

- Error: 1152 SQLSTATE: 08S01 (ER\_ABORTING\_CONNECTION)

Message: Aborted connection %ld to db: '%s' user: '%s' (%s)

- Error: 1153 SQLSTATE: 08S01 (ER\_NET\_PACKET\_TOO\_LARGE)

Message: Got a packet bigger than 'max\_allowed\_packet' bytes

- Error: 1154 SQLSTATE: 08S01 (ER\_NET\_READ\_ERROR\_FROM\_PIPE)

Message: Got a read error from the connection pipe

- Error: 1155 SQLSTATE: 08S01 (ER\_NET\_FCNTL\_ERROR)

Message: Got an error from fcntl()

- Error: 1156 SQLSTATE: 08S01 (ER\_NET\_PACKETS\_OUT\_OF\_ORDER)

Message: Got packets out of order

- Error: 1157 SQLSTATE: 08S01 (ER\_NET\_UNCOMPRESS\_ERROR)

Message: Couldn't uncompress communication packet

- Error: 1158 SQLSTATE: 08S01 (ER\_NET\_READ\_ERROR)

Message: Got an error reading communication packets

- Error: 1159 SQLSTATE: 08S01 (ER\_NET\_READ\_INTERRUPTED)

Message: Got timeout reading communication packets

- Error: 1160 SQLSTATE: 08S01 (ER\_NET\_ERROR\_ON\_WRITE)

Message: Got an error writing communication packets

- Error: 1161 SQLSTATE: 08S01 (ER\_NET\_WRITE\_INTERRUPTED)

Message: Got timeout writing communication packets

- Error: 1162 SQLSTATE: 42000 (ER\_TOO\_LONG\_STRING)

Message: Result string is longer than 'max\_allowed\_packet' bytes

- Error: 1163 SQLSTATE: 42000 (ER\_TABLE\_CANT\_HANDLE\_BLOB)

Message: The used table type doesn't support BLOB/TEXT columns

- Error: [1164 SQLSTATE: 42000 \(ER\\_TABLE\\_CANT\\_HANDLE\\_AUTO\\_INCREMENT\)](#)  
Message: The used table type doesn't support AUTO\_INCREMENT columns
- Error: [1165 SQLSTATE: HY000 \(ER\\_DELAYED\\_INSERT\\_TABLE\\_LOCKED\)](#)  
Message: INSERT DELAYED can't be used with table '%s' because it is locked with LOCK TABLES
- Error: [1166 SQLSTATE: 42000 \(ER\\_WRONG\\_COLUMN\\_NAME\)](#)  
Message: Incorrect column name '%s'
- Error: [1167 SQLSTATE: 42000 \(ER\\_WRONG\\_KEY\\_COLUMN\)](#)  
Message: The used storage engine can't index column '%s'
- Error: [1168 SQLSTATE: HY000 \(ER\\_WRONG\\_MRG\\_TABLE\)](#)  
Message: All tables in the MERGE table are not identically defined
- Error: [1169 SQLSTATE: 23000 \(ER\\_DUP\\_UNIQUE\)](#)  
Message: Can't write, because of unique constraint, to table '%s'
- Error: [1170 SQLSTATE: 42000 \(ER\\_BLOB\\_KEY\\_WITHOUT\\_LENGTH\)](#)  
Message: BLOB/TEXT column '%s' used in key specification without a key length
- Error: [1171 SQLSTATE: 42000 \(ER\\_PRIMARY\\_CANT\\_HAVE\\_NULL\)](#)  
Message: All parts of a PRIMARY KEY must be NOT NULL; if you need NULL in a key, use UNIQUE instead
- Error: [1172 SQLSTATE: 42000 \(ER\\_TOO\\_MANY\\_ROWS\)](#)  
Message: Result consisted of more than one row
- Error: [1173 SQLSTATE: 42000 \(ER\\_REQUIRES\\_PRIMARY\\_KEY\)](#)  
Message: This table type requires a primary key
- Error: [1174 SQLSTATE: HY000 \(ER\\_NO\\_RAID\\_COMPILED\)](#)  
Message: This version of MySQL is not compiled with RAID support
- Error: [1175 SQLSTATE: HY000 \(ER\\_UPDATE\\_WITHOUT\\_KEY\\_IN\\_SAFE\\_MODE\)](#)  
Message: You are using safe update mode and you tried to update a table without a WHERE that uses a KEY column
- Error: [1176 SQLSTATE: HY000 \(ER\\_KEY\\_DOES\\_NOT\\_EXISTS\)](#)  
Message: Key '%s' doesn't exist in table '%s'
- Error: [1177 SQLSTATE: 42000 \(ER\\_CHECK\\_NO\\_SUCH\\_TABLE\)](#)  
Message: Can't open table
- Error: [1178 SQLSTATE: 42000 \(ER\\_CHECK\\_NOT\\_IMPLEMENTED\)](#)

Message: The storage engine for the table doesn't support %s

- Error: 1179 SQLSTATE: 25000 (ER\_CANT\_DO\_THIS\_DURING\_AN\_TRANSACTION)

Message: You are not allowed to execute this command in a transaction

- Error: 1180 SQLSTATE: HY000 (ER\_ERROR\_DURING\_COMMIT)

Message: Got error %d during COMMIT

- Error: 1181 SQLSTATE: HY000 (ER\_ERROR\_DURING\_ROLLBACK)

Message: Got error %d during ROLLBACK

- Error: 1182 SQLSTATE: HY000 (ER\_ERROR\_DURING\_FLUSH\_LOGS)

Message: Got error %d during FLUSH\_LOGS

- Error: 1183 SQLSTATE: HY000 (ER\_ERROR\_DURING\_CHECKPOINT)

Message: Got error %d during CHECKPOINT

- Error: 1184 SQLSTATE: 08S01 (ER\_NEW\_ABORTING\_CONNECTION)

Message: Aborted connection %ld to db: '%s' user: '%s' host: '%s' (%s)

- Error: 1185 SQLSTATE: HY000 (ER\_DUMP\_NOT\_IMPLEMENTED)

Message: The storage engine for the table does not support binary table dump

- Error: 1186 SQLSTATE: HY000 (ER\_FLUSH\_MASTER\_BINLOG\_CLOSED)

Message: Binlog closed, cannot RESET MASTER

- Error: 1187 SQLSTATE: HY000 (ER\_INDEX\_REBUILD)

Message: Failed rebuilding the index of dumped table '%s'

- Error: 1188 SQLSTATE: HY000 (ER\_MASTER)

Message: Error from master: '%s'

- Error: 1189 SQLSTATE: 08S01 (ER\_MASTER\_NET\_READ)

Message: Net error reading from master

- Error: 1190 SQLSTATE: 08S01 (ER\_MASTER\_NET\_WRITE)

Message: Net error writing to master

- Error: 1191 SQLSTATE: HY000 (ER\_FT\_MATCHING\_KEY\_NOT\_FOUND)

Message: Can't find FULLTEXT index matching the column list

- Error: 1192 SQLSTATE: HY000 (ER\_LOCK\_OR\_ACTIVE\_TRANSACTION)

Message: Can't execute the given command because you have active locked tables or an active transaction

- Error: 1193 SQLSTATE: HY000 (ER\_UNKNOWN\_SYSTEM\_VARIABLE)  
Message: Unknown system variable '%s'
- Error: 1194 SQLSTATE: HY000 (ER\_CRASHED\_ON\_USAGE)  
Message: Table '%s' is marked as crashed and should be repaired
- Error: 1195 SQLSTATE: HY000 (ER\_CRASHED\_ON\_REPAIR)  
Message: Table '%s' is marked as crashed and last (automatic?) repair failed
- Error: 1196 SQLSTATE: HY000 (ER\_WARNING\_NOT\_COMPLETE\_ROLLBACK)  
Message: Some non-transactional changed tables couldn't be rolled back
- Error: 1197 SQLSTATE: HY000 (ER\_TRANS\_CACHE\_FULL)  
Message: Multi-statement transaction required more than 'max\_binlog\_cache\_size' bytes of storage; increase this mysqld variable and try again
- Error: 1198 SQLSTATE: HY000 (ER\_SLAVE\_MUST\_STOP)  
Message: This operation cannot be performed with a running slave; run STOP SLAVE first
- Error: 1199 SQLSTATE: HY000 (ER\_SLAVE\_NOT\_RUNNING)  
Message: This operation requires a running slave; configure slave and do START SLAVE
- Error: 1200 SQLSTATE: HY000 (ER\_BAD\_SLAVE)  
Message: The server is not configured as slave; fix in config file or with CHANGE MASTER TO
- Error: 1201 SQLSTATE: HY000 (ER\_MASTER\_INFO)  
Message: Could not initialize master info structure; more error messages can be found in the MySQL error log
- Error: 1202 SQLSTATE: HY000 (ER\_SLAVE\_THREAD)  
Message: Could not create slave thread; check system resources
- Error: 1203 SQLSTATE: 42000 (ER\_TOO\_MANY\_USER\_CONNECTIONS)  
Message: User %s has already more than 'max\_user\_connections' active connections
- Error: 1204 SQLSTATE: HY000 (ER\_SET\_CONSTANTS\_ONLY)  
Message: You may only use constant expressions with SET
- Error: 1205 SQLSTATE: HY000 (ER\_LOCK\_WAIT\_TIMEOUT)  
Message: Lock wait timeout exceeded; try restarting transaction
- Error: 1206 SQLSTATE: HY000 (ER\_LOCK\_TABLE\_FULL)  
Message: The total number of locks exceeds the lock table size
- Error: 1207 SQLSTATE: 25000 (ER\_READ\_ONLY\_TRANSACTION)

Message: Update locks cannot be acquired during a READ UNCOMMITTED transaction

- Error: 1208 SQLSTATE: HY000 (ER\_DROP\_DB\_WITH\_READ\_LOCK)

Message: DROP DATABASE not allowed while thread is holding global read lock

- Error: 1209 SQLSTATE: HY000 (ER\_CREATE\_DB\_WITH\_READ\_LOCK)

Message: CREATE DATABASE not allowed while thread is holding global read lock

- Error: 1210 SQLSTATE: HY000 (ER\_WRONG\_ARGUMENTS)

Message: Incorrect arguments to %s

- Error: 1211 SQLSTATE: 42000 (ER\_NO\_PERMISSION\_TO\_CREATE\_USER)

Message: '%s'@'%s' is not allowed to create new users

- Error: 1212 SQLSTATE: HY000 (ER\_UNION\_TABLES\_IN\_DIFFERENT\_DIR)

Message: Incorrect table definition; all MERGE tables must be in the same database

- Error: 1213 SQLSTATE: 40001 (ER\_LOCK\_DEADLOCK)

Message: Deadlock found when trying to get lock; try restarting transaction

- Error: 1214 SQLSTATE: HY000 (ER\_TABLE\_CANT\_HANDLE\_FT)

Message: The used table type doesn't support FULLTEXT indexes

- Error: 1215 SQLSTATE: HY000 (ER\_CANNOT\_ADD\_FOREIGN)

Message: Cannot add foreign key constraint

- Error: 1216 SQLSTATE: 23000 (ER\_NO\_REFERENCED\_ROW)

Message: Cannot add or update a child row: a foreign key constraint fails

- Error: 1217 SQLSTATE: 23000 (ER\_ROW\_IS\_REFERENCED)

Message: Cannot delete or update a parent row: a foreign key constraint fails

- Error: 1218 SQLSTATE: 08S01 (ER\_CONNECT\_TO\_MASTER)

Message: Error connecting to master: %s

- Error: 1219 SQLSTATE: HY000 (ER\_QUERY\_ON\_MASTER)

Message: Error running query on master: %s

- Error: 1220 SQLSTATE: HY000 (ER\_ERROR\_WHEN\_EXECUTING\_COMMAND)

Message: Error when executing command %s: %s

- Error: 1221 SQLSTATE: HY000 (ER\_WRONG\_USAGE)

Message: Incorrect usage of %s and %s

- Error: 1222 SQLSTATE: 21000 (ER\_WRONG\_NUMBER\_OF\_COLUMNS\_IN\_SELECT)

Message: The used SELECT statements have a different number of columns

- Error: 1223 SQLSTATE: HY000 (ER\_CANT\_UPDATE\_WITH\_READLOCK)

Message: Can't execute the query because you have a conflicting read lock

- Error: 1224 SQLSTATE: HY000 (ER\_MIXING\_NOT\_ALLOWED)

Message: Mixing of transactional and non-transactional tables is disabled

- Error: 1225 SQLSTATE: HY000 (ER\_DUP\_ARGUMENT)

Message: Option '%s' used twice in statement

- Error: 1226 SQLSTATE: 42000 (ER\_USER\_LIMIT\_REACHED)

Message: User '%s' has exceeded the '%s' resource (current value: %ld)

- Error: 1227 SQLSTATE: HY000 (ER\_SPECIFIC\_ACCESS\_DENIED\_ERROR)

Message: Access denied; you need the %s privilege for this operation

- Error: 1228 SQLSTATE: HY000 (ER\_LOCAL\_VARIABLE)

Message: Variable '%s' is a SESSION variable and can't be used with SET GLOBAL

- Error: 1229 SQLSTATE: HY000 (ER\_GLOBAL\_VARIABLE)

Message: Variable '%s' is a GLOBAL variable and should be set with SET GLOBAL

- Error: 1230 SQLSTATE: 42000 (ER\_NO\_DEFAULT)

Message: Variable '%s' doesn't have a default value

- Error: 1231 SQLSTATE: 42000 (ER\_WRONG\_VALUE\_FOR\_VAR)

Message: Variable '%s' can't be set to the value of '%s'

- Error: 1232 SQLSTATE: 42000 (ER\_WRONG\_TYPE\_FOR\_VAR)

Message: Incorrect argument type to variable '%s'

- Error: 1233 SQLSTATE: HY000 (ER\_VAR\_CANT\_BE\_READ)

Message: Variable '%s' can only be set, not read

- Error: 1234 SQLSTATE: 42000 (ER\_CANT\_USE\_OPTION\_HERE)

Message: Incorrect usage/placement of '%s'

- Error: 1235 SQLSTATE: 42000 (ER\_NOT\_SUPPORTED\_YET)

Message: This version of MySQL doesn't yet support '%s'

- Error: 1236 SQLSTATE: HY000 (ER\_MASTER\_FATAL\_ERROR\_READING\_BINLOG)

Message: Got fatal error %d: '%s' from master when reading data from binary log

- Error: 1237 SQLSTATE: HY000 (ER\_SLAVE\_IGNORED\_TABLE)

- Message: Slave SQL thread ignored the query because of replicate-*\**-table rules
- Error: 1238 SQLSTATE: HY000 (ER\_INCORRECT\_GLOBAL\_LOCAL\_VAR)  
Message: Variable '%s' is a %s variable
  - Error: 1239 SQLSTATE: 42000 (ER\_WRONG\_FK\_DEF)  
Message: Incorrect foreign key definition for '%s': %s
  - Error: 1240 SQLSTATE: HY000 (ER\_KEY\_REF\_DO\_NOT\_MATCH\_TABLE\_REF)  
Message: Key reference and table reference don't match
  - Error: 1241 SQLSTATE: 21000 (ER\_OPERAND\_COLUMNS)  
Message: Operand should contain %d column(s)
  - Error: 1242 SQLSTATE: 21000 (ER\_SUBQUERY\_NO\_1\_ROW)  
Message: Subquery returns more than 1 row
  - Error: 1243 SQLSTATE: HY000 (ER\_UNKNOWN\_STMT\_HANDLER)  
Message: Unknown prepared statement handler (%.\*s) given to %s
  - Error: 1244 SQLSTATE: HY000 (ER\_CORRUPT\_HELP\_DB)  
Message: Help database is corrupt or does not exist
  - Error: 1245 SQLSTATE: HY000 (ER\_CYCLIC\_REFERENCE)  
Message: Cyclic reference on subqueries
  - Error: 1246 SQLSTATE: HY000 (ER\_AUTO\_CONVERT)  
Message: Converting column '%s' from %s to %s
  - Error: 1247 SQLSTATE: 42S22 (ER\_ILLEGAL\_REFERENCE)  
Message: Reference '%s' not supported (%s)
  - Error: 1248 SQLSTATE: 42000 (ER\_DERIVED\_MUST\_HAVE\_ALIAS)  
Message: Every derived table must have its own alias
  - Error: 1249 SQLSTATE: 01000 (ER\_SELECT\_REDUCED)  
Message: Select %u was reduced during optimization
  - Error: 1250 SQLSTATE: 42000 (ER\_TABLENAME\_NOT\_ALLOWED\_HERE)  
Message: Table '%s' from one of the SELECTs cannot be used in %s
  - Error: 1251 SQLSTATE: 08004 (ER\_NOT\_SUPPORTED\_AUTH\_MODE)  
Message: Client does not support authentication protocol requested by server; consider upgrading MySQL client



- Error: 1252 SQLSTATE: 42000 (ER\_SPATIAL\_CANT\_HAVE\_NULL)  
Message: All parts of a SPATIAL index must be NOT NULL
- Error: 1253 SQLSTATE: 42000 (ER\_COLLATION\_CHARSET\_MISMATCH)  
Message: COLLATION '%s' is not valid for CHARACTER SET '%s'
- Error: 1254 SQLSTATE: HY000 (ER\_SLAVE\_WAS\_RUNNING)  
Message: Slave is already running
- Error: 1255 SQLSTATE: HY000 (ER\_SLAVE\_WAS\_NOT\_RUNNING)  
Message: Slave has already been stopped
- Error: 1256 SQLSTATE: HY000 (ER\_TOO\_BIG\_FOR\_UNCOMPRESS)  
Message: Uncompressed data size too large; the maximum size is %d (probably, length of uncompressed data was corrupted)
- Error: 1257 SQLSTATE: HY000 (ER\_ZLIB\_Z\_MEM\_ERROR)  
Message: ZLIB: Not enough memory
- Error: 1258 SQLSTATE: HY000 (ER\_ZLIB\_Z\_BUF\_ERROR)  
Message: ZLIB: Not enough room in the output buffer (probably, length of uncompressed data was corrupted)
- Error: 1259 SQLSTATE: HY000 (ER\_ZLIB\_Z\_DATA\_ERROR)  
Message: ZLIB: Input data corrupted
- Error: 1260 SQLSTATE: HY000 (ER\_CUT\_VALUE\_GROUP\_CONCAT)  
Message: %d line(s) were cut by GROUP\_CONCAT()
- Error: 1261 SQLSTATE: 01000 (ER\_WARN\_TOO\_FEW\_RECORDS)  
Message: Row %ld doesn't contain data for all columns
- Error: 1262 SQLSTATE: 01000 (ER\_WARN\_TOO\_MANY\_RECORDS)  
Message: Row %ld was truncated; it contained more data than there were input columns
- Error: 1263 SQLSTATE: 01000 (ER\_WARN\_NULL\_TO\_NOTNULL)  
Message: Data truncated; NULL supplied to NOT NULL column '%s' at row %ld
- Error: 1264 SQLSTATE: 01000 (ER\_WARN\_DATA\_OUT\_OF\_RANGE)  
Message: Data truncated; out of range for column '%s' at row %ld
- Error: 1265 SQLSTATE: 01000 (ER\_WARN\_DATA\_TRUNCATED)  
Message: Data truncated for column '%s' at row %ld
- Error: 1266 SQLSTATE: HY000 (ER\_WARN\_USING\_OTHER\_HANDLER)

Message: Using storage engine %s for table '%s'

- Error: 1267 SQLSTATE: HY000 (ER\_CANT\_AGGREGATE\_2COLLATIONS)

Message: Illegal mix of collations (%s,%s) and (%s,%s) for operation '%s'

- Error: 1268 SQLSTATE: HY000 (ER\_DROP\_USER)

Message: Can't drop one or more of the requested users

- Error: 1269 SQLSTATE: HY000 (ER\_REVOKE\_GRANTS)

Message: Can't revoke all privileges, grant for one or more of the requested users

- Error: 1270 SQLSTATE: HY000 (ER\_CANT\_AGGREGATE\_3COLLATIONS)

Message: Illegal mix of collations (%s,%s), (%s,%s), (%s,%s) for operation '%s'

- Error: 1271 SQLSTATE: HY000 (ER\_CANT\_AGGREGATE\_NCOLLATIONS)

Message: Illegal mix of collations for operation '%s'

- Error: 1272 SQLSTATE: HY000 (ER\_VARIABLE\_IS\_NOT\_STRUCT)

Message: Variable '%s' is not a variable component (can't be used as XXXX.variable\_name)

- Error: 1273 SQLSTATE: HY000 (ER\_UNKNOWN\_COLLATION)

Message: Unknown collation: '%s'

- Error: 1274 SQLSTATE: HY000 (ER\_SLAVE\_IGNORED\_SSL\_PARAMS)

Message: SSL parameters in CHANGE MASTER are ignored because this MySQL slave was compiled without SSL support; they can be used later if MySQL slave with SSL is started

- Error: 1275 SQLSTATE: HY000 (ER\_SERVER\_IS\_IN\_SECURE\_AUTH\_MODE)

Message: Server is running in --secure-auth mode, but '%s'@'%s' has a password in the old format; please change the password to the new format

- Error: 1276 SQLSTATE: HY000 (ER\_WARN\_FIELD\_RESOLVED)

Message: Field or reference '%s%s%s%s%s' of SELECT #d was resolved in SELECT #d

- Error: 1277 SQLSTATE: HY000 (ER\_BAD\_SLAVE\_UNTIL\_COND)

Message: Incorrect parameter or combination of parameters for START SLAVE UNTIL

- Error: 1278 SQLSTATE: HY000 (ER\_MISSING\_SKIP\_SLAVE)

Message: It is recommended to use --skip-slave-start when doing step-by-step replication with START SLAVE UNTIL; otherwise, you will get problems if you get an unexpected slave's mysqld restart

- Error: 1279 SQLSTATE: HY000 (ER\_UNTIL\_COND\_IGNORED)

Message: SQL thread is not to be started so UNTIL options are ignored

- Error: 1280 SQLSTATE: 42000 (ER\_WRONG\_NAME\_FOR\_INDEX)

Message: Incorrect index name '%s'

- Error: 1281 SQLSTATE: 42000 (ER\_WRONG\_NAME\_FOR\_CATALOG)

Message: Incorrect catalog name '%s'

- Error: 1282 SQLSTATE: HY000 (ER\_WARN\_QC\_RESIZE)

Message: Query cache failed to set size %lu; new query cache size is %lu

- Error: 1283 SQLSTATE: HY000 (ER\_BAD\_FT\_COLUMN)

Message: Column '%s' cannot be part of FULLTEXT index

- Error: 1284 SQLSTATE: HY000 (ER\_UNKNOWN\_KEY\_CACHE)

Message: Unknown key cache '%s'

- Error: 1285 SQLSTATE: HY000 (ER\_WARN\_HOSTNAME\_WONT\_WORK)

Message: MySQL is started in --skip-name-resolve mode; you must restart it without this switch for this grant to work

- Error: 1286 SQLSTATE: 42000 (ER\_UNKNOWN\_STORAGE\_ENGINE)

Message: Unknown table engine '%s'

- Error: 1287 SQLSTATE: HY000 (ER\_WARN\_DEPRECATED\_SYNTAX)

Message: '%s' is deprecated; use '%s' instead

- Error: 1288 SQLSTATE: HY000 (ER\_NON\_UPDATABLE\_TABLE)

Message: The target table %s of the %s is not updatable

- Error: 1289 SQLSTATE: HY000 (ER\_FEATURE\_DISABLED)

Message: The '%s' feature is disabled; you need MySQL built with '%s' to have it working

- Error: 1290 SQLSTATE: HY000 (ER\_OPTION\_PREVENTS\_STATEMENT)

Message: The MySQL server is running with the %s option so it cannot execute this statement

- Error: 1291 SQLSTATE: HY000 (ER\_DUPLICATED\_VALUE\_IN\_TYPE)

Message: Column '%s' has duplicated value '%s' in %s

- Error: 1292 SQLSTATE: HY000 (ER\_TRUNCATED\_WRONG\_VALUE)

Message: Truncated incorrect %s value: '%s'

- Error: 1293 SQLSTATE: HY000 (ER\_TOO\_MUCH\_AUTO\_TIMESTAMP\_COLS)

Message: Incorrect table definition; there can be only one TIMESTAMP column with CURRENT\_TIMESTAMP in DEFAULT or ON UPDATE clause

- Error: 1294 SQLSTATE: HY000 (ER\_INVALID\_ON\_UPDATE)

Message: Invalid ON UPDATE clause for '%s' column

- Error: [1295](#) SQLSTATE: [HY000](#) ([ER\\_UNSUPPORTED\\_PS](#))  
Message: This command is not supported in the prepared statement protocol yet
- Error: [1296](#) SQLSTATE: [HY000](#) ([ER\\_GET\\_ERRMSG](#))  
Message: Got error %d '%s' from %s
- Error: [1297](#) SQLSTATE: [HY000](#) ([ER\\_GET\\_TEMPORARY\\_ERRMSG](#))  
Message: Got temporary error %d '%s' from %s
- Error: [1298](#) SQLSTATE: [HY000](#) ([ER\\_UNKNOWN\\_TIME\\_ZONE](#))  
Message: Unknown or incorrect time zone: '%s'
- Error: [1299](#) SQLSTATE: [HY000](#) ([ER\\_WARN\\_INVALID\\_TIMESTAMP](#))  
Message: Invalid `TIMESTAMP` value in column '%s' at row %ld
- Error: [1300](#) SQLSTATE: [HY000](#) ([ER\\_INVALID\\_CHARACTER\\_STRING](#))  
Message: Invalid %s character string: '%s'
- Error: [1301](#) SQLSTATE: [HY000](#) ([ER\\_WARN\\_ALLOWED\\_PACKET\\_OVERFLOWED](#))  
Message: Result of %s() was larger than `max_allowed_packet` (%ld) - truncated
- Error: [1302](#) SQLSTATE: [HY000](#) ([ER\\_CONFLICTING\\_DECLARATIONS](#))  
Message: Conflicting declarations: '%s%s' and '%s%s'

## B.2. Client Error Codes and Messages

Client error information comes from the following source files:

- The Error values and the symbols in parentheses correspond to definitions in the [include/errormsg.h](#) MySQL source file.
- The Message values correspond to the error messages that are listed in the [libmysql/errmsg.c](#) file. %d and %s represent numbers and strings, respectively, that are substituted into the messages when they are displayed.

Because updates are frequent, it is possible that those files will contain additional error information not listed here.

- Error: [2000](#) ([CR\\_UNKNOWN\\_ERROR](#))  
Message: Unknown MySQL error
- Error: [2001](#) ([CR\\_SOCKET\\_CREATE\\_ERROR](#))  
Message: Can't create UNIX socket (%d)
- Error: [2002](#) ([CR\\_CONNECTION\\_ERROR](#))

- Message: Can't connect to local MySQL server through socket '%s' (%d)
- Error: 2003 (CR\_CONN\_HOST\_ERROR)  
Message: Can't connect to MySQL server on '%s' (%d)
- Error: 2004 (CR\_IPSOCK\_ERROR)  
Message: Can't create TCP/IP socket (%d)
- Error: 2005 (CR\_UNKNOWN\_HOST)  
Message: Unknown MySQL server host '%s' (%d)
- Error: 2006 (CR\_SERVER\_GONE\_ERROR)  
Message: MySQL server has gone away
- Error: 2007 (CR\_VERSION\_ERROR)  
Message: Protocol mismatch; server version = %d, client version = %d
- Error: 2008 (CR\_OUT\_OF\_MEMORY)  
Message: MySQL client ran out of memory
- Error: 2009 (CR\_WRONG\_HOST\_INFO)  
Message: Wrong host info
- Error: 2010 (CR\_LOCALHOST\_CONNECTION)  
Message: Localhost via UNIX socket
- Error: 2011 (CR\_TCP\_CONNECTION)  
Message: %s via TCP/IP
- Error: 2012 (CR\_SERVER\_HANDSHAKE\_ERR)  
Message: Error in server handshake
- Error: 2013 (CR\_SERVER\_LOST)  
Message: Lost connection to MySQL server during query
- Error: 2014 (CR\_COMMANDS\_OUT\_OF\_SYNC)  
Message: Commands out of sync; you can't run this command now
- Error: 2015 (CR\_NAMEDPIPE\_CONNECTION)  
Message: Named pipe: %s
- Error: 2016 (CR\_NAMEDPIPEWAIT\_ERROR)  
Message: Can't wait for named pipe to host: %s pipe: %s (%lu)
- Error: 2017 (CR\_NAMEDPIPEOPEN\_ERROR)

- Message: Can't open named pipe to host: %s pipe: %s (%lu)
- Error: 2018 (CR\_NAMEDPIPESETSTATE\_ERROR)  
Message: Can't set state of named pipe to host: %s pipe: %s (%lu)
- Error: 2019 (CR\_CANT\_READ\_CHARSET)  
Message: Can't initialize character set %s (path: %s)
- Error: 2020 (CR\_NET\_PACKET\_TOO\_LARGE)  
Message: Got packet bigger than 'max\_allowed\_packet' bytes
- Error: 2021 (CR\_EMBEDDED\_CONNECTION)  
Message: Embedded server
- Error: 2022 (CR\_PROBE\_SLAVE\_STATUS)  
Message: Error on SHOW SLAVE STATUS:
- Error: 2023 (CR\_PROBE\_SLAVE\_HOSTS)  
Message: Error on SHOW SLAVE HOSTS:
- Error: 2024 (CR\_PROBE\_SLAVE\_CONNECT)  
Message: Error connecting to slave:
- Error: 2025 (CR\_PROBE\_MASTER\_CONNECT)  
Message: Error connecting to master:
- Error: 2026 (CR\_SSL\_CONNECTION\_ERROR)  
Message: SSL connection error
- Error: 2027 (CR\_MALFORMED\_PACKET)  
Message: Malformed packet
- Error: 2028 (CR\_WRONG\_LICENSE)  
Message: This client library is licensed only for use with MySQL servers having '%s' license
- Error: 2029 (CR\_NULL\_POINTER)  
Message: Invalid use of null pointer
- Error: 2030 (CR\_NO\_PREPARE\_STMT)  
Message: Statement not prepared
- Error: 2031 (CR\_PARAMS\_NOT\_BOUND)  
Message: No data supplied for parameters in prepared statement
- Error: 2032 (CR\_DATA\_TRUNCATED)

Message: Data truncated

- Error: 2033 (CR\_NO\_PARAMETERS\_EXISTS)

Message: No parameters exist in the statement

- Error: 2034 (CR\_INVALID\_PARAMETER\_NO)

Message: Invalid parameter number

- Error: 2035 (CR\_INVALID\_BUFFER\_USE)

Message: Can't send long data for non-string/non-binary data types (parameter: %d)

- Error: 2036 (CR\_UNSUPPORTED\_PARAM\_TYPE)

Message: Using unsupported buffer type: %d (parameter: %d)

- Error: 2037 (CR\_SHARED\_MEMORY\_CONNECTION)

Message: Shared memory: %s

- Error: 2038 (CR\_SHARED\_MEMORY\_CONNECT\_REQUEST\_ERROR)

Message: Can't open shared memory; client could not create request event (%lu)

- Error: 2039 (CR\_SHARED\_MEMORY\_CONNECT\_ANSWER\_ERROR)

Message: Can't open shared memory; no answer event received from server (%lu)

- Error: 2040 (CR\_SHARED\_MEMORY\_CONNECT\_FILE\_MAP\_ERROR)

Message: Can't open shared memory; server could not allocate file mapping (%lu)

- Error: 2041 (CR\_SHARED\_MEMORY\_CONNECT\_MAP\_ERROR)

Message: Can't open shared memory; server could not get pointer to file mapping (%lu)

- Error: 2042 (CR\_SHARED\_MEMORY\_FILE\_MAP\_ERROR)

Message: Can't open shared memory; client could not allocate file mapping (%lu)

- Error: 2043 (CR\_SHARED\_MEMORY\_MAP\_ERROR)

Message: Can't open shared memory; client could not get pointer to file mapping (%lu)

- Error: 2044 (CR\_SHARED\_MEMORY\_EVENT\_ERROR)

Message: Can't open shared memory; client could not create %s event (%lu)

- Error: 2045 (CR\_SHARED\_MEMORY\_CONNECT\_ABANDONED\_ERROR)

Message: Can't open shared memory; no answer from server (%lu)

- Error: 2046 (CR\_SHARED\_MEMORY\_CONNECT\_SET\_ERROR)

Message: Can't open shared memory; cannot send request event to server (%lu)

- Error: 2047 (CR\_CONN\_UNKNOW\_PROTOCOL)

Message: Wrong or unknown protocol

- Error: 2048 (CR\_INVALID\_CONN\_HANDLE)

Message: Invalid connection handle

- Error: 2049 (CR\_SECURE\_AUTH)

Message: Connection using old (pre-4.1.1) authentication protocol refused (client option 'secure\_auth' enabled)

- Error: 2050 (CR\_FETCH\_CANCELED)

Message: Row retrieval was canceled by mysql\_stmt\_close() call

- Error: 2051 (CR\_NO\_DATA)

Message: Attempt to read column without prior row fetch

- Error: 2052 (CR\_NO\_STMT\_METADATA)

Message: Prepared statement contains no metadata



---

# Appendix C. Credits

This appendix lists the developers, contributors, and supporters that have helped to make MySQL what it is today.

## C.1. Developers at MySQL AB

These are the developers that are or have been employed by MySQL AB to work on the [MySQL](#) database software, roughly in the order they started to work with us. Following each developer is a small list of the tasks that the developer is responsible for, or the accomplishments they have made. All developers are involved in support.

- Michael (Monty) Widenius
  - Lead developer and main author of the MySQL server ([mysqld](#)).
  - New functions for the string library.
  - Most of the [mysys](#) library.
  - The [ISAM](#) and [MyISAM](#) libraries (B-tree index file handlers with index compression and different record formats).
  - The [HEAP](#) library. A memory table system with our superior full dynamic hashing. In use since 1981 and published around 1984.
  - The [replace](#) program (take a look at it, it's **COOL!**).
  - Connector/ODBC (MyODBC), the ODBC driver for Windows.
  - Fixing bugs in MIT-pthreads to get it to work for MySQL Server. And also Unireg, a curses-based application tool with many utilities.
  - Porting of [mSQL](#) tools like [msqlperl](#), [DBD/DBI](#), and [DB2mysql](#).
  - Most of [crash-me](#) and the foundation for the MySQL benchmarks.
- David Axmark
  - Initial main writer of the **Reference Manual**, including enhancements to [texi2html](#).
  - Automatic Web site updating from the manual.
  - Initial Autoconf, Automake, and Libtool support.
  - Licensing.
  - Parts of all the text files. (Nowadays only the [README](#) is left. The rest ended up in the manual.)
  - Lots of testing of new features.
  - Our in-house Free Software legal expert.
  - Mailing list maintainer (who never has the time to do it right...).
  - Our original portability code (now more than 10 years old). Nowadays only some parts of [mysys](#) are left.

- Someone for Monty to call in the middle of the night when he just got that new feature to work.
- Chief "Open Sourcerer" (MySQL community relations).
- Jani Tolonen
  - `mysqlimport`
  - A lot of extensions to the command-line clients.
  - `PROCEDURE ANALYSE()`
- Sinisa Milivojevic (now in support)
  - Compression (with `zlib`) in the client/server protocol.
  - Perfect hashing for the lexical analyzer phase.
  - Multi-row `INSERT`
  - `mysqldump -e` option
  - `LOAD DATA LOCAL INFILE`
  - `SQL_CALC_FOUND_ROWS SELECT` option
  - `--max-user-connections=...` option
  - `net_read` and `net_write_timeout`
  - `GRANT/REVOKE` and `SHOW GRANTS FOR`
  - New client/server protocol for 4.0
  - `UNION` in 4.0
  - Multiple-table `DELETE/UPDATE`
  - Subqueries in the `FROM` clause (4.1).
  - User resources management
  - Initial developer of the `MySQL++` C++ API and the `MySQLGUI` client.
- Tonu Samuel (past developer)
  - VIO interface (the foundation for the encrypted client/server protocol).
  - MySQL Filesystem (a way to use MySQL databases as files and directories).
  - The `CASE` expression.
  - The `MD5()` and `COALESCE()` functions.
  - `RAID` support for `MyISAM` tables.
- Sasha Pachev (past developer)
  - Initial implementation of replication (up to version 4.0).

- `SHOW CREATE TABLE`.
- `mysql-bench`
- Matt Wagner
  - MySQL test suite.
  - Webmaster (until 2002).
- Miguel Solorzano (now in support)
  - Win32 development and release builds.
  - Windows NT server code.
  - WinMySQLAdmin
- Timothy Smith (now in support)
  - Dynamic character sets support.
  - configure, RPMs and other parts of the build system.
  - Initial developer of `libmysqld`, the embedded server.
- Sergei Golubchik
  - Full-text search.
  - Added keys to the `MERGE` library.
  - Precision math.
- Jeremy Cole (past developer)
  - Proofreading and editing this fine manual.
  - `ALTER TABLE ... ORDER BY ....`
  - `UPDATE ... ORDER BY ....`
  - `DELETE ... ORDER BY ....`
- Indrek Siitan
  - Designing/programming of our Web interface.
  - Author of our newsletter management system.
- Jorge del Conde (now in support)
  - `MySQLCC (MySQL Control Center)`
  - Win32 development
  - Initial implementation of the Web site portals.
- Venu Anuganti (past developer)

- MyODBC 3.51
- New client/server protocol for 4.1 (for prepared statements).
- Arjen Lentz (also handled community, 2004-2006; now works in Support)
  - Maintainer of the MySQL Reference Manual (2001-2004).
  - Preparing the O'Reilly printed edition of the manual (2002).
- Alexander (Bar) Barkov, Alexey (Holyfoot) Botchkov, and Ramil Kalimullin
  - Spatial data (GIS) and R-Trees implementation for 4.1
  - Unicode and character sets for 4.1; documentation for same
- Oleksandr (Sanja) Byelkin
  - Query cache in 4.0
  - Implementation of subqueries (4.1).
  - Implementation of views (5.0).
- Aleksey (Walrus) Kishkin and Alexey (Ranger) Stroganov
  - Benchmarks design and analysis.
  - Maintenance of the MySQL test suite.
- Zak Greant (past employee)
  - Open Source advocate, MySQL community relations.
- Carsten Pedersen
  - The MySQL Certification program.
- Lenz Grimmer
  - Production (build and release) engineering.
- Peter Zaitsev
  - [SHA1\(\)](#), [AES\\_ENCRYPT\(\)](#) and [AES\\_DECRYPT\(\)](#) functions.
  - Debugging, cleaning up various features.
- Alexander (Salle) Keremidarski
  - Support.
  - Debugging.
- Per-Erik Martin
  - Lead developer for stored procedures (5.0).
- Jim Winstead

- Former lead Web developer.
- Improving server, fixing bugs.
- Mark Matthews
  - Connector/J driver (Java).
- Peter Gulutzan
  - SQL standards compliance.
  - Documentation of existing MySQL code/algorithms.
  - Character set documentation.
- Guilhem Bichot
  - Replication, from [MySQL](#) version 4.0.
  - Fixed handling of exponents for [DECIMAL](#).
  - Author of [mysql\\_tableinfo](#).
  - Backup (in 5.1).
- Antony T. Curtis
  - Porting of the MySQL Database software to OS/2.
- Mikael Ronstrom
  - Much of the initial work on NDB Cluster until 2000. Roughly half the code base at that time. Transaction protocol, node recovery, system restart and restart code and parts of the API functionality.
  - Lead Architect, developer, debugger of NDB Cluster 1994-2004
  - Lots of optimizations
- Jonas Oreland
  - On-line Backup
  - The automatic test environment of MySQL Cluster
  - Portability Library for NDB Cluster
  - Lots of other things
- Pekka Nouisiainen
  - Ordered index implementation of MySQL Cluster
  - BLOB support in MySQL Cluster
  - Charset support in MySQL Cluster
- Martin Skold

- Unique index implementation of MySQL Cluster
- Integration of NDB Cluster into MySQL
- Magnus Svensson
  - The test framework for MySQL Cluster
  - Integration of NDB Cluster into MySQL
- Tomas Ulin
  - Lots of work on configuration changes for simple installation and use of MySQL Cluster
- Konstantin Osipov
  - Prepared statements.
  - Cursors.
- Dmitri Lenev
  - Time zone support.
  - Triggers (in 5.0).

## C.2. Contributors to MySQL

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

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

The initial port to Win32/NT.

- Per Eric Olsson

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

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

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

- David J. Hughes

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

- Patrick Lynch  
For helping us acquire <http://www.mysql.com/>.
- Fred Lindberg  
For setting up qmail to handle the MySQL mailing list and for the incredible help we got in managing the MySQL mailing lists.
- Igor Romanenko <[igor@frog.kiev.ua](mailto:igor@frog.kiev.ua)>  
[mysqldump](#) (previously [msqldump](#), but ported and enhanced by Monty).
- Yuri Dario  
For keeping up and extending the MySQL OS/2 port.
- Tim Bunce  
Author of [mysqlhotcopy](#).
- Zarko Mocnik <[zarko.mocnik@dem.si](mailto:zarko.mocnik@dem.si)>  
Sorting for Slovenian language.
- "TAMITO" <[tommy@valley.ne.jp](mailto:tommy@valley.ne.jp)>  
The [\\_MB](#) character set macros and the [ujis](#) and [sjis](#) character sets.
- Joshua Chamas <[joshua@chamas.com](mailto:joshua@chamas.com)>  
Base for concurrent insert, extended date syntax, debugging on NT, and answering on the MySQL mailing list.
- Yves Carlier <[Yves.Carlier@rug.ac.be](mailto:Yves.Carlier@rug.ac.be)>  
[mysqlaccess](#), a program to show the access rights for a user.
- Rhys Jones <[rhys@wales.com](mailto:rhys@wales.com)> (And GWE Technologies Limited)  
For one of the early JDBC drivers.
- Dr Xiaokun Kelvin ZHU <[X.Zhu@brad.ac.uk](mailto:X.Zhu@brad.ac.uk)>  
Further development of one of the early JDBC drivers and other MySQL-related Java tools.
- James Cooper <[pixel@organic.com](mailto:pixel@organic.com)>  
For setting up a searchable mailing list archive at his site.
- Rick Mehalick <[Rick\\_Mehalick@i-o.com](mailto:Rick_Mehalick@i-o.com)>  
For [xmysql](#), a graphical X client for MySQL Server.
- Doug Sisk <[sisk@wix.com](mailto:sisk@wix.com)>  
For providing RPM packages of MySQL for Red Hat Linux.
- Diemand Alexander V. <[axeld@vial.ethz.ch](mailto:axeld@vial.ethz.ch)>

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

- Antoni Pamies Olive <toni@readysoft.es>

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

- Jay Bloodworth <jay@pathways.sde.state.sc.us>

For providing RPM versions for MySQL 3.21.

- David Sacerdote <davids@secnet.com>

Ideas for secure checking of DNS hostnames.

- Wei-Jou Chen <jou@nematic.ieo.nctu.edu.tw>

Some support for Chinese(BIG5) characters.

- Wei He <hewei@mail.ied.ac.cn>

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

- Jan Pazdziora <adelton@fi.muni.cz>

Czech sorting order.

- Zeev Suraski <bourbon@netvision.net.il>

`FROM_UNIXTIME()` time formatting, `ENCRYPT()` functions, and `bison` advisor. Active mailing list member.

- Luuk de Boer <luuk@wxs.nl>

Ported (and extended) the benchmark suite to `DBI/DBD`. Have been of great help with `crash-me` and running benchmarks. Some new date functions. The `mysql_setpermission` script.

- Alexis Mikhailov <root@medinf.chuvashia.su>

User-defined functions (UDFs); `CREATE FUNCTION` and `DROP FUNCTION`.

- Andreas F. Bobak <bobak@relog.ch>

The `AGGREGATE` extension to user-defined functions.

- Ross Wakelin <R.Wakelin@march.co.uk>

Help to set up InstallShield for MySQL-Win32.

- Jethro Wright III <jetman@li.net>

The `libmysql.dll` library.

- James Pereria <jpereira@iafrica.com>

MySQLmanager, a Win32 GUI tool for administering MySQL Servers.

- Curt Sampson <cjs@portal.ca>

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



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

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

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

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

- Daniel Koch <[dkoch@amcity.com](mailto:dkoch@amcity.com)>  
Irix setup.

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

### C.3. Documenters and translators

The following people have helped us with writing the MySQL documentation and translating the docu-

mentation or error messages in MySQL.

- Paul DuBois

Ongoing help with making this manual correct and understandable. That includes rewriting Monty's and David's attempts at English into English as other people know it.

- Kim Aldale

Helped to rewrite Monty's and David's early attempts at English into English.

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

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

- Yan Cailin

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

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

Big parts of the Perl `DBI/DBD` section in the manual.

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

Proof-reading of the Reference Manual.

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

French error messages.

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

Czech error messages.

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

Polish error messages.

- Miguel Angel Fernandez Roiz

Spanish error messages.

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

Norwegian error messages and testing of MySQL 3.21.xx.

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

Russian error messages.

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

Italian error messages.

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

German error messages.

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

Slovak error messages.

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

Romanian error messages.

- Peter Feher

Hungarian error messages.

- Roberto M. Serqueira

Portuguese error messages.

- Carsten H. Pedersen

Danish error messages.

- Arjen Lentz

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

## C.4. Libraries used by and included with MySQL

The following is a list of the creators of the libraries we have included with the MySQL server source to make it easy to compile and install MySQL. We are very thankful to all individuals that have created these and it has made our life much easier.

- Fred Fish

For his excellent C debugging and trace library. Monty has made a number of smaller improvements to the library (speed and additional options).

- Richard A. O'Keefe

For his public domain string library.

- Henry Spencer

For his regex library, used in `WHERE column REGEXP regexp`.

- Chris Provenzano

Portable user level pthreads. From the copyright: This product includes software developed by Chris Provenzano, the University of California, Berkeley, and contributors. We are currently using version 1\_60\_beta6 patched by Monty (see [mit-pthreads/Changes-mysql](#)).

- Jean-loup Gailly and Mark Adler

For the zlib library (used on MySQL on Windows).

- Bjorn Benson

For his `safe_malloc` (memory checker) package which is used in when you configure MySQL with `-debug`.

- Free Software Foundation

The `readline` library (used by the `mysql` command-line client).

- The NetBSD foundation

The `libedit` package (optionally used by the `mysql` command-line client).

## C.5. Packages that support MySQL

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

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

- Tim Bunce, Alligator Descartes

For the `DBD` (Perl) interface.

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

For the Perl interface for MySQL Server.

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

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

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

For porting PHP for MySQL Server.

- Georg Richter

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

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

For porting iODBC (Unix ODBC).

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

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

## C.6. Tools that were used to create MySQL

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

- Free Software Foundation

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

- Free Software Foundation & The XEmacs development team

For a really great editor/environment used by almost everybody at MySQL AB.

- Julian Seward

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

- Dorothea Lütkehaus and Andreas Zeller

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

## C.7. Supporters of MySQL

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

- VA Linux / Andover.net

Funded replication.

- NuSphere

Editing of the MySQL manual.

- Stork Design studio

The MySQL Web site in use between 1998-2000.

- Intel

Contributed to development on Windows and Linux platforms.

- Compaq

Contributed to Development on Linux/Alpha.

- SWSOft

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

- FutureQuest

[--skip-show-database](#)

---

# Appendix D. MySQL Change History

This appendix lists the changes from version to version in the MySQL source code through the latest version of MySQL 4.1.

We are working actively on MySQL 4.1 and 5.0, and provide only critical bugfixes for MySQL 4.0. We update this section as we add new features, so that everybody can follow the development.

Note that we tend to update the manual at the same time we make changes to MySQL. If you find a recent version of MySQL listed here that you can't find on our download page (<http://dev.mysql.com/downloads/>), it means that the version has not yet been released.

The date mentioned with a release version is the date of the last BitKeeper ChangeSet on which the release was based, not the date when the packages were made available. The binaries are usually made available a few days after the date of the tagged ChangeSet, because building and testing all packages takes some time.

The manual included in the source and binary distributions may not be fully accurate when it comes to the release changelog entries, because the integration of the manual happens at build time. For the most up-to-date release changelog, please refer to the online version instead.

## D.1. Changes in release 4.1.x (Production)

Version 4.1 of the MySQL server includes many enhancements and new features. Binaries for this version are available for download at <http://dev.mysql.com/downloads/mysql-4.1.html>.

- The `SUBSTRING()` function can now take a negative value for the `pos` (position) argument. See [Section 12.3, “String Functions”](#).
- Subqueries and derived tables (unnamed views). See [Section 13.2.8, “Subquery Syntax”](#).
- `INSERT ... ON DUPLICATE KEY UPDATE ...` syntax. This allows you to `UPDATE` an existing row if the insert would cause a duplicate value in a `PRIMARY` or `UNIQUE` key. (`REPLACE` allows you to overwrite an existing row, which is something entirely different.) See [Section 13.2.4, “INSERT Syntax”](#).
- A newly designed `GROUP_CONCAT()` aggregate function. See [Section 12.10, “Functions and Modifiers for Use with GROUP BY Clauses”](#).
- Extensive Unicode (UTF8) support.
- Table names and column names now are stored in `UTF8`. This makes MySQL more flexible, but might cause some problems upgrading if you have table or column names that use characters outside of the standard 7-bit US-ASCII range. See [Section 2.11.1, “Upgrading from MySQL 4.0 to 4.1”](#).
- Character sets can be defined per column, table, and database.
- New key cache for `MyISAM` tables with many tunable parameters. You can have multiple key caches, preload index into caches for batches...
- `BTREE` index on `HEAP` tables.
- Support for OpenGIS spatial types (geographical data). See [Chapter 16, \*Spatial Extensions\*](#).
- `SHOW WARNINGS` shows warnings for the last command. See [Section 13.5.4.21, “SHOW WARNINGS Syntax”](#).

- Faster binary protocol with prepared statements and parameter binding. See [Section 17.2.4, “C API Prepared Statements”](#).
- You can now issue multiple statements with a single C API call and then read the results in one go. See [Section 17.2.9, “C API Handling of Multiple Statement Execution”](#).
- Create Table: `CREATE [TEMPORARY] TABLE [IF NOT EXISTS] table2 LIKE table1`.
- Server based `HELP` command that can be used in the `mysql` command-line client (and other clients) to get help for SQL statements.

For a full list of changes, please refer to the changelog sections for each individual 4.1.x release.

### D.1.1. Changes in release 4.1.22 (Not yet released)

This is a bugfix release for the recent production release family.

This section documents all changes and bug fixes that have been applied since the last official MySQL release. If you would like to receive more fine-grained and personalized *update alerts* about fixes that are relevant to the version and features you use, please consider subscribing to *MySQL Network* (a commercial MySQL offering). For more details please see <http://www.mysql.com/network/advisors.html>.

Functionality added or changed:

- The `mysqld` manpage has been reclassified from volume 1 to volume 8. ([Bug#21220](#) [<http://bugs.mysql.com/21220>])
- MySQL now can do stack dumps on `x86_64` and `i386/NTPL` systems. ([Bug#21250](#) [<http://bugs.mysql.com/21250>])
- The `LOAD DATA FROM MASTER` and `LOAD TABLE FROM MASTER` statements are deprecated. See [Section 13.6.2.2, “LOAD DATA FROM MASTER Syntax”](#), for recommended alternatives. ([Bug#18822](#) [<http://bugs.mysql.com/18822>])

Bugs fixed:

- For `INSERT ... ON DUPLICATE KEY UPDATE`, use of `VALUES(col_name)` within the `UPDATE` clause sometimes was handled incorrectly. ([Bug#21555](#) [<http://bugs.mysql.com/21555>])
- Table aliases in multiple-table `DELETE` statements sometimes were not resolved. ([Bug#21392](#) [<http://bugs.mysql.com/21392>])
- `EXPORT_SET()` did not accept arguments with coercible character sets. ([Bug#21531](#) [<http://bugs.mysql.com/21531>])
- The `--collation-server` server option was being ignored. With the fix for this problem, if you choose a non-default character set with `--character-set-server`, you should also use `--collation-server` to specify the collation. ([Bug#15276](#) [<http://bugs.mysql.com/15276>])
- A subquery that uses an index for both the `WHERE` and `ORDER BY` clauses produced an empty result. ([Bug#21180](#) [<http://bugs.mysql.com/21180>])
- Queries containing a subquery that used aggregate functions could return incorrect results. ([Bug#16792](#) [<http://bugs.mysql.com/16792>])



- The `MD5()`, `SHA1()`, and `ENCRYPT()` functions should return a binary string, but the result sometimes was converted to the character set of the argument. `MAKE_SET()` and `EXPORT_SET()` now use the correct character set for their default separators, resulting in consistent result strings which can be coerced according to normal character set rules. (Bug#20536 [http://bugs.mysql.com/20536])
- Use of `mysampack` or `mysamchk` on a table with `FULLTEXT` indexing resulted in table corruption. (Bug#19702 [http://bugs.mysql.com/19702])
- The optimizer could produce an incorrect result after `AND` with collations such as `latin1_german2_ci`, `utf8_czech_ci`, and `utf8_lithianian_ci`. (Bug#9509 [http://bugs.mysql.com/9509])
- `character_set_results` can be `NULL` to signify “no conversion,” but some code did not check for `NULL`, resulting in a server crash. (Bug#21913 [http://bugs.mysql.com/21913])
- The `ndb_mgm` program was included in both the `MySQL-ndb-tools` and `MySQL-ndb-management` RPM packages, resulting in a conflict if both were installed. Now `ndb_mgm` is included only in `MySQL-ndb-tools`. (Bug#21058 [http://bugs.mysql.com/21058])
- The `mysam_stats_method` variable was mishandled when set from an option file or on the command line. (Bug#21054 [http://bugs.mysql.com/21054])
- `libmysqld` produced some warnings to `stderr` which could not be silenced. These warnings now are suppressed. (Bug#13717 [http://bugs.mysql.com/13717])
- If a column definition contained a character set declaration, but a `DEFAULT` value began with an introducer, the introducer character set was used as the column character set. (Bug#20695 [http://bugs.mysql.com/20695])
- Some Linux-x86\_64-icc packages (of previous releases) mistakenly contained 32-bit binaries. Only `ICC` builds are affected, not `gcc` builds. Solaris and FreeBSD x86\_64 builds are not affected. (Bug#22238 [http://bugs.mysql.com/22238])
- For `TIME_FORMAT()`, the `%H` and `%k` format specifiers can return values larger than two digits (if the hour is greater than 99), but for some query results that contained three-character hours, column values were truncated. (Bug#19844 [http://bugs.mysql.com/19844])
- For table-format output, `mysql` did not always calculate columns widths correctly for columns containing multi-byte characters in the column name or contents. (Bug#17939 [http://bugs.mysql.com/17939])
- Views could not be updated within a stored function or trigger. (Bug#17591 [http://bugs.mysql.com/17591])
- Usernames have a maximum length of 16 characters (even if they contain multi-byte characters), but were being truncated to 16 bytes. (Bug#20393 [http://bugs.mysql.com/20393])
- Database and table names have a maximum length of 64 characters (even if they contain multi-byte characters), but were being truncated to 64 bytes. (Bug#21432 [http://bugs.mysql.com/21432])
- When using tables created under MySQL 4.1 with a 5.0 server, if the tables contained `VARCHAR` columns, for some queries the metadata sent to the client could have an empty column name. (Bug#14897 [http://bugs.mysql.com/14897])
- On 64-bit systems, use of the `cp1250` character set with a primary key column in a `LIKE` clause caused a server crash for patterns having letters in the range 128..255. (Bug#19741 [http://bugs.mysql.com/19741])
- A subquery in the `WHERE` clause of the outer query and using `IN` and `GROUP BY` returned an incor-

rect result. ([Bug#16255](http://bugs.mysql.com/16255) [<http://bugs.mysql.com/16255>])

- `COUNT( * )` queries with `ORDER BY` and `LIMIT` could return the wrong result. ([Bug#21787](http://bugs.mysql.com/21787) [<http://bugs.mysql.com/21787>])

**Note:** This problem was introduced by the fix for [Bug#9676](http://bugs.mysql.com/9676) [<http://bugs.mysql.com/9676>], which limited the rows stored in a temporary table to the `LIMIT` clause. This optimization is not applicable to non-group queries with aggregate functions. The current fix disables the optimization in such cases.

- Running `SHOW MASTER LOGS` at the same time as binary log files were being switched would cause `mysqld` to hang. ([Bug#21965](http://bugs.mysql.com/21965) [<http://bugs.mysql.com/21965>])
- Adding `ORDER BY` to a `SELECT DISTINCT(expr)` query could produce incorrect results. ([Bug#21456](http://bugs.mysql.com/21456) [<http://bugs.mysql.com/21456>])
- For `InnoDB` tables, the server could crash when executing `NOT IN ( )` subqueries. ([Bug#21077](http://bugs.mysql.com/21077) [<http://bugs.mysql.com/21077>])
- `mysqld --flush` failed to flush changes to disk following an `UPDATE` statement for which no updated column had an index. ([Bug#20060](http://bugs.mysql.com/20060) [<http://bugs.mysql.com/20060>])
- Setting `myisam_repair_threads` caused any repair operation on the table to fail to update the cardinality of indexes, instead making them always equal to 1. ([Bug#18874](http://bugs.mysql.com/18874) [<http://bugs.mysql.com/18874>])
- The `--with-collation` option was not honored for client connections. ([Bug#7192](http://bugs.mysql.com/7192) [<http://bugs.mysql.com/7192>])
- `NDB Cluster`: `ndb_size.pl` and `ndb_error_reporter` were missing from RPM packages. ([Bug#20426](http://bugs.mysql.com/20426) [<http://bugs.mysql.com/20426>])
- `NDB Cluster`: Setting `TransactionDeadlockDetectionTimeout` to a value greater than 12000 would cause scans to deadlock, time out, fail to release scan records, until the cluster ran out of scan records and stopped processing. ([Bug#21800](http://bugs.mysql.com/21800) [<http://bugs.mysql.com/21800>])
- `NDB Cluster`: The server provided a non-descriptive error message when encountering a fatally corrupted REDO log. ([Bug#21615](http://bugs.mysql.com/21615) [<http://bugs.mysql.com/21615>])
- `NDB Cluster`: A partial rollback could lead to node restart failures. ([Bug#21536](http://bugs.mysql.com/21536) [<http://bugs.mysql.com/21536>])
- `NDB Cluster`: The failure of a unique index read due to an invalid schema version could be handled incorrectly in some cases, leading to unpredictable results. ([Bug#21384](http://bugs.mysql.com/21384) [<http://bugs.mysql.com/21384>])
- `NDB Cluster`: In a cluster with more than 2 replicas, a manual restart of one of the data nodes could fail and cause the other nodes in its nodegroup to shut down. ([Bug#21213](http://bugs.mysql.com/21213) [<http://bugs.mysql.com/21213>])
- `NDB Cluster`: When the redo buffer ran out of space, a Pointer too large error was raised and the cluster could become unusable until restarted with `--initial`. ([Bug#20892](http://bugs.mysql.com/20892) [<http://bugs.mysql.com/20892>])
- `NDB Cluster`: In some situations with a high disk-load, writing of the redo log could hang, causing a crash with the error message GCP STOP detected. ([Bug#20904](http://bugs.mysql.com/20904) [<http://bugs.mysql.com/20904>])
- `NDB Cluster`: `ndb_size.pl` and `ndb_error_reporter` were missing from RPM pack-

- ages. ([Bug#20426](http://bugs.mysql.com/20426) [<http://bugs.mysql.com/20426>])
- **NDB Cluster**: The server failed with a non-descriptive error message when out of data memory. ([Bug#18475](http://bugs.mysql.com/18475) [<http://bugs.mysql.com/18475>])
  - **NDB Cluster**: `SELECT ... FOR UPDATE` failed to lock the selected rows. ([Bug#18184](http://bugs.mysql.com/18184) [<http://bugs.mysql.com/18184>])
  - **NDB Cluster**: Some queries involving joins on very large **NDB** tables could crash the MySQL server. ([Bug#21059](http://bugs.mysql.com/21059) [<http://bugs.mysql.com/21059>])
  - Character set collation was ignored in `GROUP BY` clauses. ([Bug#20709](http://bugs.mysql.com/20709) [<http://bugs.mysql.com/20709>])
  - A query using `WHERE column = constant OR column IS NULL` did not return consistent results on successive invocations. The `column` in each part of the `WHERE` clause could be either the same column, or two different columns, for the effect to be observed. ([Bug#21091](http://bugs.mysql.com/21091) [<http://bugs.mysql.com/21091>])
  - A query using `WHERE NOT (column < ANY (subquery))` yielded a different result from the same query using the same `column` and `subquery` with `WHERE (column > ANY (subquery))`. ([Bug#20975](http://bugs.mysql.com/20975) [<http://bugs.mysql.com/20975>])
  - Using the extended syntax for `TRIM()` — that is, `TRIM(... FROM ...)` — caused erroneous output from `EXPLAIN EXTENDED` statements. ([Bug#17526](http://bugs.mysql.com/17526) [<http://bugs.mysql.com/17526>])
  - `DELETE` with `WHERE` condition on a **BTREE**-indexed column for a **MEMORY** table deleted only the first matched row. ([Bug#9719](http://bugs.mysql.com/9719) [<http://bugs.mysql.com/9719>])
  - For cross-database multiple-table `UPDATE` statements, a user with all privileges for the default database could update tables in another database for which the user did not have `UPDATE` privileges. ([Bug#7391](http://bugs.mysql.com/7391) [<http://bugs.mysql.com/7391>])
  - `mysql_install_db` incorrectly had a blank first line. ([Bug#20721](http://bugs.mysql.com/20721) [<http://bugs.mysql.com/20721>])
  - Under heavy load (executing more than 1024 simultaneous complex queries), a problem in the code that handles internal temporary tables could lead to writing beyond allocated space and memory corruption. ([Bug#21206](http://bugs.mysql.com/21206) [<http://bugs.mysql.com/21206>])
  - Multiple invocations of the `REVERSE()` function could return different results. ([Bug#18243](http://bugs.mysql.com/18243) [<http://bugs.mysql.com/18243>])
  - Conversion of `TIMESTAMP` values between UTC and the local time zone resulted in some values having the year 2069 rather than 1969. ([Bug#16327](http://bugs.mysql.com/16327) [<http://bugs.mysql.com/16327>])
  - Under certain circumstances, `AVG(key_val)` returned a value but `MAX(key_val)` returned an empty set due to incorrect application of `MIN()/MAX()` optimization. ([Bug#20954](http://bugs.mysql.com/20954) [<http://bugs.mysql.com/20954>])
  - Using aggregate functions in subqueries yielded incorrect results under certain circumstances due to incorrect application of `MIN()/MAX()` optimization. ([Bug#20792](http://bugs.mysql.com/20792) [<http://bugs.mysql.com/20792>])
  - Using `> ALL` with subqueries that return no rows yielded incorrect results under certain circumstances due to incorrect application of `MIN()/MAX()` optimization. ([Bug#18503](http://bugs.mysql.com/18503) [<http://bugs.mysql.com/18503>])
  - Using `ANY` with “non-table” subqueries such as `SELECT 1` yielded incorrect results under certain circumstances due to incorrect application of `MIN()/MAX()` optimization. ([Bug#16302](http://bugs.mysql.com/16302) [<http://bugs.mysql.com/16302>])

[<http://bugs.mysql.com/16302>])

- The use of `WHERE col_name IS NULL` in `SELECT` statements reset the value of `LAST_INSERT_ID( )` to zero. (Bug#14553 [<http://bugs.mysql.com/14553>])
- Use of the join cache in favor of an index for `ORDER BY` operations could cause incorrect result sorting. (Bug#17212 [<http://bugs.mysql.com/17212>])
- `libmysqld` returned `TEXT` columns to the client as number of bytes, not number of characters (which can be different for multi-byte character sets). (Bug#19983 [<http://bugs.mysql.com/19983>])

## D.1.2. Changes in release 4.1.21 (19 July 2006)

This is a bugfix release for the recent production release family.

This section documents all changes and bug fixes that have been applied since the last official MySQL release. If you would like to receive more fine-grained and personalized *update alerts* about fixes that are relevant to the version and features you use, please consider subscribing to *MySQL Network* (a commercial MySQL offering). For more details please see <http://www.mysql.com/network/advisors.html>.

Functionality added or changed:

- For spatial data types, the server formerly returned these as `VARSTRING` values with a binary collation. Now the server returns spatial values as `BLOB` values. (Bug#10166 [<http://bugs.mysql.com/10166>])
- Added the `--set-charset` option to `mysqlbinlog` to allow the character set to be specified for processing binary log files. (Bug#18351 [<http://bugs.mysql.com/18351>])
- For a table with an `AUTO_INCREMENT` column, `SHOW CREATE TABLE` now shows the next `AUTO_INCREMENT` value to be generated. (Bug#19025 [<http://bugs.mysql.com/19025>])
- A warning now is issued if the client attempts to set the `SQL_LOG_OFF` variable without the `SUPER` privilege. (Bug#16180 [<http://bugs.mysql.com/16180>])
- The `mysqldumpslow` script has been moved from client RPM packages to server RPM packages. This corrects a problem where `mysqldumpslow` could not be used with a client-only RPM install, because it depends on `my_print_defaults` which is in the server RPM. (Bug#20216 [<http://bugs.mysql.com/20216>])

Bugs fixed:

- **Security fix:** On Linux, and possibly other platforms using case-sensitive filesystems, it was possible for a user granted rights on a database to create or access a database whose name differed only from that of the first by the case of one or more letters. (CVE-2006-4226 [<http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2006-4226>], Bug#17647 [<http://bugs.mysql.com/17647>])
- **Security fix:** If a user has access to `MyISAM` table `t`, that user can create a `MERGE` table `m` that accesses `t`. However, if the user's privileges on `t` are subsequently revoked, the user can continue to access `t` by doing so through `m`. If this behavior is undesirable, you can start the server with the new `--skip-merge` option to disable the `MERGE` storage engine. (Bug#15195 [<http://bugs.mysql.com/15195>])
- **Security fix:** Invalid arguments to `DATE_FORMAT( )` caused a server crash. (CVE-2006-3469

[<http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2006-3469>], [Bug#20729](#) [<http://bugs.mysql.com/20729>]) Thanks to Jean-David Maillefer for discovering and reporting this problem to the Debian project and to Christian Hammers from the Debian Team for notifying us of it.

- Closing of temporary tables failed if binary logging was not enabled. ([Bug#20919](#) [<http://bugs.mysql.com/20919>])
- Repeated `DROP TABLE` statements in a stored procedure could sometimes cause the server to crash. ([Bug#19399](#) [<http://bugs.mysql.com/19399>])
- `DATE_ADD()` and `DATE_SUB()` returned `NULL` when the result date was on the day '9999-12-31'. ([Bug#12356](#) [<http://bugs.mysql.com/12356>])
- For a `DATE` parameter sent via a `MYSQL_TIME` data structure, `mysql_stmt_execute()` zeroed the hour, minute, and second members of the structure rather than treating them as read-only. ([Bug#20152](#) [<http://bugs.mysql.com/20152>])
- The `DATA DIRECTORY` table option did not work for `TEMPORARY` tables. ([Bug#8706](#) [<http://bugs.mysql.com/8706>])
- The `mysql` client did not ignore client-specific commands (such as `use` or `help`) that occurred as the first word on a line within multiple-line `/* ... */` comments. ([Bug#20432](#) [<http://bugs.mysql.com/20432>])
- The `mysql` client did not understand `help` commands that had spaces at the end. ([Bug#20328](#) [<http://bugs.mysql.com/20328>])
- Failure to account for a `NULL` table pointer on big-endian machines could cause a server crash during type conversion. ([Bug#21135](#) [<http://bugs.mysql.com/21135>])
- Some memory leaks in the `libmysqld` embedded server were corrected. ([Bug#16017](#) [<http://bugs.mysql.com/16017>])
- When `mysqldump` disabled keys and locked a `MyISAM` table, the lock operation happened second. If another client performed a query on the table in the interim, it could take a long time due to indexes not being used. Now the lock operation happens first. ([Bug#15977](#) [<http://bugs.mysql.com/15977>])
- The length of the pattern string prefix for `LIKE` operations was calculated incorrectly for multi-byte character sets. As a result, the scanned range was wider than necessary if the prefix contained any multi-byte characters. ([Bug#16674](#) [<http://bugs.mysql.com/16674>], [Bug#18359](#) [<http://bugs.mysql.com/18359>])
- For very complex `SELECT` statements could create temporary tables that were too big, but for which the temporary files did not get removed, causing subsequent queries to fail. ([Bug#11824](#) [<http://bugs.mysql.com/11824>])
- Using `SELECT` and a table join while running a concurrent `INSERT` operation would join incorrect rows. ([Bug#14400](#) [<http://bugs.mysql.com/14400>])
- Using `SELECT` on a corrupt table using the dynamic record format could cause a server crash. ([Bug#19835](#) [<http://bugs.mysql.com/19835>])
- Checking a spatial table (using `CHECK TABLE`) with an index and only one row would indicate a table corruption. ([Bug#17877](#) [<http://bugs.mysql.com/17877>])
- For `SELECT ... FOR UPDATE` statements that used `DISTINCT` or `GROUP BY` over all key parts of a unique index (or primary key), the optimizer unnecessarily created a temporary table, thus

losing the linkage to the underlying unique index values. This caused a `Result set not updatable` error. (The temporary table is unnecessary because under these circumstances the distinct or grouped columns must also be unique.) ([Bug#16458](http://bugs.mysql.com/16458) [<http://bugs.mysql.com/16458>])

- Concatenating the results of multiple constant subselects produced incorrect results. ([Bug#16716](http://bugs.mysql.com/16716) [<http://bugs.mysql.com/16716>])
- The use of `MIN()` and `MAX()` on columns with a partial index produced incorrect results in some queries. ([Bug#18206](http://bugs.mysql.com/18206) [<http://bugs.mysql.com/18206>])
- Use of `MIN()` or `MAX()` with `GROUP BY` on a `ucs2` column could cause a server crash. ([Bug#20076](http://bugs.mysql.com/20076) [<http://bugs.mysql.com/20076>])
- `INSERT INTO ... SELECT ... LIMIT 1` could be slow because the `LIMIT` was ignored when selecting candidate rows. ([Bug#9676](http://bugs.mysql.com/9676) [<http://bugs.mysql.com/9676>])
- **NDB Cluster:** A Cluster whose storage nodes were installed from the `MySQL-ndb-storage-*` RPMs could not perform `CREATE` or `ALTER` operations that made use of non-default character sets or collations. ([Bug#14918](http://bugs.mysql.com/14918) [<http://bugs.mysql.com/14918>])
- **NDB Cluster:** The repeated creating and dropping of a table would eventually lead to `NDB Error 826, Too many tables and attributes ... Insufficient space.` ([Bug#20847](http://bugs.mysql.com/20847) [<http://bugs.mysql.com/20847>])
- **NDB Cluster:** When attempting to restart the cluster following a data import, the cluster would fail during Phase 4 of the restart with Error 2334: Job buffer congestion. ([Bug#20774](http://bugs.mysql.com/20774) [<http://bugs.mysql.com/20774>])
- **NDB Cluster:** A node failure during a scan could sometime cause the node to crash when restarting too quickly following the failure. ([Bug#20197](http://bugs.mysql.com/20197) [<http://bugs.mysql.com/20197>])
- **NDB Cluster:** It was possible to use port numbers greater than 65535 for `ServerPort` in the `config.ini` file. ([Bug#19164](http://bugs.mysql.com/19164) [<http://bugs.mysql.com/19164>])
- The omission of leading zeros in dates could lead to erroneous results when these were compared with the output of certain date and time functions. ([Bug#16377](http://bugs.mysql.com/16377) [<http://bugs.mysql.com/16377>])
- Certain queries having a `WHERE` clause that included conditions on multi-part keys with more than 2 key parts could produce incorrect results and send [Note] Use\_count: Wrong count for key at... messages to `STDERR`. ([Bug#16168](http://bugs.mysql.com/16168) [<http://bugs.mysql.com/16168>])
- An invalid comparison between keys in partial indexes over multi-byte character fields could lead to incorrect result sets if the selected query execution plan used a range scan by a partial index over a `UTF8` character field. This also caused incorrect results under similar circumstances with many other character sets. ([Bug#14896](http://bugs.mysql.com/14896) [<http://bugs.mysql.com/14896>])
- **NDB Cluster:** The cluster's data nodes would fail while trying to load data when `NoOfFragmentLogFiles` was equal to 1. ([Bug#19894](http://bugs.mysql.com/19894) [<http://bugs.mysql.com/19894>])
- **NDB Cluster:** A problem with error handling when `ndb_use_exact_count` was enabled could lead to incorrect values returned from queries using `COUNT()`. A warning is now returned in such cases. ([Bug#19202](http://bugs.mysql.com/19202) [<http://bugs.mysql.com/19202>])
- **NDB Cluster:** `LOAD DATA LOCAL` failed to ignore duplicate keys in Cluster tables. ([Bug#19496](http://bugs.mysql.com/19496) [<http://bugs.mysql.com/19496>])
- **NDB Cluster:** Repeated `CREATE - INSERT - DROP` operations tables could in some circumstances cause the MySQL table definition cache to become corrupt, so that some `mysqld` processes could access table information but others could not. ([Bug#18595](http://bugs.mysql.com/18595) [<http://bugs.mysql.com/18595>])

- **NDB Cluster:** The `mgm` client command `ALL CLUSTERLOG STATISTICS=15;` had no effect. (Bug#20336 [http://bugs.mysql.com/20336])
- **NDB Cluster:** `TRUNCATE TABLE` failed to reset the `AUTO_INCREMENT` counter. (Bug#18864 [http://bugs.mysql.com/18864])
- **NDB Cluster:** The failure of a data node when preparing to commit a transaction (that is, while the node's status was `CS_PREPARE_TO_COMMIT`) could cause the failure of other cluster data nodes. (Bug#20185 [http://bugs.mysql.com/20185])
- **NDB Cluster:** Renaming a table in such a way as to move it to a different database failed to move the table's indexes. (Bug#19967 [http://bugs.mysql.com/19967])
- **NDB Cluster:** Resources for unique indexes on Cluster table columns were incorrectly allocated, so that only one-fourth as many unique indexes as indicated by the value of `UniqueHashIndexes` could be created. (Bug#19623 [http://bugs.mysql.com/19623])
- **NDB Cluster (NDBAPI):** On big-endian platforms, `NdbOperation::write_attr()` did not update 32-bit fields correctly. (Bug#19537 [http://bugs.mysql.com/19537])
- **NDB Cluster:** Some queries having a `WHERE` clause of the form `c1=val1 OR c2 LIKE 'val2'` were not evaluated correctly. (Bug # 17421)
- **NDB Cluster:** Using “stale” `mysqld .FRM` files could cause a newly-restored cluster to fail. This situation could arise when restarting a MySQL Cluster using the `--initial` option while leaving connected `mysqld` processes running. (Bug#16875 [http://bugs.mysql.com/16875])
- **NDB Cluster:** Repeated use of the `SHOW` and `ALL STATUS` commands in the `ndb_mgm` client could cause the `mgmd` process to crash. (Bug#18591 [http://bugs.mysql.com/18591])
- **NDB Cluster:** An issue with `ndb_mgmd` prevented more than 27 `mysqld` processes from connecting to a single cluster at one time. (Bug#17150 [http://bugs.mysql.com/17150])
- **NDB Cluster:** Data node failures could cause excessive CPU usage by `ndb_mgmd`. (Bug#13987 [http://bugs.mysql.com/13987])
- **NDB Cluster:** `TRUNCATE` failed on tables having `BLOB` or `TEXT` columns with the error Lock wait timeout exceeded. (Bug#19201 [http://bugs.mysql.com/19201])
- A cast problem caused incorrect results for prepared statements that returned float values when MySQL was compiled with `gcc 4.0`. (Bug#19694 [http://bugs.mysql.com/19694])
- Improper character set initialization in the embedded server could result in a server crash. (Bug#20318 [http://bugs.mysql.com/20318])
- Some queries that used `ORDER BY` and `LIMIT` performed quickly in MySQL 3.23, but slowly in MySQL 4.x/5.x due to an optimizer problem. (Bug#4981 [http://bugs.mysql.com/4981])
- Queries using an indexed column as the argument for the `MIN()` and `MAX()` functions following an `ALTER TABLE ... DISABLE KEYS` statement returned Got error 124 from storage engine until `ALTER TABLE ... ENABLE KEYS` was run on the table. (Bug#20357 [http://bugs.mysql.com/20357])
- A number of dependency issues in the RPM `bench` and `test` packages caused installation of these packages to fail. (Bug#20078 [http://bugs.mysql.com/20078])
- The `MD5()` and `SHA()` functions treat their arguments as case-sensitive strings. But when they are compared, their arguments were compared as case-insensitive strings, which leads to two function calls with different arguments (and thus different results) compared as being identical. This can lead

- to a wrong decision made in the range optimizer and thus to an incorrect result set. (Bug#15351 [<http://bugs.mysql.com/15351>])
- InnoDB unlocked its data directory before committing a transaction, potentially resulting in non-recoverable tables if a server crash occurred before the commit. (Bug#19727 [<http://bugs.mysql.com/19727>])
  - Multiple-table DELETE statements containing a subquery that selected from one of the tables being modified caused a server crash. (Bug#19225 [<http://bugs.mysql.com/19225>])
  - The ARCHIVE storage engine does not support TRUNCATE TABLE, but the server was not returning an appropriate error when truncation of an ARCHIVE table was attempted. (Bug#15558 [<http://bugs.mysql.com/15558>])
  - An update that used a join of a table to itself and modified the table on both sides of the join reported the table as crashed. (Bug#18036 [<http://bugs.mysql.com/18036>])
  - The fill\_help\_tables.sql file did not load properly if the ANSI\_QUOTES SQL mode was enabled. (Bug#20542 [<http://bugs.mysql.com/20542>])
  - The fill\_help\_tables.sql file did not contain a SET NAMES 'utf8' statement to indicate its encoding. This caused problems for some settings of the MySQL character set such as big5. (Bug#20551 [<http://bugs.mysql.com/20551>])
  - The MySQL server startup script /etc/init.d/mysql (created from mysql.server) is now marked to ensure that the system services ypbind, nscd, ldap, and NTP are started first (if these are configured on the machine). (Bug#18810 [<http://bugs.mysql.com/18810>])
  - For a reference to a non-existent index in FORCE INDEX, the error message referred to a column, not an index. (Bug#17873 [<http://bugs.mysql.com/17873>])
  - In a multiple-row INSERT statement, LAST\_INSERT\_ID() should return the same value for each row. However, in some cases, the value could change if the table being inserted into had its own AUTO\_INCREMENT column. (Bug#6880 [<http://bugs.mysql.com/6880>])
  - Invalid escape sequences in option files caused MySQL programs that read them to abort. (Bug#15328 [<http://bugs.mysql.com/15328>])
  - Binary log lacked character set information for table name when dropping temporary tables. (Bug#14157 [<http://bugs.mysql.com/14157>])
  - InnoDB did not increment the handler\_read\_prev counter. (Bug#19542 [<http://bugs.mysql.com/19542>])
  - Slave SQL thread cleanup was not handled properly on Mac OS X when a statement was killed, resulting in a slave crash. (Bug#16900 [<http://bugs.mysql.com/16900>])
  - mysqldump did not respect the order of tables named with the --tables option. (Bug#18536 [<http://bugs.mysql.com/18536>])
  - The server no longer uses a signal handler for signal 0 because it could cause a crash on some platforms. (Bug#15869 [<http://bugs.mysql.com/15869>])
  - LOAD\_FILE() returned an error if the file did not exist, rather than NULL as it should according to the manual. (Bug#10418 [<http://bugs.mysql.com/10418>])
  - Use of uninitialized user variables in a subquery in the FROM clause results in bad entries in the binary log. (Bug#19136 [<http://bugs.mysql.com/19136>])



- `IS_USED_LOCK()` could return an incorrect connection identifier. (Bug#16501 [<http://bugs.mysql.com/16501>])
- Concurrent reading and writing of privilege structures could crash the server. (Bug#16372 [<http://bugs.mysql.com/16372>])
- A statement containing `GROUP BY` and `HAVING` clauses could return incorrect results when the `HAVING` clause contained logic that returned `FALSE` for every row. (Bug#14927 [<http://bugs.mysql.com/14927>])
- `MONTHNAME(STR_TO_DATE(NULL, '%m'))` could cause a server crash. (Bug#18501 [<http://bugs.mysql.com/18501>])
- The `ref` optimizer could choose the `ref_or_null` access method in cases where it was not applicable. This could cause inconsistent `EXPLAIN` or `SELECT` results for a given statement. (Bug#16798 [<http://bugs.mysql.com/16798>])
- `ANALYZE TABLE` for `TEMPORARY` tables had no effect. (Bug#15225 [<http://bugs.mysql.com/15225>])
- When `myisamchk` needed to rebuild a table, `AUTO_INCREMENT` information was lost. (Bug#10405 [<http://bugs.mysql.com/10405>])
- No error message was being issued for storage engines that do not support `ALTER TABLE`. Now an `ER_NOT_SUPPORTED_YET` error occurs. (Bug#7643 [<http://bugs.mysql.com/7643>])
- The binary log would create an incorrect `DROP` query when creating temporary tables during replication. (Bug#17263 [<http://bugs.mysql.com/17263>])
- `SHOW CREATE TABLE` did not display the `AUTO_INCREMENT` column attribute if the SQL mode was `MYSQL323` or `MYSQL40`. This also affected `mysqldump`, which uses `SHOW CREATE TABLE` to get table definitions. (Bug#14515 [<http://bugs.mysql.com/14515>])

### D.1.3. Changes in release 4.1.20 (24 May 2006)

This is a security fix release for the previous production release family.

This release includes the security fix described later in this section and a few other changes to resolve build problems, relative to the last official MySQL release (4.1.19). If you would like to receive more fine-grained and personalized *update alerts* about fixes that are relevant to the version and features you use, please consider subscribing to *MySQL Network* (a commercial MySQL offering). For more details please see <http://www.mysql.com/network/advisors.html>.

Bugs fixed:

- **Security fix:** An SQL-injection security hole has been found in multi-byte encoding processing. The bug was in the server, incorrectly parsing the string escaped with the `mysql_real_escape_string()` C API function. (CVE-2006-2753 [<http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2006-2753>], Bug#8378 [<http://bugs.mysql.com/8378>])

This vulnerability was discovered and reported by Josh Berkus <[josh@postgresql.org](mailto:josh@postgresql.org)> and Tom Lane <[tgl@sss.pgh.pa.us](mailto:tgl@sss.pgh.pa.us)> as part of the inter-project security collaboration of the OS-DB consortium. For more information about SQL injection, please see the following text.

**Discussion:** An SQL-injection security hole has been found in multi-byte encoding processing. An

SQL-injection security hole can include a situation whereby when a user supplied data to be inserted into a database, the user might inject SQL statements into the data that the server will execute. With regards to this vulnerability, when character set unaware-escaping is used (for example, `addslashes()` in PHP), it is possible to bypass the escaping in some multi-byte character sets (for example, SJIS, BIG5 and GBK). As a result, a function such as `addslashes()` is not able to prevent SQL-injection attacks. It is impossible to fix this on the server side. The best solution is for applications to use character set-aware escaping offered by a function such `mysql_real_escape_string()`.

However, a bug was detected in how the MySQL server parses the output of `mysql_real_escape_string()`. As a result, even when the character set-aware function `mysql_real_escape_string()` was used, SQL injection was possible. This bug has been fixed.

**Workarounds:** If you are unable to upgrade MySQL to a version that includes the fix for the bug in `mysql_real_escape_string()` parsing, but run MySQL 5.0.1 or higher, you can use the `NO_BACKSLASH_ESCAPES` SQL mode as a workaround. (This mode was introduced in MySQL 5.0.1.) `NO_BACKSLASH_ESCAPES` enables an SQL standard compatibility mode, where backslash is not considered a special character. The result will be that queries will fail.

To set this mode for the current connection, enter the following SQL statement:

```
SET sql_mode='NO_BACKSLASH_ESCAPES';
```

You can also set the mode globally for all clients:

```
SET GLOBAL sql_mode='NO_BACKSLASH_ESCAPES';
```

This SQL mode also can be enabled automatically when the server starts by using the command-line option `--sql-mode=NO_BACKSLASH_ESCAPES` or by setting `sql_mode=NO_BACKSLASH_ESCAPES` in the server option file (for example, `my.cnf` or `my.ini`, depending on your system).

- The patch for [Bug#8303](http://bugs.mysql.com/8303) [http://bugs.mysql.com/8303] broke the fix for [Bug#8378](http://bugs.mysql.com/8378) [http://bugs.mysql.com/8378] and was undone. (In string literals with an escape character (`\`) followed by a multi-byte character that has a second byte of (`\`), the literal was not interpreted correctly. The next byte now is escaped, not the entire multi-byte character. This means it a strict reverse of the `mysql_real_escape_string()` function.)
- The dropping of a temporary table whose name contained a backtick (```) character was not correctly replicated. ([Bug#19188](http://bugs.mysql.com/19188) [http://bugs.mysql.com/19188])
- The client libraries had not been compiled for position-independent code on Solaris-SPARC and AMD x86\_64 platforms. ([Bug#13159](http://bugs.mysql.com/13159) [http://bugs.mysql.com/13159], [Bug#14202](http://bugs.mysql.com/14202) [http://bugs.mysql.com/14202], [Bug#18091](http://bugs.mysql.com/18091) [http://bugs.mysql.com/18091])
- Running `myisampack` followed by `myisamchk` with the `--unpack` option would corrupt the `auto_increment` key. ([Bug#12633](http://bugs.mysql.com/12633) [http://bugs.mysql.com/12633])
- RPM packages had spurious dependencies on Perl modules and other programs. ([Bug#13634](http://bugs.mysql.com/13634) [http://bugs.mysql.com/13634])

## D.1.4. Changes in release 4.1.19 (29 April 2006)

This release includes the patches for recently reported security vulnerabilities in the MySQL client-server protocol. We would like to thank Stefano Di Paola <[stefano.dipaola@wisec.it](mailto:stefano.dipaola@wisec.it)> for finding

and reporting these to us.

Functionality added or changed:

- **Security enhancement:** Added the global `max_prepared_stmt_count` system variable to limit the total number of prepared statements in the server. This limits the potential for denial-of-service attacks based on running the server out of memory by preparing huge numbers of statements. The current number of prepared statements is available through the `prepared_stmt_count` system variable. (Bug#16365 [<http://bugs.mysql.com/16365>])
- The `MySQL-shared-compat-4.1.X-.i386.rpm` shared compatibility RPMs no longer contain libraries for MySQL 5.0 and up. It contains libraries for 3.23, 4.0, and 4.1.1. (Bug#19288 [<http://bugs.mysql.com/19288>])
- Creating a table in an InnoDB database with a column name that matched the name of an internal InnoDB column (including `DB_ROW_ID`, `DB_TRX_ID`, `DB_ROLL_PTR` and `DB_MIX_ID`) would cause a crash. MySQL now returns error 1005 (cannot create table) with `errno` set to -1. (Bug#18934 [<http://bugs.mysql.com/18934>])
- InnoDB now caches a list of unflushed files instead of scanning for unflushed files during a table flush operation. This improves performance when `--innodb-file-per-table` is set on a system with a large number of InnoDB tables. (Bug#15653 [<http://bugs.mysql.com/15653>])
- New `charset` command added to `mysql` command-line client. By typing `charset name` or `\C name` (such as `\C UTF8`), the client character set can be changed without reconnecting. (Bug#16217 [<http://bugs.mysql.com/16217>])
- Large file support was re-enabled for the MySQL server binary for the AIX 5.2 platform. (Bug#13571 [<http://bugs.mysql.com/13571>])
- When using the `GROUP_CONCAT()` function where the `group_concat_max_len` system variable was greater than 512, the type of the result was `BLOB` only if the query included an `ORDER BY` clause; otherwise the result was a `VARCHAR`.

The result type of the `GROUP_CONCAT()` function is now `VARCHAR` only if the value of the `group_concat_max_len` system variable is less than or equal to 512. Otherwise, this function returns a `BLOB`. (Bug#14169 [<http://bugs.mysql.com/14169>])

Bugs fixed:

- **Security fix:** A malicious client, using specially crafted invalid login or `COM_TABLE_DUMP` packets was able to read uninitialized memory, which potentially, though unlikely in MySQL, could have led to an information disclosure. (CVE-2006-1516 [<http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2006-1516>], CVE-2006-1517 [<http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2006-1517>]) Thanks to Stefano Di Paola <[stefano.dipaola@wisec.it](mailto:stefano.dipaola@wisec.it)> for finding and reporting this bug.
- **NDB Cluster:** A simultaneous `DROP TABLE` and table update operation utilising a table scan could trigger a node failure. (Bug#18597 [<http://bugs.mysql.com/18597>])
- `MySQL-shared-compat-4.1.15-0.i386.rpm`, `MySQL-shared-compat-4.1.16-0.i386.rpm`, and `MySQL-shared-compat-4.1.18-0.i386.rpm` incorrectly depended on `glibc` 2.3 and could not be installed on a `glibc` 2.2 system. (Bug#16539 [<http://bugs.mysql.com/16539>])
- IA-64 RPM packages for Red Hat and SuSE Linux that were built with the `icc` compiler incorrectly depended on `icc` runtime libraries. (Bug#16662 [<http://bugs.mysql.com/16662>])

- Index prefixes for `utf8 VARCHAR` columns did not work for `UPDATE` statements. (Bug#19080 [http://bugs.mysql.com/19080])
- MySQL would not compile on Linux distributions that use the `tinfo` library. (Bug#18912 [http://bugs.mysql.com/18912])
- **NDB Cluster**: Backups could fail for large clusters with many tables, where the number of tables approached `MaxNoOfTables`. (Bug#17607 [http://bugs.mysql.com/17607])
- For single-`SELECT` union constructs of the form `(SELECT ... ORDER BY order_list1 [LIMIT n]) ORDER BY order_list2`, the `ORDER BY` lists were concatenated and the `LIMIT` clause was ignored. (Bug#18767 [http://bugs.mysql.com/18767])
- The `IN-to-EXISTS` transformation was making a reference to a parse tree fragment that was left out of the parse tree. This caused problems with prepared statements. (Bug#18492 [http://bugs.mysql.com/18492])
- Attempting to set the default value of an `ENUM` or `SET` column to `NULL` caused a server crash. (Bug#19145 [http://bugs.mysql.com/19145])
- Index corruption could occur in cases when `key_cache_block_size` was not a multiple of `myisam_block_size` (for example, with `key_cache_block_size=1536` and `myisam_block_size=1024`). (Bug#19079 [http://bugs.mysql.com/19079])
- `UNCOMPRESS(NULL)` could cause subsequent `UNCOMPRESS()` calls to return `NULL` for legal non-`NULL` arguments. (Bug#18643 [http://bugs.mysql.com/18643])
- Conversion of a number to a `CHAR UNICODE` string returned an invalid result. (Bug#18691 [http://bugs.mysql.com/18691])
- A call to `MIN()` with a `CASE` expression as its argument could return a non-minimum value. (Bug#17896 [http://bugs.mysql.com/17896])
- A `LOCK TABLES` statement that failed could cause `MyISAM` not to update table statistics properly, causing a subsequent `CHECK TABLE` to report table corruption. (Bug#18544 [http://bugs.mysql.com/18544])
- Avoid trying to include `<asm/atomic.h>` when it doesn't work in C++ code. (Bug#13621 [http://bugs.mysql.com/13621])
- Executing `SELECT` on a large table that had been compressed within `mysampack` could cause a crash. (Bug#17917 [http://bugs.mysql.com/17917])
- **NDB Cluster**: In a 2-node cluster with a node failure, restarting the node with a low value for `StartPartialTimeout` could cause the cluster to come up partitioned (“split-brain” issue). (Bug#16447 [http://bugs.mysql.com/16447])  
  
A similar issue could occur when the cluster was first started with a sufficiently low value for this parameter. (Bug#18612 [http://bugs.mysql.com/18612])
- **NDB Cluster**: On systems with multiple network interfaces, data nodes would get “stuck” in startup phase 2 if the interface connecting them to the management server was working on node startup while the interface interconnecting the data nodes experienced a temporary outage. (Bug#15695 [http://bugs.mysql.com/15695])
- `mysql_config` returned incorrect libraries on `x86_64` systems. (Bug#13158 [http://bugs.mysql.com/13158])
- `mysql_reconnect()` sent a `SET NAMES` statement to the server, even for pre-4.1 servers that

do not understand the statement. ([Bug#18830](http://bugs.mysql.com/18830) [<http://bugs.mysql.com/18830>])

- During conversion from one character set to `ucs2`, multi-byte characters with no `ucs2` equivalent were converted to multiple characters, rather than to `0x003F QUESTION MARK`. ([Bug#15375](http://bugs.mysql.com/15375) [<http://bugs.mysql.com/15375>])
- The `mysql_close()` C API function leaked handles for share-memory connections on Windows. ([Bug#15846](http://bugs.mysql.com/15846) [<http://bugs.mysql.com/15846>])
- The euro sign (€) was not stored correctly in columns using the `latin1_german1_ci` or `latin1_general_ci` collation. ([Bug#18321](http://bugs.mysql.com/18321) [<http://bugs.mysql.com/18321>])
- The server was always built as though `--with-extra-charsets=complex` had been specified. ([Bug#12076](http://bugs.mysql.com/12076) [<http://bugs.mysql.com/12076>])
- `SELECT ... WHERE column LIKE 'A%'`, when `column` had a key and used the `latin2_czech_cs` collation, caused the wrong number of rows to be returned. ([Bug#17374](http://bugs.mysql.com/17374) [<http://bugs.mysql.com/17374>])
- A query using `WHERE (column_1, column_2) IN ((value_1, value_2)[, (..., ...), ...])` would return incorrect results. ([Bug#16248](http://bugs.mysql.com/16248) [<http://bugs.mysql.com/16248>])
- The `-lmtmalloc` library was removed from the output of `mysql_config` on Solaris, as it caused problems when building `DBD::mysql` (and possibly other applications) on that platform that tried to use `dlopen()` to access the client library. ([Bug#18322](http://bugs.mysql.com/18322) [<http://bugs.mysql.com/18322>])
- When running a query that contained a `GROUP_CONCAT( SELECT GROUP_CONCAT(...) )`, the result was `NULL` except in the `ROLLUP` part of the result, if there was one. ([Bug#15560](http://bugs.mysql.com/15560) [<http://bugs.mysql.com/15560>])
- `CAST(double AS SIGNED INT)` for large `double` values outside the signed integer range truncates the result to be within range, but the result sometimes had the wrong sign. ([Bug#15098](http://bugs.mysql.com/15098) [<http://bugs.mysql.com/15098>])
- `SET` value definitions containing commas were not rejected. Now a definition such as `SET( 'a,b', 'c,d' )` results in an error. ([Bug#15316](http://bugs.mysql.com/15316) [<http://bugs.mysql.com/15316>])
- **MyISAM:** Keys for which the first part of the key was a `CHAR` or `VARCHAR` column using the UTF-8 character set and longer than 254 bytes could become corrupted. ([Bug#17705](http://bugs.mysql.com/17705) [<http://bugs.mysql.com/17705>])
- **NDB Cluster:** A timeout in the handling of an `ABORT` condition with more than 32 operations could yield a node failure. ([Bug#18414](http://bugs.mysql.com/18414) [<http://bugs.mysql.com/18414>])
- **NDB Cluster:** A node restart immediately following a `CREATE TABLE` would fail. **Important:** This fix supports 2-node Clusters only. ([Bug#18385](http://bugs.mysql.com/18385) [<http://bugs.mysql.com/18385>])
- **NDB Cluster:** In event of a node failure during a rollback, a “false” lock could be established on the backup for that node, which lock could not be removed without restarting the node. ([Bug#18352](http://bugs.mysql.com/18352) [<http://bugs.mysql.com/18352>])
- **NDB Cluster:** The cluster created a crashed replica of a table having an ordered index — or when logging was not enabled, of a table having a table or unique index — leading to a crash of the cluster following 8 successive restarts. ([Bug#18298](http://bugs.mysql.com/18298) [<http://bugs.mysql.com/18298>])
- **NDB Cluster:** When replacing a failed master node, the replacement node could cause the cluster to crash from a buffer overflow if it had an excessively large amount of data to write to the cluster log. ([Bug#18118](http://bugs.mysql.com/18118) [<http://bugs.mysql.com/18118>])

- **NDB Cluster**: Restarting nodes were allowed to start and join the cluster too early. ([Bug#16772](http://bugs.mysql.com/16772) [<http://bugs.mysql.com/16772>])
- If **InnoDB** encountered a `HA_ERR_LOCK_TABLE_FULL` error and rolled back a transaction, the transaction was still written to the binary log. ([Bug#18283](http://bugs.mysql.com/18283) [<http://bugs.mysql.com/18283>])
- Connecting to a server with a UCS2 default character set with a client using a non-UCS2 character set crashed the server. ([Bug#18004](http://bugs.mysql.com/18004) [<http://bugs.mysql.com/18004>])
- Character set conversion of string constants for **UNION** of constant and table column was not done when it was safe to do so. ([Bug#15949](http://bugs.mysql.com/15949) [<http://bugs.mysql.com/15949>])
- Use of `TRUNCATE TABLE` for a **TEMPORARY** table on a master server was propagated to slaves properly, but slaves did not decrement the `Slave_open_temp_tables` counter properly. ([Bug#17137](http://bugs.mysql.com/17137) [<http://bugs.mysql.com/17137>])
- `SELECT COUNT( *)` for a **MyISAM** table could return different results depending on whether an index was used. ([Bug#14980](http://bugs.mysql.com/14980) [<http://bugs.mysql.com/14980>])
- Large file support did not work in AIX server binaries. ([Bug#10776](http://bugs.mysql.com/10776) [<http://bugs.mysql.com/10776>])
- Security Improvement: GRANTS to users with wildcards in their host information could be erroneously applied to similar users with the same username and similar wildcards. For example, a privilege granted to `foo@%` are also applied to user `foo@192.%`. ([Bug#14385](http://bugs.mysql.com/14385) [<http://bugs.mysql.com/14385>])
- **NDB Cluster**: Inserting and deleting **BLOB** column values while a backup was in process could cause the loss of an `ndbd` node. ([Bug#14028](http://bugs.mysql.com/14028) [<http://bugs.mysql.com/14028>])
- Setting the `myisam_repair_threads` system variable to a value larger than 1 could cause corruption of large **MyISAM** tables. ([Bug#11527](http://bugs.mysql.com/11527) [<http://bugs.mysql.com/11527>])
- Security improvement: In grant table comparisons, improper use of a `latin1` collation caused some hostname matches to be true that should have been false. Thanks to Deomid Ryabkov for finding this bug and proposing a solution. ([Bug#15756](http://bugs.mysql.com/15756) [<http://bugs.mysql.com/15756>])
- **NDB Cluster**: `ndb_delete_all` would run out of memory on tables containing **BLOB** columns. ([Bug#16693](http://bugs.mysql.com/16693) [<http://bugs.mysql.com/16693>])
- `mysqldump` tried to dump data from a view. (In MySQL 4.1, this applies when connecting to a server from MySQL 5.0 or higher.) ([Bug#16389](http://bugs.mysql.com/16389) [<http://bugs.mysql.com/16389>])
- **NDB Cluster**: An `UPDATE` with an inner join failed to match any records if both tables in the join did not have a primary key. ([Bug#17257](http://bugs.mysql.com/17257) [<http://bugs.mysql.com/17257>])
- **NDB Cluster**: A `DELETE` with a join in the `WHERE` clause failed to retrieve any records if both tables in the join did not have a primary key. ([Bug#17249](http://bugs.mysql.com/17249) [<http://bugs.mysql.com/17249>])
- **NDB Cluster**: In some cases, `LOAD DATA INFILE` did not load all data into **NDB** tables. ([Bug#17081](http://bugs.mysql.com/17081) [<http://bugs.mysql.com/17081>])
- **NDB Cluster**: The `REDO` log would become corrupted (and thus unreadable) in some circumstances, due to a failure in the query handler. ([Bug#17295](http://bugs.mysql.com/17295) [<http://bugs.mysql.com/17295>])
- **NDB Cluster**: No error message was generated for setting `NoOfFragmentLogFiles` too low. ([Bug#13966](http://bugs.mysql.com/13966) [<http://bugs.mysql.com/13966>])
- **NDB Cluster**: No error message was generated for setting `MaxNoOfAttributes` too low. ([Bug#13965](http://bugs.mysql.com/13965) [<http://bugs.mysql.com/13965>])

- Binary distributions for Solaris contained files with group ownership set to the non-existing `wheel` group. Now the `bin` group is used. ([Bug#15562](#) [<http://bugs.mysql.com/15562>])
- Killing a long-running query containing a subquery could cause a server crash. ([Bug#14851](#) [<http://bugs.mysql.com/14851>])
- Repeated invocation of `my_init()` and `my_end()` caused corruption of character set data and connection failure. ([Bug#6536](#) [<http://bugs.mysql.com/6536>])
- A `FULLTEXT` query in a prepared statement could result in unexpected behavior. ([Bug#14496](#) [<http://bugs.mysql.com/14496>])
- A `FULLTEXT` query in a `UNION` could result in unexpected behavior. ([Bug#16893](#) [<http://bugs.mysql.com/16893>])
- Server crash when dropping InnoDB constraints named `TABLENAME_ibfk_0`. ([Bug#16387](#) [<http://bugs.mysql.com/16387>])
- Corrected race condition when dropping the adaptive hash index for a B-tree page in InnoDB. ([Bug#16582](#) [<http://bugs.mysql.com/16582>])
- `LOAD DATA FROM MASTER` produced invalid warnings and Packet out of order errors when the database already existed on the slave. ([Bug#15302](#) [<http://bugs.mysql.com/15302>])
- A key on a `MEMORY` table would sometimes fail to match a row. ([Bug#12796](#) [<http://bugs.mysql.com/12796>])
- `MYSQL_STMT` objects were not preserved following a connection reset. Attempting to operate on them afterwards caused the server to crash. ([Bug#12744](#) [<http://bugs.mysql.com/12744>])

## D.1.5. Changes in release 4.1.18 (27 January 2006)

Functionality added or changed:

- `NDB Cluster`: More descriptive warnings are now issued when inappropriate logging parameters are set in `config.ini`. (Formerly, the warning issued was simply Could not add logfile destination.) ([Bug#11331](#) [<http://bugs.mysql.com/11331>])
- `libmysqlclient` now uses versioned symbols with GNU ld. ([Bug#3074](#) [<http://bugs.mysql.com/3074>])

Bugs fixed:

- The length of a `VARCHAR()` column that used the `utf8` character set would increase each time the table was re-created in a stored procedure or prepared statement, eventually causing the `CREATE TABLE` statement to fail. ([Bug#13134](#) [<http://bugs.mysql.com/13134>])
- RPM packages had an incorrect `zlib` dependency. ([Bug#15223](#) [<http://bugs.mysql.com/15223>])
- `STR_TO_DATE(1, NULL)` caused a server crash. ([CVE-2006-3081](#) [<http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2006-3081>], [Bug#15828](#) [<http://bugs.mysql.com/15828>])
- The `--replicate-do` and `--replicate-ignore` options were not being enforced on multiple-table statements. ([Bug#15699](#) [<http://bugs.mysql.com/15699>], [Bug#16487](#))

[<http://bugs.mysql.com/16487>])

- Running out of disk space in the location specified by the `tmpdir` option resulted in incorrect error message. (Bug#14634 [<http://bugs.mysql.com/14634>])
- Test suite `func_math` test returned warnings when server not compiled with InnoDB support. (Bug#15429 [<http://bugs.mysql.com/15429>])
- The `MBROverlaps` GIS function returned incorrect results. (Bug#14320 [<http://bugs.mysql.com/14320>])
- A `CREATE TABLE ... SELECT ...` on an equation involving `DOUBLE` values could result in the table being created with columns too small to hold the equation result. (Bug#9855 [<http://bugs.mysql.com/9855>])
- `UPDATE` statement crashed multi-byte character set `FULLTEXT` index if update value was almost identical to initial value only differing in some spaces being changed to `&nbsp;`. (Bug#16489 [<http://bugs.mysql.com/16489>])
- Single table `UPDATE` statements without `ORDER BY` clauses which updated the same indexed column that was being filtered on were optimized with a full index scan instead of a more appropriate index range scan. (Bug#15935 [<http://bugs.mysql.com/15935>])
- A prepared statement created from a `SELECT ... LIKE` query (such as `PREPARE stmt1 FROM 'SELECT col_1 FROM tadd_test WHERE col_1 LIKE ?';`) would begin to produce erratic results after being executed repeatedly numerous (thousands) of times. (Bug#12734 [<http://bugs.mysql.com/12734>])

## D.1.6. Changes in release 4.1.17 (Not released)

Functionality added or changed:

- In the `latin5_turkish_ci` collation, the order of the characters `A WITH CIRCUMFLEX`, `I WITH CIRCUMFLEX`, and `U WITH CIRCUMFLEX` was changed. If you have used these characters in any indexed columns, you should rebuild those indexes. (Bug#13421 [<http://bugs.mysql.com/13421>])
- Support files for compiling with Visual Studio 6 have been removed. (Bug#15094 [<http://bugs.mysql.com/15094>])
- Internal `sha1_result` function renamed to `mysql_sha1_result` to prevent conflicts with other projects. (Bug#13944 [<http://bugs.mysql.com/13944>])

Bugs fixed:

- `CAST(... AS TIME)` operations returned different results when using versus not using prepared-statement protocol. (Bug#15805 [<http://bugs.mysql.com/15805>])
- Performing a `RENAME TABLE` on an InnoDB table when the server is started with the `-innodb-file-per-table` and the data directory is a symlink caused a server crash. (Bug#15991 [<http://bugs.mysql.com/15991>])
- Characters in the `gb2312` and `euckr` character sets which did not have Unicode mappings were truncated. (Bug#15377 [<http://bugs.mysql.com/15377>])



- Piping the `fill_help_tables.sql` file into `mysqld` resulted in a syntax error. ([Bug#15965](http://bugs.mysql.com/15965) [<http://bugs.mysql.com/15965>])
- **NDBCluster**: Upon the completion of a scan where a key request remained outstanding on the primary replica and a starting node died, the scan did not terminate. This caused incompleted error handling of the failed node. ([Bug#15908](http://bugs.mysql.com/15908) [<http://bugs.mysql.com/15908>])
- Using `CAST()` to convert values with long fractional and/or exponent parts to `TIME` returned wrong results. ([Bug#12440](http://bugs.mysql.com/12440) [<http://bugs.mysql.com/12440>])
- An `INSERT ... SELECT` statement between tables in a `MERGE` set can return errors when statement involves insert into child table from merge table or vice-versa. ([Bug#5390](http://bugs.mysql.com/5390) [<http://bugs.mysql.com/5390>])
- Certain permission management statements could create a `NULL` hostname for a user, resulting in a server crash. ([Bug#15598](http://bugs.mysql.com/15598) [<http://bugs.mysql.com/15598>])
- For **InnoDB** tables, using a column prefix for a `utf8` column in a primary key caused `Cannot find record` errors when attempting to locate records. ([Bug#14056](http://bugs.mysql.com/14056) [<http://bugs.mysql.com/14056>])
- Certain `CREATE TABLE ... AS ...` statements involving `ENUM` columns could cause server crash. ([Bug#12913](http://bugs.mysql.com/12913) [<http://bugs.mysql.com/12913>])
- Using an aggregate function as the argument for a `HAVING` clause would result in the aggregate function always returning `FALSE`. ([Bug#14274](http://bugs.mysql.com/14274) [<http://bugs.mysql.com/14274>])
- The `COALESCE()` function truncated data in a `TINYTEXT` column. ([Bug#15581](http://bugs.mysql.com/15581) [<http://bugs.mysql.com/15581>])
- **InnoDB**: Comparison of indexed `VARCHAR CHARACTER SET ucs2 COLLATE ucs2_bin` columns using `LIKE` could fail. ([Bug#14583](http://bugs.mysql.com/14583) [<http://bugs.mysql.com/14583>])
- Issuing a `DROP USER` command could cause some users to encounter a `hostname is not allowed to connect to this MySQL server` error. ([Bug#15775](http://bugs.mysql.com/15775) [<http://bugs.mysql.com/15775>])
- `Access Denied` error could be erroneously returned with specific grant combinations under high load. ([Bug#7209](http://bugs.mysql.com/7209) [<http://bugs.mysql.com/7209>])
- Symbolic links did not function properly on Windows platforms. ([Bug#14960](http://bugs.mysql.com/14960) [<http://bugs.mysql.com/14960>], [Bug#14310](http://bugs.mysql.com/14310) [<http://bugs.mysql.com/14310>])
- **BDB**: A `DELETE`, `INSERT`, or `UPDATE` of a **BDB** table could cause the server to crash where the query contained a subquery using an index read. ([Bug#15536](http://bugs.mysql.com/15536) [<http://bugs.mysql.com/15536>])
- `DELETE` could report full-text index corruption (`Invalid key for table ...`) if the index was built with repair-by-sort, the data in the full-text index used `UCA` collation, and some word appeared in the data terminated by a `0xC2A0` character as well as by other non-letter characters. ([Bug#11336](http://bugs.mysql.com/11336) [<http://bugs.mysql.com/11336>])
- A race condition when creating temporary files caused a deadlock on Windows with threads in `Opening tables` or `Waiting for table` states. ([Bug#12071](http://bugs.mysql.com/12071) [<http://bugs.mysql.com/12071>])
- **InnoDB**: If `FOREIGN_KEY_CHECKS` was 0, **InnoDB** allowed inconsistent foreign keys to be created. ([Bug#13778](http://bugs.mysql.com/13778) [<http://bugs.mysql.com/13778>])
- **NDB Cluster**: A memory leak occurred when performing ordered index scans using indexes a

columns larger than 32 bytes, which would eventually lead to the forced shutdown of all `mysqld` server processes used with the cluster. (Bug#13078 [<http://bugs.mysql.com/13078>])

- **NDB Cluster:** Under some circumstances, it was possible for a restarting node to undergo a forced shutdown. (Bug#15632 [<http://bugs.mysql.com/15632>])
- **NDB Cluster:** If an abort by the Transaction Coordinator timed out, the abort condition was incorrectly handled, causing the transaction record to be released prematurely. (Bug#15685 [<http://bugs.mysql.com/15685>])
- **NDB Cluster:** A node which failed during cluster startup was sometimes not removed from the internal list of active nodes. (Bug#15587 [<http://bugs.mysql.com/15587>])
- **NDB Cluster:** There was a small window for a node failure to occur during a backup without an error being reported. (Bug#15425 [<http://bugs.mysql.com/15425>])
- Multiple-table update operations were counting updates and not updated rows. As a result, if a row had several updates it was counted several times for the “rows matched” value but updated only once. (Bug#15028 [<http://bugs.mysql.com/15028>])
- **SELECT** queries that began with an opening parenthesis were not being placed in the query cache. (Bug#14652 [<http://bugs.mysql.com/14652>])

## D.1.7. Changes in release 4.1.16 (29 November 2005)

Functionality added or changed:

- The `CHAR ( )` function now takes an optional `USING charset` clause that may be used to produce a result in a specific character set rather than in the connection character set.
- MySQL 4.1 now supports character set conversion for seven additional `cp950` characters into the `big5` character set: `0xF9D6`, `0xF9D7`, `0xF9D8`, `0xF9D9`, `0xF9DA`, `0xF9DB`, and `0xF9DC`. **Note:** If you move data containing these additional characters to an older MySQL installation which does not support them, you may encounter errors. (Bug#12476 [<http://bugs.mysql.com/12476>])
- **NDBCluster:** The `pererror` utility included with the `MySQL-Server` RPM now provides support for the `--ndb` option, and so can be used to obtain error message text for MySQL Cluster error codes. (Bug#13740 [<http://bugs.mysql.com/13740>])
- When executing single-table `UPDATE` or `DELETE` queries containing an `ORDER BY ... LIMIT N` clause, but not having any `WHERE` clause, MySQL can now take advantage of an index to read the first `N` rows in the ordering specified in the query. If an index is used, only the first `N` records will be read, as opposed to scanning the entire table. (Bug#12915 [<http://bugs.mysql.com/12915>])
- The `MySQL-server` RPM now explicitly assigns the `mysql` system user to the `mysql` user group during the postinstallation process. This corrects an issue with upgrading the server on some Linux distributions whereby a previously existing `mysql` user was not changed to the `mysql` group, resulting in wrong groups for files created following the installation. (Bug#12823 [<http://bugs.mysql.com/12823>])
- When a date column is set `NOT NULL` and contains `0000-00-00`, it will be updated for `UPDATE` statements that contains `columnname IS NULL` in the `WHERE` clause. (Bug#14186 [<http://bugs.mysql.com/14186>])

Bugs fixed:

- **NDB Cluster**: `REPLACE` failed when attempting to update a primary key value in a Cluster table. (Bug#14007 [http://bugs.mysql.com/14007])
- When the `DATE_FORMAT( )` function appeared in both the `SELECT` and `ORDER BY` clauses of a query but with arguments that differ by case (i.e. %m and %M), incorrect sorting may have occurred. (Bug#14016 [http://bugs.mysql.com/14016])
- `Make` failed when attempting to build MySQL in different directory than source. (Bug#11827 [http://bugs.mysql.com/11827])
- `PROCEDURE ANALYSE( )` could suggest a data type with a negative display width. (Bug#10716 [http://bugs.mysql.com/10716])
- **InnoDB**: During replication, There was a failure to record events in the binary log that still occurred even in the event of a `ROLLBACK`. For example, this sequence of commands:

```
BEGIN;
CREATE TEMPORARY TABLE t1 (a INT) ENGINE=INNODB;
ROLLBACK;
INSERT INTO t1 VALUES (1);
```

would succeed on the replication master as expected. However, the `INSERT` would fail on the slave because the `ROLLBACK` would (erroneously) cause the `CREATE TEMPORARY TABLE` statement not to be written to the binlog. (Bug#7947 [http://bugs.mysql.com/7947])

- **NDB Cluster**: Creating a table with packed keys failed silently. **NDB** now supports the `PACK_KEYS` option to `CREATE TABLE` correctly. (Bug#14514 [http://bugs.mysql.com/14514])
- `Non-latin1` object names were written with wrong character set to grant tables. (Bug#14406 [http://bugs.mysql.com/14406])
- Issuing `STOP SLAVE` after having acquired a global read lock with `FLUSH TABLES WITH READ LOCK` caused a deadlock. Now `STOP SLAVE` is generates an error in such circumstances. (Bug#10942 [http://bugs.mysql.com/10942])
- Portability fixes to support OpenSSL 0.9.8a. (Bug#14221 [http://bugs.mysql.com/14221])
- Closed a memory leak in the SSL code. (Bug#14780 [http://bugs.mysql.com/14780])
- Perform character set conversion of constant values whenever possible without data loss. (Bug#10446 [http://bugs.mysql.com/10446])
- A `UNION` of `DECIMAL` columns could produce incorrect results. (Bug#14216 [http://bugs.mysql.com/14216])
- **InnoDB**: Pad UTF-8 `VARCHAR` columns with `0x20`. Pad UCS2 `CHAR` columns with `0x0020`. (Bug#10511 [http://bugs.mysql.com/10511])
- Full-text indexing/searching failed for words that end with more than one apostrophe. (Bug#5686 [http://bugs.mysql.com/5686])
- Selecting from a table in both an outer query and a subquery could cause a server crash. (Bug#14482 [http://bugs.mysql.com/14482])
- Creating a table containing an `ENUM` or `SET` column from within a stored procedure or prepared statement caused a server crash later when executing the procedure or statement. (Bug#14410 [http://bugs.mysql.com/14410])
- `mysql_fix_privilege_tables.sql` contained an erroneous comment that resulted in an er-

- ror when the file contents were processed. ([Bug#14469](http://bugs.mysql.com/14469) [<http://bugs.mysql.com/14469>])
- On Windows, the server could crash during shutdown if both replication threads and normal client connection threads were active. (Re-fix of [Bug#11796](http://bugs.mysql.com/11796) [<http://bugs.mysql.com/11796>])
  - A `LIMIT`-related optimization failed to take into account that `MyISAM` table indexes can be disabled, causing Error 124 when it tried to use such an index. ([Bug#14616](http://bugs.mysql.com/14616) [<http://bugs.mysql.com/14616>])
  - For a table that had been opened with `HANDLER OPEN`, issuing `OPTIMIZE TABLE`, `ALTER TABLE`, or `REPAIR TABLE` caused a server crash. ([Bug#14397](http://bugs.mysql.com/14397) [<http://bugs.mysql.com/14397>])
  - `CREATE TABLE tbl_name (...) SELECT ...` could crash the server and write invalid data into the `.frm` file if the `CREATE TABLE` and `SELECT` both contained a column with the same name. Also, if a default value is specified in the column definition, it is now actually used. ([Bug#14480](http://bugs.mysql.com/14480) [<http://bugs.mysql.com/14480>])
  - For `MyISAM` tables, incorrect query results or incorrect updates could occur under these conditions: There is a multiple-column index that includes a `BLOB` column that is not the last column in the index, and the statement performs a lookup on the index using key column values that have `NULL` for the `BLOB` column and that provide values for all columns up to the `BLOB` column and at least the next column in the index. ([Bug#13814](http://bugs.mysql.com/13814) [<http://bugs.mysql.com/13814>])
  - Deletes from a `CSV` table could cause table corruption. ([Bug#14672](http://bugs.mysql.com/14672) [<http://bugs.mysql.com/14672>])
  - An update of a `CSV` table could cause a server crash. ([Bug#13894](http://bugs.mysql.com/13894) [<http://bugs.mysql.com/13894>])
  - `mysqld_safe` did not correctly start the `-max` version of the server (if it was present) if the `-ledir` option was given. ([Bug#13774](http://bugs.mysql.com/13774) [<http://bugs.mysql.com/13774>])
  - The endian byte in for spatial values in WKB format was not consistently respected. ([Bug#12839](http://bugs.mysql.com/12839) [<http://bugs.mysql.com/12839>])
  - An expression in an `ORDER BY` clause failed with `Unknown column 'col_name' in 'order clause'` if the expression referred to a column alias. ([Bug#11694](http://bugs.mysql.com/11694) [<http://bugs.mysql.com/11694>])
  - Statements of the form `CREATE TABLE ... SELECT ...` that created a column with a multi-byte character set could incorrectly calculate the maximum length of the column, resulting in a `Specified key was too long` error. ([Bug#14139](http://bugs.mysql.com/14139) [<http://bugs.mysql.com/14139>])
  - Use of `col_name = VALUES(col_name)` in the `ON DUPLICATE KEY UPDATE` clause of an `INSERT` statement failed with an `Column 'col_name' in field list is ambiguous` error. ([Bug#13392](http://bugs.mysql.com/13392) [<http://bugs.mysql.com/13392>])
  - On Windows, the server was not ignoring hidden or system directories that Windows may have created in the data directory, and would treat them as available databases. ([Bug#4375](http://bugs.mysql.com/4375) [<http://bugs.mysql.com/4375>])
  - `LIKE` operations did not work reliably for the `cp1250` character set. ([Bug#13347](http://bugs.mysql.com/13347) [<http://bugs.mysql.com/13347>])
  - Maximum values were handled incorrectly for command-line options of type `GET_LL`. ([Bug#12925](http://bugs.mysql.com/12925) [<http://bugs.mysql.com/12925>])
  - Use of `WITH ROLLUP PROCEDURE ANALYSE()` could hang the server. ([Bug#14138](http://bugs.mysql.com/14138) [<http://bugs.mysql.com/14138>])
  - `TIMEDIFF()`, `ADDTIME()`, and `STR_TO_DATE()` were not reporting that they could return

- NULL, so functions that invoked them might misinterpret their results. ([Bug#14009](#) [<http://bugs.mysql.com/14009>])
- The example configuration files supplied with MySQL distributions listed the `thread_cache_size` variable as `thread_cache`. ([Bug#13811](#) [<http://bugs.mysql.com/13811>])
  - `LOAD DATA INFILE` would not accept the same character for both the `ESCAPED BY` and the `ENCLOSED BY` clauses. ([Bug#11203](#) [<http://bugs.mysql.com/11203>])
  - `NDB Cluster`: Repeated transactions using unique index lookups could cause a memory leak leading to error 288, `Out of index operations in transaction coordinator`. ([Bug#14199](#) [<http://bugs.mysql.com/14199>])
  - `SELECT DISTINCT CHAR(col_name)` returned incorrect results after `SET NAMES utf8`. ([Bug#13233](#) [<http://bugs.mysql.com/13233>])
  - Character set conversion was not being done for `FIND_IN_SET()`. ([Bug#13751](#) [<http://bugs.mysql.com/13751>])
  - The default value of `query_prealloc_size` was set to 8192, lower than its minimum of 16384. The minimum has been lowered to 8192. ([Bug#13334](#) [<http://bugs.mysql.com/13334>])
  - The server did not take character set into account in checking the width of the `mysql.user.Password` column. As a result, it could incorrectly generate long password hashes even if the column was not long enough to hold them. ([Bug#13064](#) [<http://bugs.mysql.com/13064>])
  - `CAST(1E+300 TO SIGNED INT)` produced an incorrect result on little-endian machines. ([Bug#13344](#) [<http://bugs.mysql.com/13344>])
  - `mysqladmin` and `mysqldump` would hang on SCO OpenServer. ([Bug#13238](#) [<http://bugs.mysql.com/13238>])
  - Specifying `--default-character-set=cp-932` for `mysqld` would cause SQL scripts containing comments written using that character set to fail with a syntax error. ([Bug#13487](#) [<http://bugs.mysql.com/13487>])
  - Given a column `col_name` defined as `NOT NULL`, a `SELECT ... FROM ... WHERE col_name IS NULL` query following `SHOW TABLE STATUS` would erroneously return a non-empty result. ([Bug#13535](#) [<http://bugs.mysql.com/13535>])
  - Corrected a memory-copying problem for `big5` values when using `icc` compiler on Linux IA-64 systems. ([Bug#10836](#) [<http://bugs.mysql.com/10836>])
  - On BSD systems, the system `crypt()` call could return an error for some salt values. The error was not handled, resulting in a server crash. ([Bug#13619](#) [<http://bugs.mysql.com/13619>])
  - Character set file parsing during `mysql_real_connect()` read past the end of a memory buffer. ([Bug#6413](#) [<http://bugs.mysql.com/6413>])
  - The `--interactive-timeout` and `--slave-net-timeout` options for `mysqld` were not being obeyed on Mac OS X and other BSD-based platforms. ([Bug#8731](#) [<http://bugs.mysql.com/8731>])
  - Queries of the form `(SELECT ...) ORDER BY ...` were being treated as a `UNION`. This improperly resulted in only distinct values being returned (because `UNION` by default eliminates duplicate results). Also, references to column aliases in `ORDER BY` clauses following parenthesized `SELECT` statements were not resolved properly. ([Bug#7672](#) [<http://bugs.mysql.com/7672>])

- Multiple update queries using any type of subquery would be ignored by a replication slave when a condition such as `--replicate-ignore-table` like condition was used. (Bug#13236 [<http://bugs.mysql.com/13236>])
- An `UPDATE` query using a join would be executed incorrectly on a replication slave. (Bug#12618 [<http://bugs.mysql.com/12618>])
- `NDBCluster`: Placing multiple `[TCP DEFAULT]` sections in the cluster `config.ini` file crashed `ndb_mgmd`. (The process now exits gracefully with an appropriate error message.) (Bug#13611 [<http://bugs.mysql.com/13611>])
- Multiple race conditions existed in OpenSSL, particularly noticeable on Solaris. (Bug#9270 [<http://bugs.mysql.com/9270>])
- With `--log-slave-updates` `Exec_master_log_pos` of SQL thread lagged IO (Bug#13023 [<http://bugs.mysql.com/13023>])
- `PURGE MASTER LOGS` statement that used subquery for date crashed server. (Bug#10308 [<http://bugs.mysql.com/10308>])

## D.1.8. Changes in release 4.1.15 (13 October 2005)

Functionality added or changed:

The limit of 255 characters on the input buffer for `mysql` on Windows has been lifted. The exact limit depends on what the system allows, but can be up to 64K characters. A typical limit is 16K characters. (Bug#12929 [<http://bugs.mysql.com/12929>])

Added the `myisam_stats_method`, which controls whether `NULL` values in indexes are considered the same or different when collecting statistics for `MyISAM` tables. This influences the query optimizer as described in Section 7.4.7, “`MyISAM` Index Statistics Collection”. (Bug#12232 [<http://bugs.mysql.com/12232>])

- Better detection of connection timeout for replication servers on Windows allows elimination of extraneous `Lost connection` errors in the error log. (Bug#5588 [<http://bugs.mysql.com/5588>])
- When using `IF NOT EXISTS` with `CREATE DATABASE` or `CREATE TABLE`, a warning now is generated if the database or table already exists. (Bug#6008 [<http://bugs.mysql.com/6008>])
- A new command line argument was added to `mysqld` to ignore client character set information sent during handshake, and use server side settings instead, to reproduce 4.0 behavior (Bug#9948 [<http://bugs.mysql.com/9948>]):

```
mysqld --skip-character-set-client-handshake
```

Bugs fixed:

- A prepared statement failed with `Illegal mix of collations` if the client character set was `utf8` and the statement used a table that had a character set of `latin1`. (Bug#12371 [<http://bugs.mysql.com/12371>])
- If special characters such as `'_'`, `'%'`, or the escape character were included within the prefix of a column index, `LIKE` pattern matching on the indexed column did not return the correct result. (Bug#13046 [<http://bugs.mysql.com/13046>], Bug#13919 [<http://bugs.mysql.com/13919>])

- **NDBCluster**: Updating a text-type column during a cluster backup could cause the `ndbd` process to crash, due to the incorrect use of charset-normalized reads for. This could also lead to “wrong” data in the backup if such a column was updated during the backup; for example, supposing that the column used `latin_ci`, then “aAa” might be stored in the backup as “AAA”. (Bug#12950 [<http://bugs.mysql.com/12950>])
- **NDBCluster**: When performing a delete of a great many (tens of thousands of) rows at once from a Cluster table, an improperly dereferenced pointer could cause the `mysqld` process to crash. (Bug#9282 [<http://bugs.mysql.com/9282>])
- Server could over-allocate memory when performing a `FULLTEXT` search for stopwords only. (Bug#13582 [<http://bugs.mysql.com/13582>])
- `UNION` of two `DECIMAL` columns returned wrong field type. (Bug#13372 [<http://bugs.mysql.com/13372>])
- `CHECKSUM TABLE` locked `InnoDB` tables and did not use a consistent read. (Bug#12669 [<http://bugs.mysql.com/12669>])
- `SHOW CREATE TABLE` did not display any `FOREIGN KEY` clauses if a temporary file could not be created. Now `SHOW CREATE TABLE` displays an error message in an SQL comment if this occurs. (Bug#13002 [<http://bugs.mysql.com/13002>])
- Display of the `AUTO_INCREMENT` attribute by `SHOW CREATE TABLE` was not controlled by the `NO_FIELD_OPTIONS` SQL mode as it should have been. (Bug#7977 [<http://bugs.mysql.com/7977>])
- For `VARCHAR` columns with the `ucs2` character set, `InnoDB` trimmed trailing `0x20` bytes rather than `0x0020` words, resulting in incorrect index lookups later. (Bug#12178 [<http://bugs.mysql.com/12178>])
- MySQL programs in binary distributions for Solaris 8/9/10 x86 systems would not run on Pentium III machines. (Bug#6772 [<http://bugs.mysql.com/6772>])
- **NDB Cluster**: Multiple `ndb_mgmd` processes in a cluster would not know each other's IP addresses. (Bug#12037 [<http://bugs.mysql.com/12037>])
- **NDB Cluster**: With two `mgmd` processes in a cluster, `ndb_mgm` output for `SHOW` would display the same IP address for both processes, even when they were on different hosts. (Bug#11595 [<http://bugs.mysql.com/11595>])
- The `--replicate-rewrite-db` and `--replicate-do-table` options did not work for statements in which tables were aliased to names other than those listed by the options. (Bug#11139 [<http://bugs.mysql.com/11139>])
- After running `configure` with the `--with-embedded-privilege-control` option, the embedded server failed to build. (Bug#13501 [<http://bugs.mysql.com/13501>])
- Queries against a `MERGE` table that has a composite index could produce incorrect results. (Bug#9112 [<http://bugs.mysql.com/9112>])
- Comparisons involving row constructors containing constants could cause a server crash. (Bug#13356 [<http://bugs.mysql.com/13356>])
- MySQL would pass an incorrect key length to storage engines for `MIN()`. This could cause warnings such as `InnoDB: Warning: using a partial-field key prefix in search.` in the `.err` log. (Bug#13218 [<http://bugs.mysql.com/13218>], same as Bug#11039 [<http://bugs.mysql.com/11039>] in MySQL 5.0.7)

- `NDB Cluster: LOAD DATA INFILE` with a large data file failed. ([Bug#10694](http://bugs.mysql.com/10694))
- `NDB Cluster`: Adding an index to a table with a large number of columns (more than 100) crashed the storage node. ([Bug#13316](http://bugs.mysql.com/13316))
- Aggregate functions sometimes incorrectly were allowed in the `WHERE` clause of `UPDATE` and `DELETE` statements. ([Bug#13180](http://bugs.mysql.com/13180))
- `MIN()` and `MAX()` sometimes returned a non-`NULL` value for an empty row set (for example, `SELECT MAX(1) FROM empty_table`). ([Bug#12882](http://bugs.mysql.com/12882))
- `LOAD DATA INFILE` did not respect the `NO_AUTO_VALUE_ON_ZERO` SQL mode setting. ([Bug#12053](http://bugs.mysql.com/12053))
- Use of a user-defined function within the `HAVING` clause of a query resulted in an `Unknown column` error. ([Bug#11553](http://bugs.mysql.com/11553))
- The data type for `DECIMAL` columns was not respected when updating the column from another column. For example, updating a `DECIMAL(10,1)` column with the value from a `DECIMAL(10,5)` column resulted in a `DECIMAL(10,5)` value being stored. Similarly, altering a column with a `DECIMAL(10,5)` datatype to a `DECIMAL(10,1)` data type did not properly convert data values. ([Bug#7598](http://bugs.mysql.com/7598))
- `NDB`: A cluster shutdown following the crash of a data node would fail to terminate the remaining node processes, even though `ndb_mgm` showed the shutdown request as having been completed. ([Bug#10938](http://bugs.mysql.com/10938), [Bug#9996](http://bugs.mysql.com/9996), [Bug#11623](http://bugs.mysql.com/11623))
- For queries with `DISTINCT` and `WITH ROLLUP`, the `DISTINCT` should be applied after the rollup operation, but was not always. ([Bug#12887](http://bugs.mysql.com/12887))
- The counters for the `Key_read_requests`, `Key_reads`, `Key_write_requests`, and `Key_writes` status variables were changed from `unsigned long` to `unsigned longlong` to accommodate larger values before the variables roll over and restart from 0. ([Bug#12920](http://bugs.mysql.com/12920))
- A column that can be `NULL` was not handled properly for `WITH ROLLUP` in a subquery or view. ([Bug#12885](http://bugs.mysql.com/12885))
- Shared-memory connections were not working on Windows. ([Bug#12723](http://bugs.mysql.com/12723))
- A concurrency problem for `CREATE ... SELECT` could cause a server crash. ([Bug#12845](http://bugs.mysql.com/12845))
- Performing an `IS NULL` check on the `MIN()` or `MAX()` of an indexed column in a complex query could produce incorrect results. ([Bug#12695](http://bugs.mysql.com/12695))
- The `NDB START BACKUP` command could be interrupted by a `SHOW` command. ([Bug#13054](http://bugs.mysql.com/13054))
- The `LIKE ... ESCAPE` syntax produced invalid results when escape character was larger than one byte. ([Bug#12611](http://bugs.mysql.com/12611))
- A client connection thread cleanup problem caused the server to crash when closing the connection if the binary log was enabled. ([Bug#12517](http://bugs.mysql.com/12517))
- `CHECKSUM TABLE` command returned incorrect results for tables with deleted rows. After upgrad-



- ing, users who used stored checksum information to detect table changes should rebuild their checksum data. (Bug#12296 [<http://bugs.mysql.com/12296>])
- The value of `character_set_results` could be set to `NULL`, but returned the string "NULL" when retrieved. (Bug#12363 [<http://bugs.mysql.com/12363>])
  - `GROUP_CONCAT()` ignored an empty string if it was the first value to occur in the result. (Bug#12863 [<http://bugs.mysql.com/12863>])
  - Outer join elimination was erroneously applied for some queries that used a `NOT BETWEEN` condition, an `IN(value_list)` condition, or an `IF()` condition. (Bug#12101 [<http://bugs.mysql.com/12101>], Bug#12102 [<http://bugs.mysql.com/12102>])
  - Reverted a change introduced in MySQL 4.1.13 to fix a problem of `SHOW FIELDS` truncating the `TYPE` column to 40 characters. This fix was reverted for MySQL 4.1 because it broke existing applications. The fix will be made to MySQL 5.0 instead (5.0.13). (Bug#7142 [<http://bugs.mysql.com/7142>], Bug#12817 [<http://bugs.mysql.com/12817>])
  - On HP-UX 11.x (PA-RISC), the `-L` option caused `mysqlimport` to crash. (Bug#12958 [<http://bugs.mysql.com/12958>])
  - After changing the character set with `SET CHARACTER SET`, the result of the `GROUP_CONCAT()` function was not converted to the proper character set. (Bug#12829 [<http://bugs.mysql.com/12829>])
  - `mysisampack` did not properly pack `BLOB` values larger than  $2^{24}$  bytes. (Bug#4214 [<http://bugs.mysql.com/4214>])
  - The server crashed when one thread resized the query cache while another thread was using it. (Bug#12848 [<http://bugs.mysql.com/12848>])
  - `mysqld_multi` now quotes arguments on command lines that it constructs to avoid problems with arguments that contain shell metacharacters. (Bug#11280 [<http://bugs.mysql.com/11280>])
  - When any `--replicate-wild-*` option is used, the slave ignores `SET ONE_SHOT TIME_ZONE` statements as belonging to a non-replicated table. (Bug#12542 [<http://bugs.mysql.com/12542>])
  - Deadlock occurred when several account management statements were run (particularly between `FLUSH PRIVILEGES/SET PASSWORD` and `GRANT/REVOKE` statements). (Bug#12423 [<http://bugs.mysql.com/12423>])
  - On Windows, the server was preventing tables from being created if the table name was a prefix of a forbidden name. For example, `nul` is a forbidden name because it's the same as a Windows device name, but a table with the name of `n` or `nu` was being forbidden as well. (Bug#12325 [<http://bugs.mysql.com/12325>])
  - `InnoDB` was too permissive with `LOCK TABLE ... READ LOCAL` and allowed new inserts into the table. Now `READ LOCAL` is equivalent to `READ` for `InnoDB`. This will cause slightly more locking in `mysqldump`, but makes `InnoDB` table dumps consistent with `MyISAM` table dumps. (Bug#12410 [<http://bugs.mysql.com/12410>])
  - The `have_innodb` read-only system variable could not be selected with `SELECT @@have_innodb`. (Bug#9613 [<http://bugs.mysql.com/9613>])
  - A `UNION` of long `utf8 VARCHAR` columns was sometimes returned as a column with a `LONG-TEXT` data type rather than `VARCHAR`. This could prevent such queries from working at all if selected into a `MEMORY` table because the `MEMORY` storage engine does not support the `TEXT` data types. (Bug#12537 [<http://bugs.mysql.com/12537>])

- **NDB Cluster**: Corrected the parsing of the `CLUSTERLOG` command by `ndb_mgm` to allow multiple items. (Bug#12833 [<http://bugs.mysql.com/12833>])
- **NDB Cluster**: Improved error messages related to filesystem issues. (Bug#11218 [<http://bugs.mysql.com/11218>])
- **NDB Cluster**: When a schema was detected to be corrupt, `ndb` neglected to close it, resulting in a “file already open” error if the schema was opened again later. written. (Bug#12027 [<http://bugs.mysql.com/12027>])
- **NDB Cluster**: When it could not copy a fragment, `ndbd` exited without printing a message about the condition to the error log. Now the message is written. (Bug#12900 [<http://bugs.mysql.com/12900>])
- **NDB Cluster**: When a disk full condition occurred, `ndbd` exited without printing a message about the condition to the error log. Now the message is written. (Bug#12716 [<http://bugs.mysql.com/12716>])
- If a client has opened an **InnoDB** table for which the `.ibd` file is missing, **InnoDB** would not honor a `DROP TABLE` statement for the table. (Bug#12852 [<http://bugs.mysql.com/12852>])
- `SELECT GROUP_CONCAT( . . . ) FROM DUAL` in a subquery could cause the client to hang. (Bug#12861 [<http://bugs.mysql.com/12861>])
- **NDB Cluster**: Bad values in `config.ini` caused `ndb_mdmd` to crash. (Bug#12043 [<http://bugs.mysql.com/12043>])
- `TRUNCATE TABLE` did not work with **TEMPORARY InnoDB** tables. (Bug#11816 [<http://bugs.mysql.com/11816>])
- `ALTER TABLE db_name.t RENAME t` did not move the table to default database unless the new name was qualified with the database name. (Bug#11493 [<http://bugs.mysql.com/11493>])
- Spatial index corruption could occur during updates. (Bug#9645 [<http://bugs.mysql.com/9645>])
- If a `DROP DATABASE` fails on a master server due to the presence of a non-database file in the database directory, the master have the database tables deleted, but not the slaves. To deal with failed database drops, we now write `DROP TABLE` statements to the binary log for the tables so that they are dropped on slaves. (Bug#4680 [<http://bugs.mysql.com/4680>])
- `DELETE` or `UPDATE` for an indexed **MyISAM** table could fail. This was due to a change in end-space comparison behavior from 4.0 to 4.1. (Bug#12565 [<http://bugs.mysql.com/12565>])
- The **ARCHIVE** storage engine does not support deletes, but it was possible to delete by using `DELETE` or `TRUNCATE TABLE` with a **FEDERATED** table that points to an **ARCHIVE** table. (Bug#12836 [<http://bugs.mysql.com/12836>])
- Queries that created implicit temporary tables could return incorrect data types for some columns. (Bug#11718 [<http://bugs.mysql.com/11718>])
- An optimizer estimate of zero rows for a non-empty **InnoDB** table used in a left or right join could cause incomplete rollback for the table. (Bug#12779 [<http://bugs.mysql.com/12779>])
- A `SELECT DISTINCT` query with a constant value for one of the columns would return only a single row. (Bug#12625 [<http://bugs.mysql.com/12625>])
- Users created using an IP address or other alias rather than a hostname listed in `/etc/hosts` could not set their own passwords. (Bug#12302 [<http://bugs.mysql.com/12302>])

- **NDB Cluster**: An `ALTER TABLE` command caused loss of data stored prior to the issuing of the command. (Bug#12118 [<http://bugs.mysql.com/12118>])
- **MEMORY** tables using **B-Tree** index on 64-bit platforms could produce false table is full errors. (Bug#12460 [<http://bugs.mysql.com/12460>])
- MySQL failed to compile when `--with-ndb-ccflags` was specified. (Bug#11538 [<http://bugs.mysql.com/11538>])
- On Windows when the `--innodb_buffer_pool_ave_mem_mb` option has been given, the server detects whether AWE support is available and has been compiled into the server, and displays an appropriate error message if not. (Bug#6581 [<http://bugs.mysql.com/6581>])

## D.1.9. Changes in release 4.1.14 (17 August 2005)

Functionality added or changed:

- `SHOW CHARACTER SET` and `INFORMATION_SCHEMA` now properly report the `Latin1` character set as `cp1252`. (Bug#11216 [<http://bugs.mysql.com/11216>])
- **MySQL Cluster**: A new `-P` option is available for use with the `ndb_mgmd` client. When called with this option, `ndb_mgmd` prints all configuration data to `stdout`, then exits.
- The output of `perror --help` now displays the `--ndb` option. (Bug#11999 [<http://bugs.mysql.com/11999>])
- **NDB**: Improved handling of the configuration variables `NoOfPagesToDiskDuringRestartACC`, `NoOfPagesToDiskAfterRestartACC`, `NoOfPagesToDiskDuringRestartTUP`, and `NoOfPagesToDiskAfterRestartTUP` should result in noticeably faster start-up times for MySQL Cluster. (Bug#12149 [<http://bugs.mysql.com/12149>])
- Added support of where clause for queries with `FROM DUAL`. (Bug#11745 [<http://bugs.mysql.com/11745>])
- Added an optimization that avoids key access with `NULL` keys for the `ref` method when used in outer joins. (Bug#12144 [<http://bugs.mysql.com/12144>])
- Added new query cache test for the embedded server to the test suite, there are now specific tests for the embedded and non-embedded servers. (Bug#9508 [<http://bugs.mysql.com/9508>])

Bugs fixed:

- If a thread (connection) has tables locked, the query cache is switched off for that thread. This prevents invalid results where the locking thread inserts values between a second thread connecting and selecting from the table. (Bug#12385 [<http://bugs.mysql.com/12385>])
- Prepared statement parameters could cause errors in the binary log if the character set was `cp932`. (Bug#11338 [<http://bugs.mysql.com/11338>])
- Queries with subqueries that contain outer joins could return wrong results. (Bug#11479 [<http://bugs.mysql.com/11479>])
- Slave I/O threads were considered to be in the running state when launched (rather than after successfully connecting to the master server), resulting in incorrect `SHOW SLAVE STATUS` output. (Bug#10780 [<http://bugs.mysql.com/10780>])

- On Windows, the server could crash during shutdown if both replication threads and normal client connection threads were active. (Bug#11796 [http://bugs.mysql.com/11796])
- Some subqueries of the form `SELECT ... WHERE ROW(...) IN (subquery)` were being handled incorrectly. (Bug#11867 [http://bugs.mysql.com/11867])
- The `mysql_info()` C API function could return incorrect data when executed as part of a multi-statement that included a mix of statements that do and do not return information. (Bug#11688 [http://bugs.mysql.com/11688])
- Renamed the `rest()` macro in `my_list.h` to `list_rest()` to avoid name clashes with user code. (Bug#12327 [http://bugs.mysql.com/12327])
- `myisampack` failed to delete `.TMD` temporary files when run with `-T` option. (Bug#12235 [http://bugs.mysql.com/12235])
- Concatenating `USER()/DATABASE()` with a column produces invalid results. (Bug#12351 [http://bugs.mysql.com/12351])
- For PKG installs on Mac OS X, the preinstallation and postinstallation scripts were being run only for new installations and not for upgrade installations, resulting in an incomplete installation process. (Bug#11380 [http://bugs.mysql.com/11380])
- User variables were not automatically cast for comparisons, causing queries to fail if the column and connection character sets differed. Now when mixing strings with different character sets but the same coercibility, allow conversion if one character set is a superset of the other. (Bug#10892 [http://bugs.mysql.com/10892])
- Pathame values for options such as `---basedir` or `--datadir` didn't work on Japanese Windows machines for directory names containing multi-byte characters having a second byte of `0x5C` (`'\``). (Bug#5439 [http://bugs.mysql.com/5439])
- Mishandling of comparison for rows containing `NULL` values against rows produced by an `IN` subquery could cause a server crash. (Bug#12392 [http://bugs.mysql.com/12392])
- `INSERT ... SELECT ... ON DUPLICATE KEY UPDATE` could fail with an erroneous "Column '`col_name`' specified twice" error. (Bug#10109 [http://bugs.mysql.com/10109])
- `myisam.test` failed when server compiled using `--without-geometry` option. (Bug#11083 [http://bugs.mysql.com/11083])
- Creation of the `mysql` group account failed during the RPM installation. (Bug#12348 [http://bugs.mysql.com/12348])
- `FLUSH TABLES WITH READ LOCK` combined with `LOCK TABLE ... WRITE` caused deadlock. (Bug#9459 [http://bugs.mysql.com/9459])
- `GROUP_CONCAT` ignores the `DISTINCT` modifier when used in a query joining multiple tables where one of the tables has a single row. (Bug#12095 [http://bugs.mysql.com/12095])
- `UNION` query with `FULLTEXT` could cause server crash. (Bug#11869 [http://bugs.mysql.com/11869])
- Performing `DATE(LEFT(column, 8))` on a `DATE` column produces incorrect results. (Bug#12266 [http://bugs.mysql.com/12266])
- `max_connections_per_hour` setting was being capped by unrelated `max_user_connections` setting. (Bug#9947 [http://bugs.mysql.com/9947])

- `big5` strings were not being stored in `FULLTEXT` index. (Bug#12075 [http://bugs.mysql.com/12075])
- Updated dependency list for RPM builds to include missing dependencies such as `useradd` and `groupadd`. (Bug#12233 [http://bugs.mysql.com/12233])
- Multiplying `ABS()` output by a negative number would return incorrect results. (Bug#11402 [http://bugs.mysql.com/11402])
- `SELECT @@local...` returned `@@session...` in the column header. (Bug#10724 [http://bugs.mysql.com/10724])
- Character data truncated when GBK characters `0xA3A0` and `0xA1` are present. (Bug#11987 [http://bugs.mysql.com/11987])
- Comparisons like `SELECT "A\\" LIKE "A\\";` fail when using `SET NAMES utf8;`. (Bug#11754 [http://bugs.mysql.com/11754])
- Attempting to repair a table having a fulltext index on a column containing words whose length exceeded 21 characters and where `myisam_repair_threads` was greater than 1 would crash the server. (Bug#11684 [http://bugs.mysql.com/11684])
- Two threads could potentially initialize different characters sets and overwrite each other. (Bug#12109 [http://bugs.mysql.com/12109])
- `NDB`: Attempting to create or drop tables during a backup would cause the cluster to shut down. (Bug#11942 [http://bugs.mysql.com/11942])
- `NDB_MGMD` was leaking file descriptors. (Bug#11898 [http://bugs.mysql.com/11898])
- `NDB` ignored the `Hostname` option in the `NDBD DEFAULT` section of the `NDB` configuration file. (Bug#12028 [http://bugs.mysql.com/12028])
- The temporary tables created by an `ALTER TABLE` on a cluster table were visible to all MySQL servers. (Bug#12055 [http://bugs.mysql.com/12055])
- For prepared statements, the SQL parser did not disallow ‘?’ parameter markers immediately adjacent to other tokens, which could result in malformed statements in the binary log. (For example, `SELECT * FROM t WHERE? = 1` could become `SELECT * FROM t WHERE0 = 1.`) (Bug#11299 [http://bugs.mysql.com/11299])
- `GROUP_CONCAT()` sometimes returned a result with a different collation from that of its arguments. (Bug#10201 [http://bugs.mysql.com/10201])
- When two threads compete for the same table, a deadlock could occur if one thread has also a lock on another table through `LOCK TABLES` and the thread is attempting to remove the table in some manner and the other thread want locks on both tables. (Bug#10600 [http://bugs.mysql.com/10600])
- Incorrect error message displayed if user attempted to create a table in a non-existing database using `CREATE database_name.table_name` syntax. (Bug#10407 [http://bugs.mysql.com/10407])
- The `LPAD()` and `RPAD()` functions returned the wrong length to `mysql_fetch_fields()`. (Bug#11311 [http://bugs.mysql.com/11311])
- The C API function `mysql_statement_reset()` did not clear error information. (Bug#11183 [http://bugs.mysql.com/11183])
- Multiple-table `UPDATE` queries using `CONVERT_TZ()` would fail with an error. (Bug#9979 [http://bugs.mysql.com/9979])

- `mysql_fetch_fields()` returned incorrect length information for `MEDIUM` and `LONG TEXT` and `BLOB` columns. (Bug#9735 [<http://bugs.mysql.com/9735>])
- `LIKE` pattern matching using prefix index didn't return correct result. (Bug#11650 [<http://bugs.mysql.com/11650>])
- The MySQL server had issues with certain combinations of `basedir` and `datadir`. (Bug#7249 [<http://bugs.mysql.com/7249>])
- `mysql_next_result()` returns incorrect value if final query in a batch fails. (Bug#12001 [<http://bugs.mysql.com/12001>])
- `SHOW BINARY LOGS` displayed a file size of 0 for all log files but the current one if the files were not located in the data directory. (Bug#12004 [<http://bugs.mysql.com/12004>])
- Server-side prepared statements failed for columns with a character set of `ucs2`. (Bug#9442 [<http://bugs.mysql.com/9442>])
- References to system variables in an SQL statement prepared with `PREPARE` were evaluated during `EXECUTE` to their values at prepare time, not to their values at execution time. (Bug#9359 [<http://bugs.mysql.com/9359>])
- For server shutdown on Windows, error messages of the form `Forcing close of thread n user: 'name'` were being written to the error log. Now connections are closed more gracefully without generating error messages. (Bug#7403 [<http://bugs.mysql.com/7403>])
- Corrected a problem with the optimizer incorrectly adding `NOT NULL` constraints, producing in incorrect results for complex queries. (Bug#11482 [<http://bugs.mysql.com/11482>])
- Corrected an optimizer problem with `NOT NULL` constraints within a subquery in an `UPDATE` statement that resulted in a server crash. (Bug#11868 [<http://bugs.mysql.com/11868>])
- Creating a table with a `SET` or `ENUM` column with the `DEFAULT 0` clause caused a server crash if the table's character set was `utf8`. (Bug#11819 [<http://bugs.mysql.com/11819>])
- In SQL prepared statements, comparisons could fail for values not equally space-padded. For example, `SELECT 'a' = 'a '`; returns 1, but `PREPARE s FROM 'SELECT ?=?'; SET @a = 'a', @b = 'a '; PREPARE s FROM 'SELECT ?=?'; EXECUTE s USING @a, @b;` incorrectly returned 0. (Bug#9379 [<http://bugs.mysql.com/9379>])
- `InnoDB`: Do not flush after each write, not even before setting up the doublewrite buffer. Flushing can be extremely slow on some systems. (Bug#12125 [<http://bugs.mysql.com/12125>])
- `ISO-8601` formatted dates were not being parsed correctly. (Bug#7308 [<http://bugs.mysql.com/7308>])

## D.1.10. Changes in release 4.1.13 (15 July 2005)

Functionality added or changed:

- Security improvement: Applied a patch that addresses a `zlib` data vulnerability that could result in a buffer overflow and code execution. (CVE-2005-2096 [<http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2005-2096>]) (Bug#11844 [<http://bugs.mysql.com/11844>]). Shortly after MySQL 4.1.13 was released, a second potential `zlib` security flaw was discovered and fixed - the issue is tracked by the Mitre CVE ID (CVE-2005-1849 [<http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2005-1849>]). A patch for this flaw was ap-

plied on top of the 4.1.13 sources and published as 4.1.13a. The affected binaries have been rebuilt.

- Security improvement: Applied a patch to fix a UDF library-loading vulnerability that could result in a buffer overflow and code execution. (<http://www.appsecinc.com/resources/alerts/mysql/2005-002.html>)
- **Warning: Incompatible change:** Previously, conversion of `DATETIME` values to numeric form by adding zero produced a result in `YYYYMMDDHHMMSS` format. The result of `DATETIME+0` is now in `YYYYMMDDHHMMSS.000000` format. ([Bug#12268](http://bugs.mysql.com/12268) [<http://bugs.mysql.com/12268>])
- Added the `--add-drop-database` option to `mysqldump`. ([Bug#3716](http://bugs.mysql.com/3716) [<http://bugs.mysql.com/3716>])
- Added `mysql_set_character_set()` C API function for setting the default character set of the current connection. This allows clients to affect the character set used by `mysql_real_escape_string()`. ([Bug#8317](http://bugs.mysql.com/8317) [<http://bugs.mysql.com/8317>])
- `MEMORY` tables now support indexes of up to 500 bytes. See [Section 14.4, “The MEMORY \(HEAP\) Storage Engine”](#). ([Bug#10566](http://bugs.mysql.com/10566) [<http://bugs.mysql.com/10566>])
- System variables are now treated as having `SYSVAR` (system constant) coercibility. For example, `@@version` is now treated like `VERSION()` and `@@character_set_client` is now treated like `CHARSET( USER() )`. See [Section 10.5.4, “Some Special Cases Where the Collation Determination Is Tricky”](#). ([Bug#10904](http://bugs.mysql.com/10904) [<http://bugs.mysql.com/10904>])
- The statements `CREATE TABLE`, `TRUNCATE TABLE`, `DROP DATABASE`, and `CREATE DATABASE` cause an implicit commit. ([Bug#6883](http://bugs.mysql.com/6883) [<http://bugs.mysql.com/6883>])
- Added the `--log-slow-admin-statements` server option to request logging of slow administrative statements such as `OPTIMIZE TABLE`, `ANALYZE TABLE`, and `ALTER TABLE` to the slow query log. These statements were logged in MySQL 4.0, but not in 4.1. ([Bug#9141](http://bugs.mysql.com/9141) [<http://bugs.mysql.com/9141>])
- `SHOW BINARY LOGS` now displays a `File_size` column that indicates the size of each file.
- You can again refer to other tables in the `ON DUPLICATE KEY UPDATE` part of an `INSERT ... SELECT` statement as long as there is no `GROUP BY` in the `SELECT` part. One side effect of this is that you may have to qualify non-unique column names in the values part of `ON DUPLICATE KEY UPDATE`. ([Bug#9728](http://bugs.mysql.com/9728) [<http://bugs.mysql.com/9728>], [Bug#8147](http://bugs.mysql.com/8147) [<http://bugs.mysql.com/8147>])
- The `table`, `type`, and `rows` columns of `EXPLAIN` output can now be `NULL`. This is required for using `EXPLAIN` on `SELECT` queries that use no tables, such as `EXPLAIN SELECT 1`. ([Bug#9899](http://bugs.mysql.com/9899) [<http://bugs.mysql.com/9899>])
- Expanded on information provided in general log and slow query log for prepared statements. ([Bug#8367](http://bugs.mysql.com/8367) [<http://bugs.mysql.com/8367>], [Bug#9334](http://bugs.mysql.com/9334) [<http://bugs.mysql.com/9334>])
- **InnoDB:** When creating or extending an InnoDB data file, allocate at most one megabyte at a time for initializing the file. Previously, InnoDB used to allocate and initialize 1 or 8 megabytes of memory, even if a few 16-kilobyte pages were to be written. This fix improves the performance of `CREATE TABLE` in `innodb_file_per_table` mode.

Bugs fixed:

- **Security fix:** On Windows systems, a user with any of the following privileges

- [REFERENCES](#)
- [CREATE TEMPORARY TABLES](#)
- [GRANT OPTION](#)
- [CREATE](#)
- [SELECT](#)

on \*.\* could crash `mysqld` by issuing a `USE LPT1;` or `USE PRN;` command. In addition, any of the commands `USE NUL;`, `USE CON;`, `USE COM1;`, or `USE AUX;` would report success even though the database was not in fact changed. ([Bug#9148](#) [<http://bugs.mysql.com/9148>], [CVE-2005-0799](#) [<http://cve.mitre.org/cvename.cgi?name=CVE-2005-0799>])

- A `CREATE TABLE db_name.tbl_name LIKE ...` statement would crash the server when no database was selected. ([Bug#11028](#) [<http://bugs.mysql.com/11028>])
- IP addresses not shown in `ndb_mgm SHOW` command on second `ndb_mgmd` (or on `ndb_mgmd` restart). ([Bug#11596](#) [<http://bugs.mysql.com/11596>])
- `SHOW FIELDS` truncated the `TYPE` column to 40 characters. ([Bug#7142](#) [<http://bugs.mysql.com/7142>], [Bug#12817](#) [<http://bugs.mysql.com/12817>]) (Note: This fix was reverted in MySQL 4.1.15 because it broke existing applications.)
- Prepared statement with subqueries returned corrupt data. ([Bug#11458](#) [<http://bugs.mysql.com/11458>])
- A `ROLLUP` query could return a wrong result set when its `GROUP BY` clause contained references to the same column. ([Bug#11543](#) [<http://bugs.mysql.com/11543>])
- Queries with subqueries in the `FROM` clause were not being added to the query cache. ([Bug#11522](#) [<http://bugs.mysql.com/11522>])
- Possible crash on Windows when performing `GROUP BY` on a calculated field. ([Bug#11414](#) [<http://bugs.mysql.com/11414>])
- The `mysql_config` script did not handle symbolic linking properly. ([Bug#10986](#) [<http://bugs.mysql.com/10986>])
- When used within a subquery, `SUBSTRING()` returned an empty string. ([Bug#10269](#) [<http://bugs.mysql.com/10269>])
- On Mac OS X, `libmysqlclient_r.a` now is built with `--fno-common` to make it possible to link a shared two-level namespace library against `libmysqlclient_r.a`. ([Bug#10638](#) [<http://bugs.mysql.com/10638>])
- The handling by the `HEX()` function of numbers larger than  $2^{64}$  was improved. ([Bug#9854](#) [<http://bugs.mysql.com/9854>])
- Added a missing mutex when rotating the relay logs. Also, the server now logs an error message if the size of a relay log cannot be read. ([Bug#6987](#) [<http://bugs.mysql.com/6987>])
- Added a missing mutex when rotating the relay logs. Also, the server now logs an error message if the size of a relay log cannot be read. ([Bug#6987](#) [<http://bugs.mysql.com/6987>])
- `mysqldump` could crash for illegal or non-existent table names. ([Bug#9358](#) [<http://bugs.mysql.com/9358>])



- The `--no-data` option for `mysqldump` was being ignored if table names were given after the database name. (Bug#9558 [<http://bugs.mysql.com/9558>])
- `mysqldump` now exports `HASH` index definitions using `USING` rather than `TYPE` when the index name is optional. This corrects a problem when reloading the output for `PRIMARY KEY` definition, because `TYPE` must be preceded an index name, which is not given for a `PRIMARY KEY`. (Bug#11635 [<http://bugs.mysql.com/11635>])
- The `--master-data` option for `mysqldump` resulted in no error if the binary log was not enabled. Now an error occurs unless the `--force` option is given. (Bug#11678 [<http://bugs.mysql.com/11678>])
- Corrected an optimization failure where a query returned an incorrect result for use of a newly populated table until the table was flushed. (Bug#11700 [<http://bugs.mysql.com/11700>])
- Modifying a `CHAR` column with the `utf8` character set to a shorter length did not properly truncate values due to not computing their length in `utf8` character units. (Bug#11591 [<http://bugs.mysql.com/11591>])
- `DES_ENCRYPT()` and `DES_DECRYPT()` require SSL support to be enabled, but were not checking for it. Checking for incorrect arguments or resource exhaustion was also improved for these functions. (Bug#10589 [<http://bugs.mysql.com/10589>])
- Invoking the `DES_ENCRYPT()` function could cause a server crash if the server was started without the `--des-key-file` option. (Bug#11643 [<http://bugs.mysql.com/11643>])
- Selecting the result of an aggregate function for an `ENUM` or `SET` column within a subquery could result in a server crash. (Bug#11821 [<http://bugs.mysql.com/11821>])
- When used in joins, `SUBSTRING()` failed to truncate to zero those string values that could not be converted to numbers. (Bug#10124 [<http://bugs.mysql.com/10124>])
- Error when performing `GROUP BY` on calculated values of a single row table. (Bug#11414 [<http://bugs.mysql.com/11414>])
- `mysqldump --xml` did not format `NULL` column values correctly. (Bug#9657 [<http://bugs.mysql.com/9657>])
- `SHOW WARNINGS` did not properly display warnings generated by executing a cached query. (Bug#9414 [<http://bugs.mysql.com/9414>])
- Temporary tables were created in the data directory instead of `tmpdir`. (Bug#11440 [<http://bugs.mysql.com/11440>])
- `LOAD DATA ... REPLACE INTO ...` on a replication slave failed for an `InnoDB` table having a unique index in addition to the primary key. (Bug#11401 [<http://bugs.mysql.com/11401>])
- MySQL would not compile correctly on QNX due to missing `rint()` function. (Bug#11544 [<http://bugs.mysql.com/11544>])
- Incorrect results when searching using `IN()` where search items included `NULL` and `0`. (Bug#9393 [<http://bugs.mysql.com/9393>])
- NDB Cluster: When trying to open a table that could not be discovered or unpacked, cluster would return error codes which the MySQL server falsely interpreted as operating system errors. (Bug#103651 [<http://bugs.mysql.com/103651>])
- Manually inserting a row with `host=''` into `mysql.tables_priv` and performing a `FLUSH PRIVILEGES` would cause the server to crash. (Bug#11330 [<http://bugs.mysql.com/11330>])

- MySQL sometimes reported erroneously that certain character values had crashed a table when trying to convert other character sets to UTF-8. (Bug#9557 [http://bugs.mysql.com/9557])
- Using `CONCAT_WS()` on a column set `NOT NULL` caused incorrect results when used in a `LEFT JOIN`. (Bug#11469 [http://bugs.mysql.com/11469])
- `mysqld_safe` would sometimes fail to remove the pid file for the old `mysql` process after a crash. As a result, the server would fail to start due to a false `A mysqld process already exists...` error. (Bug#11122 [http://bugs.mysql.com/11122])
- For `MEMORY` tables, it was possible for updates to be performed using outdated key statistics when the updates involved only very small changes in a very few rows. This resulted in the random failures of queries such as `UPDATE t SET col = col + 1 WHERE col_key = 2;` where the same query with no `WHERE` clause would succeed. (Bug#10178 [http://bugs.mysql.com/10178])
- The `NULLIF()` function could produce incorrect results if the first argument was `NULL`. (Bug#11142 [http://bugs.mysql.com/11142])
- Optimizer performed range check when comparing unsigned integers to negative constants, could cause errors. (Bug#11185 [http://bugs.mysql.com/11185])
- Cluster failed to build with `gcc` 4.0. (Bug#11377 [http://bugs.mysql.com/11377])
- The `LAST_DAY()` failed to return `NULL` when supplied with an invalid argument. See Section 12.5, “Date and Time Functions”. (Bug#10568 [http://bugs.mysql.com/10568])
- Setting `@@SQL_MODE = NULL` caused an erroneous error message. (Bug#10732 [http://bugs.mysql.com/10732])
- Server crashed when using `GROUP BY` on the result of a `DIV` operation on a `DATETIME` value. (Bug#11385 [http://bugs.mysql.com/11385])
- Possible `NULL` values in `BLOB` columns could crash server when `BLOB` used in `GROUP BY`. (Bug#11295 [http://bugs.mysql.com/11295])
- Fixed 64 bit compiler warning for packet length in replication. (Bug#11064 [http://bugs.mysql.com/11064])
- `CASE` function returns incorrect result when its arguments are not constants and its return value is put into a regular or temporary table (temporary == created by SQL engine for `UNION`/non-indexed `GROUP BY` and such operations). (Bug#10151 [http://bugs.mysql.com/10151])
- A problem with the `my_global.h` file caused compilation of MySQL to fail on single-processor Linux systems running 2.6 kernels. (Bug#10364 [http://bugs.mysql.com/10364])
- Queries against a table using a compound index based on the length of a UTF-8 text column produced incorrect results. For example, given a table with an index defined as shown:

```
CREATE TABLE t (
  id INT NOT NULL,
  city VARCHAR(20) NOT NULL,
  KEY (city(7),id)
) TYPE=MYISAM CHARACTER SET=utf8;
```

Assuming that suitable data has been inserted into the table, then a query such as `SELECT * FROM t WHERE city = 'Durban';` would fail. (Bug#10253 [http://bugs.mysql.com/10253])

- The `mysqlhotcopy` script was not parsing the output of `SHOW SLAVE STATUS` correctly when called with the `--record_log_pos` option. (Bug#7967 [http://bugs.mysql.com/7967])

- An `UPDATE` query containing a subquery caused replication to fail. (Bug#9361 [<http://bugs.mysql.com/9361>])
- Last insert expected from a query of the form `INSERT ... SELECT ... ON DUPLICATE KEY UPDATE` would fail. (Bug#9728 [<http://bugs.mysql.com/9728>])
- `INSERT ... SELECT ... ON DUPLICATE KEY UPDATE` produced inaccurate results. (Bug#10886 [<http://bugs.mysql.com/10886>])
- `SELECT DISTINCT ... GROUP BY constant` returned multiple rows (it should return a single row). (Bug#8614 [<http://bugs.mysql.com/8614>])
- Queries of the form `UPDATE ... (SELECT ... ) SET ...` run on a replication master would crash all the slaves. (CVE-2004-4380 [<http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2004-4380>], Bug#10442 [<http://bugs.mysql.com/10442>])
- `OPTIMIZE` of InnoDB table did not return 'Table is full' if out of tablespace. (Bug#8135 [<http://bugs.mysql.com/8135>])
- Queries with `ROLLUP` returned wrong results for expressions containing `GROUP BY` columns. (Bug#7894 [<http://bugs.mysql.com/7894>])
- Fixed hang/crash with Boolean full-text search where a query contained more query terms than one-third of the query length (it could be achieved with truncation operator: 'a\*b\*c\*d\*'). (Bug#7858 [<http://bugs.mysql.com/7858>])
- The `mysql` client would output a prompt twice following input of very long strings, because it incorrectly assumed that a call to the `_cgets()` function would clear the input buffer. (Bug#10840 [<http://bugs.mysql.com/10840>])
- A bug in `FIELD()` function caused the value list to contain `NULL`. (Bug#10944 [<http://bugs.mysql.com/10944>])
- A three byte buffer overflow in the client functions could cause improper exiting of the client when reading a command from the user. (Bug#10841 [<http://bugs.mysql.com/10841>])
- Fixed a problem with the `cp1250_czech_cs` collation that caused some `LIKE` comparisons to fail. (Bug#9759 [<http://bugs.mysql.com/9759>])
- Fixed a problem resolving table names with `lower_case_table_names=2` when the table name lettercase differed in the `FROM` and `WHERE` clauses. (Bug#9500 [<http://bugs.mysql.com/9500>])
- Fixed server crash due to some internal functions not taking into account that for multi-byte character sets, `CHAR` and `VARCHAR` columns could exceed 255 bytes. (Bug#11167 [<http://bugs.mysql.com/11167>])
- Fixed a portability problem testing for `crypt()` support that caused compilation problems when using OpenSSL/yaSSL on HP-UX and Mac OS X. (Bug#10675 [<http://bugs.mysql.com/10675>], Bug#11150 [<http://bugs.mysql.com/11150>])
- The hostname cache was not working. (Bug#10931 [<http://bugs.mysql.com/10931>])
- For a `MERGE` table with `MyISAM` tables in other, symlinked, databases, `SHOW CREATE TABLE` reported the `MyISAM` tables using the name of the symlinked directory rather than the database name. (Bug#8183 [<http://bugs.mysql.com/8183>])
- Fixed a server crash resulting from an attempt to allocate too much memory when `GROUP BY blob_col` and `COUNT(DISTINCT)` were used. (Bug#11088 [<http://bugs.mysql.com/11088>])

- The incorrect sequence of statements `HANDLER tbl_name READ index_name NEXT` without a preceding `HANDLER tbl_name READ index_name = (value_list)` for an InnoDB table resulted in a server crash rather than an error. (Bug#5373 [<http://bugs.mysql.com/5373>])
- `SHOW WARNINGS` with a `LIMIT 0` clause returned all messages rather than an empty result set. (Bug#11095 [<http://bugs.mysql.com/11095>])
- On Windows, with `lower_case_table_names` set to 2, using `ALTER TABLE` to alter a MEMORY or InnoDB table that had a mixed-case name also improperly changed the name to lower-case. (Bug#9660 [<http://bugs.mysql.com/9660>])
- The server timed out SSL connections too quickly on Windows. (Bug#8572 [<http://bugs.mysql.com/8572>])
- Inserting a `DOUBLE` value into a `utf8` string column crashed the server on Windows. (Bug#10714 [<http://bugs.mysql.com/10714>])
- Executing `LOAD INDEX INTO CACHE` for a table while other threads were selecting from the table caused a deadlock. (Bug#10602 [<http://bugs.mysql.com/10602>])
- Fixed a server crash resulting from `CREATE TABLE ... SELECT` that selected from a table being altered by `ALTER TABLE`. (Bug#10224 [<http://bugs.mysql.com/10224>])
- Fixed a server crash resulting from invalid string pointer when inserting into the `mysql.host` table. (Bug#10181 [<http://bugs.mysql.com/10181>])
- `GROUP_CONCAT()` with `DISTINCT` and `WITH ROLLUP` ignored `DISTINCT` for some rows. (Bug#7405 [<http://bugs.mysql.com/7405>])
- Fixed a problem creating the result set for a `UNION` that involved long string values. Values were not being converted correctly to `TEXT` values. (Bug#10025 [<http://bugs.mysql.com/10025>])
- Locking for `CREATE TABLE ... SELECT` for InnoDB tables was too weak. It allowed `INSERT` statements issued for the created table while the `CREATE TABLE` statement was still running to appear in the binary log before the `CREATE TABLE` statement. (Bug#6678 [<http://bugs.mysql.com/6678>])
- InnoDB: In `DROP DATABASE`, check for all referencing tables from other databases before dropping any tables. (Bug#10335 [<http://bugs.mysql.com/10335>])
- InnoDB: Fix bug: InnoDB wrongly complained in the `.err` log that MySQL is trying to drop a non-existent table, if tablespace ran out. (Bug#10607 [<http://bugs.mysql.com/10607>])
- Fixed an overly strict debugging assertion that caused debug server builds to fail for some `col_name = const_expr`, where `const_expr` was a constant expression such as a subquery. (Bug#10020 [<http://bugs.mysql.com/10020>])
- `SUBSTR()` did not work properly for input in the `ucs2` character set. (Bug#10344 [<http://bugs.mysql.com/10344>])
- Fixed a problem causing an incorrect result for columns that include an aggregate function as part of an expression when `WITH ROLLUP` is added to `GROUP BY`. (Bug#7914 [<http://bugs.mysql.com/7914>])
- Fixed a server crash for `INSERT ... ON DUPLICATE KEY UPDATE` with `MERGE` tables, which do not have unique indexes. (Bug#10400 [<http://bugs.mysql.com/10400>])
- `CREATE TABLE t AS SELECT UUID()` created a `VARCHAR(12)` column, which is too small

to hold the 36-character result from `UUID()`. (Bug#9535 [http://bugs.mysql.com/9535])

- Portability fix for Cygwin: Don't use `#pragma interface` or `#pragma implementation` in source files. (Bug#10241 [http://bugs.mysql.com/10241])
- Fixed a `mysqldump` crash that occurred with the `--complete-insert` option when dumping tables with a large number of long column names. (Bug#10286 [http://bugs.mysql.com/10286])
- `ALTER TABLE ... ENABLE INDEXES` treated `NULL` values as equal when collecting index statistics for `MyISAM` tables, resulting in different statistics from those generated by `ANALYZE TABLE` and causing the optimizer to make poor index choices later. The same problem occurred for bulk insert statistics collection. Now `NULL` values are treated as unequal, just as for `ANALYZE TABLE`. (Bug#9622 [http://bugs.mysql.com/9622])

## D.1.11. Changes in release 4.1.12 (13 May 2005)

**Note:** The fix for interpretation of `MERGE` table `.MRG` files (Bug#10687 [http://bugs.mysql.com/10687]) was made for Windows builds after MySQL 4.1.12 was released and is present in MySQL 4.1.12a.

Functionality added or changed:

- **Incompatible change:** The behavior of `LOAD DATA INFILE` and `SELECT ... INTO OUTFILE` has changed when the `FIELDS TERMINATED BY` and `FIELDS ENCLOSED BY` values both are empty. Formerly, a column was read or written the display width of the column. For example, `INT(4)` was read or written using a field with a width of 4. Now columns are read and written using a field width wide enough to hold all values in the field. However, data files written before this change was made might not be reloaded correctly with `LOAD DATA INFILE` for MySQL 4.1.12 and up. This change also affects data files read by `mysqlimport` and written by `mysqldump --tab`, which use `LOAD DATA INFILE` and `SELECT ... INTO OUTFILE`. For more information, see Section 13.2.5, “`LOAD DATA INFILE` Syntax”. (Bug#12564 [http://bugs.mysql.com/12564])
- Added `--debug` option to `my_print_defaults`.
- When the server cannot read a table because it cannot read the `.frm` file, print a message that the table was created with a different version of MySQL. (This can happen if you create tables that use new features and then downgrade to an older version of MySQL.) (Bug#10435 [http://bugs.mysql.com/10435])
- New `/*>` prompt for `mysql`. This prompt indicates that a `/* ... */` comment was begun on an earlier line and the closing `*/` sequence has not yet been seen. (Bug#9186 [http://bugs.mysql.com/9186])
- Added `cp1250_croatian_ci` collation. (Bug#6505 [http://bugs.mysql.com/6505])
- Updated version of `libedit` to 2.9. (Bug#2596 [http://bugs.mysql.com/2596])
- **InnoDB:** When the maximum length of `SHOW INNODB STATUS` output would be exceeded, truncate the beginning of the list of active transactions, instead of truncating the end of the output. (Bug#5436 [http://bugs.mysql.com/5436])
- **InnoDB:** When `FOREIGN_KEY_CHECKS=0`, `ALTER TABLE` and `RENAME TABLE` will ignore any type incompatibilities between referencing and referenced columns. Thus, it will be possible to convert the character sets of columns that participate in a foreign key. Be sure to convert all tables before modifying any data! (Bug#9802 [http://bugs.mysql.com/9802])

- **InnoDB:** Setting the initial `AUTO_INCREMENT` value for an InnoDB table using `CREATE TABLE ... AUTO_INCREMENT = n` now works, and `ALTER TABLE ... AUTO_INCREMENT = n` resets the current value. (Bug#7061 [http://bugs.mysql.com/7061])
- **InnoDB:** If `innodb_locks_unsafe_for_binlog` option set and isolation level of the transaction is not set to serializable then InnoDB uses a consistent read for select in clauses like `INSERT INTO ... SELECT` and `UPDATE ... (SELECT)` that do not specify `FOR UPDATE` or `IN SHARE MODE`. Thus no locks are set to rows read from selected table.
- Previously in MySQL 4.1, an `Illegal mix of collations` error occurred when mixing strings from same character set when one had a non-binary collation and the other a binary collation. Now the binary collation takes precedence, so that both strings are treated as having the binary collation. This restores compatibility with MySQL 4.0 behavior.
- Added the `cp932` Japanese character set.

Bugs fixed:

- **Security fix:** If `mysqld` was started with `--user=non_existent_user`, it would run using the privileges of the account it was invoked from, even if that was `root`. (Bug#9833 [http://bugs.mysql.com/9833])
- Fixed handling of floats and doubles when using prepared statement API in the embedded server. (Bug#10443 [http://bugs.mysql.com/10443])
- An error in the implementation of the `MyISAM` compression algorithm caused `myisampack` to fail with very large sets of data (total size of all the records in a single column needed to be  $\geq 3$  GB in order to trigger this issue). (Bug#8321 [http://bugs.mysql.com/8321])
- When `SELECT constant` was the final `SELECT` in a `UNION`, a trailing `LIMIT ...` worked, but a trailing `ORDER BY ...` or `ORDER BY ... LIMIT ...` did not. (Bug#10032 [http://bugs.mysql.com/10032])
- `MERGE` tables could fail on Windows due to incorrect interpretation of pathname separator characters for filenames in the `.MRG` file. (Bug#10687 [http://bugs.mysql.com/10687])
- `CHAR` and `VARCHAR` columns that used the `sjis` character set were not being saved correctly, causing the following columns to be corrupted. (Bug#10493 [http://bugs.mysql.com/10493])
- For a user-defined function invoked from within a prepared statement, the UDF's initialization routine was invoked for each execution of the statement, but the deinitialization routine was not. (It was invoked only when the statement was closed.) For UDFs that have an expensive deinit function (such as `myperl`), this bugfix will have negative performance consequences. (Bug#9913 [http://bugs.mysql.com/9913])
- Fix `CREATE TABLE ... LIKE` to work when `lower_case_table_names` is set on a case-sensitive filesystem and the source table name is not given in lowercase. (Bug#9761 [http://bugs.mysql.com/9761])
- `my_print_defaults` was ignoring the `--defaults-extra-file` option or crashing when the option was given. (Bug#9136 [http://bugs.mysql.com/9136], Bug#9851 [http://bugs.mysql.com/9851])
- For `MERGE` tables, avoid writing absolute pathnames in the `.MRG` file for the names of the constituent `MyISAM` tables so that if the data directory is moved, `MERGE` tables will not break. For `mysqld`, write just the `MyISAM` table name if it is in the same database as the `MERGE` table, and a path relative to the data directory otherwise. For the embedded servers, absolute pathnames may still be used.

- ([Bug#5964](http://bugs.mysql.com/5964) [<http://bugs.mysql.com/5964>])
- Corrected the error message for exceeding the `MAX_CONNECTIONS_PER_HOUR` limit to say `max_connections_per_hour` instead of `max_connections`. ([Bug#9947](http://bugs.mysql.com/9947) [<http://bugs.mysql.com/9947>])
  - Fixed incorrect memory block allocation for the query cache in the embedded server. ([Bug#9549](http://bugs.mysql.com/9549) [<http://bugs.mysql.com/9549>])
  - Fixed a `configure` problem in checking for capability of performing atomic operations. ([Bug#7970](http://bugs.mysql.com/7970) [<http://bugs.mysql.com/7970>])
  - `net_read_timeout` and `net_write_timeout` were not being respected on Windows. ([Bug#9721](http://bugs.mysql.com/9721) [<http://bugs.mysql.com/9721>])
  - An error occurred if you specified a default value of `TRUE` or `FALSE` for a `BOOL` column. ([Bug#9666](http://bugs.mysql.com/9666) [<http://bugs.mysql.com/9666>])
  - Corrected some failures of prepared statements for SQL (`PREPARE` plus `EXECUTE`) to return all rows for some `SELECT` statements. ([Bug#9096](http://bugs.mysql.com/9096) [<http://bugs.mysql.com/9096>], [Bug#9777](http://bugs.mysql.com/9777) [<http://bugs.mysql.com/9777>])
  - Remove extra slashes in `--tmpdir` value (for example, convert `/var//tmp` to `/var/tmp`, because they caused various errors. ([Bug#8497](http://bugs.mysql.com/8497) [<http://bugs.mysql.com/8497>])
  - Fixed a sort order problem with the `latin2_croatian_ci` collation. All tables that have indexes that use this collation will be treated as crashed. After upgrading, for each such table, you must use `CHECK TABLE` and possibly repair the table. ([Bug#6505](http://bugs.mysql.com/6505) [<http://bugs.mysql.com/6505>])
  - `mysqld` was not checking whether the PID file was successfully created. ([Bug#5843](http://bugs.mysql.com/5843) [<http://bugs.mysql.com/5843>])
  - With `DISTINCT CONCAT(col, ...)` returned incorrect results when the arguments to `CONCAT()` were columns with an integer data type declared with a display width narrower than the values in the column. (For example, if an `INT(1)` column contain `1111`.) ([Bug#4082](http://bugs.mysql.com/4082) [<http://bugs.mysql.com/4082>])
  - Fixed `configure` to properly recognize whether NTPL is available on Linux. ([Bug#2173](http://bugs.mysql.com/2173) [<http://bugs.mysql.com/2173>])
  - Fixed a portability problem in compiling `mysql.cc` with `VC++` on Windows. ([Bug#10245](http://bugs.mysql.com/10245) [<http://bugs.mysql.com/10245>])
  - `CAST(string_argument AS UNSIGNED)` didn't work for big integers above the signed range. Now this function and `CAST(string_argument AS SIGNED)` also produces warnings for wrong string arguments. ([Bug#7036](http://bugs.mysql.com/7036) [<http://bugs.mysql.com/7036>])
  - Fixed compile problem with MinGW. Thanks to Nils Durner for patch! ([Bug#8872](http://bugs.mysql.com/8872) [<http://bugs.mysql.com/8872>])
  - MySQL no longer automatically blocks IP numbers for which `gethostbyname_r()` fails when the reason is that the DNS server is down. Thanks to Jeremy Cole for patch. ([Bug#8467](http://bugs.mysql.com/8467) [<http://bugs.mysql.com/8467>])
  - Fixed a bug in the key cache that caused a core dump. ([Bug#10167](http://bugs.mysql.com/10167) [<http://bugs.mysql.com/10167>])
  - The `--delimiter` option for the `nds_select` program was non-functional. ([Bug#10287](http://bugs.mysql.com/10287) [<http://bugs.mysql.com/10287>])

- `MAX()` for an `INT UNSIGNED` (unsigned 4-byte integer) column could return negative values if the column contained values larger than  $2^{31}$ . (Bug#9298 [http://bugs.mysql.com/9298])
- Fixed a deadlock resulting from use of `FLUSH TABLES WITH READ LOCK` while an `INSERT DELAYED` statement is in progress. (Bug#7823 [http://bugs.mysql.com/7823])
- Multiple-table updates could produce spurious data-truncation warnings if they used a join across columns that are indexed using a column prefix. (Bug#9103 [http://bugs.mysql.com/9103])
- Use of a subquery that used `WITH ROLLUP` in the `FROM` clause of the main query sometimes resulted in a `Column cannot be null` error. (Bug#9681 [http://bugs.mysql.com/9681])
- `RENAME TABLE` for an `ARCHIVE` table failed if the `.arn` file was not present. (Bug#9911 [http://bugs.mysql.com/9911])
- Fixed an optimizer problem where extraneous comparisons between `NULL` values in indexed columns were being done for operators such as `=` that are never true for `NULL`. (Bug#8877 [http://bugs.mysql.com/8877])
- `SELECT ROUND(expr)` produced a different result from `CREATE TABLE ... SELECT ROUND(expr)`. (Bug#9837 [http://bugs.mysql.com/9837])
- Fixed some `awk` script portability problems in `cmd-line-utils/libedit/makelist.sh`. (Bug#9954 [http://bugs.mysql.com/9954])
- Changed metadata for result of `SHOW KEYS`: Data type for `Sub_part` column now is `SMALLINT` rather than `TINYINT` because key part length can be longer than 255. (Bug#9439 [http://bugs.mysql.com/9439])
- Fixed some problems with `myisampack` on 64-bit systems that resulted in segmentation violations. (Bug#9487 [http://bugs.mysql.com/9487])
- Fixed an optimizer bug in computing the union of two ranges for the `OR` operator. (Bug#9348 [http://bugs.mysql.com/9348])
- Fixed an index corruption problem for `MyISAM` tables that resulted from the 4.1 behavior of padding values with blanks for comparison: Dumping a table with `mysqldump`, reloading it, and then re-running the binary log against it crashed the index and necessitated a repair. (Bug#9188 [http://bugs.mysql.com/9188])
- Fixed a segmentation fault in `mysqlcheck` that occurred when the last table checked in `-auto-repair` mode returned an error (such as the table being a `MERGE` table). (Bug#9492 [http://bugs.mysql.com/9492])
- Fixed the client/server protocol for prepared statements so that reconnection works properly when the connection is killed while `reconnect` is enabled. (Bug#8866 [http://bugs.mysql.com/8866])
- `INSERT ... ON DUPLICATE KEY UPDATE` incorrectly updated a `TIMESTAMP` column to the current timestamp, even if the update list included `col_name = col_name` for that column to prevent the update. (Bug#7806 [http://bugs.mysql.com/7806])
- Starting `mysqld` with the `--skip-innodb` and `--default-storage-engine=innodb` (or `--default-table-type=innodb`) caused a server crash. (Bug#9815 [http://bugs.mysql.com/9815])
- Queries containing `CURRENT_USER()` incorrectly were registered in the query cache. (Bug#9796 [http://bugs.mysql.com/9796])
- A server installed as a Windows service and started with `--shared-memory` could not be



stopped. ([Bug#9665](http://bugs.mysql.com/9665) [<http://bugs.mysql.com/9665>])

- `mysqldump` dumped core when invoked with `--tmp` and `--single-transaction` options and a non-existent table name. ([Bug#9175](http://bugs.mysql.com/9175) [<http://bugs.mysql.com/9175>])
- Additional fix for `mysql_server_init()` and `mysql_server_end()` C API functions so that stopping and restarting the embedded server will not cause a crash. ([Bug#7344](http://bugs.mysql.com/7344) [<http://bugs.mysql.com/7344>])
- `mysql.server` no longer uses non-portable `alias` command or LSB functions. ([Bug#9852](http://bugs.mysql.com/9852) [<http://bugs.mysql.com/9852>])
- Fixed a `readline`-related crash in `mysql` when the user pressed Control-R. ([Bug#9568](http://bugs.mysql.com/9568) [<http://bugs.mysql.com/9568>])
- `TIMEDIFF()` with a negative time first argument and positive time second argument produced incorrect results. ([Bug#8068](http://bugs.mysql.com/8068) [<http://bugs.mysql.com/8068>])
- Fixed a bug that caused concurrent inserts to be allowed into the tables in the `SELECT ... UNION ...` part of `INSERT ... SELECT ... UNION ...`. This could result in the incorrect order of queries in the binary log. ([Bug#9922](http://bugs.mysql.com/9922) [<http://bugs.mysql.com/9922>])
- The warning message from `GROUP_CONCAT()` did not always indicate the correct number of lines. ([Bug#8681](http://bugs.mysql.com/8681) [<http://bugs.mysql.com/8681>])
- InnoDB: `ENUM` and `SET` columns were treated incorrectly as character strings. This bug did not manifest itself with `latin1` collations, but it caused malfunction with `utf8`. Old tables will continue to work. In new tables, `ENUM` and `SET` will be internally stored as unsigned integers. ([Bug#9526](http://bugs.mysql.com/9526) [<http://bugs.mysql.com/9526>])
- InnoDB: Avoid test suite failures caused by a locking conflict between two server instances at server shutdown/startup. This conflict on advisory locks appears to be the result of a bug in the operating system; these locks should be released when the files are closed, but somehow that does not always happen immediately in Linux. ([Bug#9381](http://bugs.mysql.com/9381) [<http://bugs.mysql.com/9381>])
- InnoDB: Prevent `ALTER TABLE` from changing the storage engine if there are foreign key constraints on the table. ([Bug#5574](http://bugs.mysql.com/5574) [<http://bugs.mysql.com/5574>], [Bug#5670](http://bugs.mysql.com/5670) [<http://bugs.mysql.com/5670>])
- InnoDB: Fixed a deadlock without any locking, simple select and update. ([Bug#7975](http://bugs.mysql.com/7975) [<http://bugs.mysql.com/7975>]) InnoDB now takes an exclusive lock when `INSERT ON DUPLICATE KEY UPDATE` is checking duplicate keys.
- InnoDB: Fix a problem in crash recovery of `.ibd` files on Windows if the user used `lower_case_table_names=0` or `2`; the directory scan in crash recovery forgot to put all paths to lower case, so that the tablespace name would be consistent with the internal data dictionary of InnoDB.
- InnoDB: Add fault tolerance in the scan of `.ibd` files at a crash recovery; formerly a single failure of `readdir_get_next` caused the rest of the directory to be skipped.
- InnoDB: Fix assertion failures of type `ut_a(cursor->old_stored == BTR_PCUR_OLD_STORED)` and `prebuilt->template_type == 0`. This bug was introduced in 4.1.10 and 4.0.24. ([Bug#9670](http://bugs.mysql.com/9670) [<http://bugs.mysql.com/9670>])
- InnoDB: Fix a performance bug: At the shutdown, write the latest lsn only to the first pages of the `ibdata` files of the system tablespace, NOT to the `.ibd` files; writing to tens of thousands `.ibd` files can take minutes.

- Fix for auto-increment not working with `INSERT . . SELECT` and NDB storage engine. (Bug#9675 [<http://bugs.mysql.com/9675>])

## D.1.12. Changes in release 4.1.11 (01 April 2005)

Functionality added or changed:

- `ONLY_FULL_GROUP_BY` no longer is included in the `ANSI` composite SQL mode. (Bug#8510 [<http://bugs.mysql.com/8510>])
- `mysqld_safe` will create the directory where the UNIX socket file is to be located if the directory does not exist. This applies only to the last component of the directory pathname. (Bug#8513 [<http://bugs.mysql.com/8513>])
- The coercibility for the return value of functions such as `USER()` or `VERSION()` now is “system constant” rather than “implicit.” This makes these functions more coercible than column values so that comparisons of the two do not result in `Illegal mix of collations` errors. `COERCIBILITY()` was modified to accommodate this new coercibility value. See Section 12.9.3, “Information Functions”.
- User variable coercibility has been changed from “coercible” to “implicit.” That is, user variables have the same coercibility as column values.
- `NULL` now is considered more coercible than string constants. This resolves some `Illegal mix of collations` conflicts.
- Modified the parser to allow `SELECT` statements following the `UNION` keyword to be subqueries in parentheses. (Bug#2435 [<http://bugs.mysql.com/2435>])
- For slave replication servers started with `--replicate-*` options, statements that should not be replicated according those options no longer are written to the slave's general query log. (Bug#8297 [<http://bugs.mysql.com/8297>])
- Added `SQL_NOTES` session variable to cause `Note`-level warnings not to be recorded. (Bug#6662 [<http://bugs.mysql.com/6662>])
- `InnoDB`: Commit after every 10,000 copied rows when executing `CREATE INDEX`, `DROP INDEX` or `OPTIMIZE TABLE`, which are internally implemented as `ALTER TABLE`. This makes it much faster to recover from an aborted operation.
- Added a new global system variable `slave_transaction_retries`: If the replication slave SQL thread fails to execute a transaction because of an `InnoDB` deadlock or exceeded `InnoDB`'s `innodb_lock_wait_timeout` or `NDBCluster`'s `TransactionDeadlockDetectionTimeout` or `TransactionInactiveTimeout`, it automatically retries `slave_transaction_retries` times before stopping with an error. The default in MySQL 4.1 is 0. You must explicitly set the value greater than 0 to enable the “retry” behavior. (In MySQL 5.0.3 or newer, the default is 10.) (Bug#8325 [<http://bugs.mysql.com/8325>])
- Added `--with-big-tables` compilation option to `configure`. (Previously it was necessary to pass `-DBIG_TABLES` to the compiler manually in order to enable large table support.) See Section 2.9.2, “Typical configure Options”, for details.
- Added configuration directives `!include` and `!includedir` for including option files and searching directories for option files. See Section 4.3.2, “Using Option Files”, for usage.

Bugs fixed:

- Disallow use of `SESSION` or `GLOBAL` for user variables. ([Bug#9286](http://bugs.mysql.com/9286) [<http://bugs.mysql.com/9286>])
- The use of `XOR` together with `NOT ISNULL( )` erroneously resulted in some outer joins being converted to inner joins by the optimizer. ([Bug#9017](http://bugs.mysql.com/9017) [<http://bugs.mysql.com/9017>])
- Fixed `utf8_spanish2_ci` and `ucs2_spanish2_ci` collations to not consider ‘r’ equal to ‘rr’. If you upgrade to this version from an earlier version, you should rebuild the indexes of affected tables. ([Bug#9269](http://bugs.mysql.com/9269) [<http://bugs.mysql.com/9269>])
- Allow extra HKSCS and cp950 characters (`big5` extension characters) to be accepted in `big5` columns. ([Bug#9357](http://bugs.mysql.com/9357) [<http://bugs.mysql.com/9357>])
- `BLOB(M)` and `TEXT(M)` columns, with `M` less than 256, were being created as `BLOB` and `TEXT` columns rather than `TINYBLOB` or `TINYTEXT` columns. ([Bug#9303](http://bugs.mysql.com/9303) [<http://bugs.mysql.com/9303>])
- Fixed a problem with `INSERT ... SELECT ... ON DUPLICATE KEY UPDATE` where a column named in the insert list and in the `ON DUPLICATE KEY UPDATE` clause was erroneously declared to be ambiguous. ([Bug#8147](http://bugs.mysql.com/8147) [<http://bugs.mysql.com/8147>])
- In prepared statements, subqueries containing parameters were erroneously treated as `const` tables during preparation, resulting in a server crash. ([Bug#8807](http://bugs.mysql.com/8807) [<http://bugs.mysql.com/8807>])
- Fixed a problem with `OPTIMIZE TABLE` for `InnoDB` tables being written twice to the binary log. ([Bug#9149](http://bugs.mysql.com/9149) [<http://bugs.mysql.com/9149>])
- Provide more informative error messages in clustered setting when a query is issued against a table that has been modified by another `mysqld` server. ([Bug#6762](http://bugs.mysql.com/6762) [<http://bugs.mysql.com/6762>])
- For `MyISAM` tables, `REPAIR TABLE` no longer discard rows that have incorrect checksum. ([Bug#9824](http://bugs.mysql.com/9824) [<http://bugs.mysql.com/9824>])
- Depending on index statistics, `GROUP BY col1, col2, ...` could return incorrect results if the first table processed for a join had several indexes that cover the grouped columns. ([Bug#9213](http://bugs.mysql.com/9213) [<http://bugs.mysql.com/9213>])
- Fixed incorrect evaluation of `ALL/ANY` subqueries that contain a `HAVING` clause. ([Bug#9350](http://bugs.mysql.com/9350) [<http://bugs.mysql.com/9350>])
- Fixed server crash when left expression of `IN/ALL/ANY` comparison was a subquery. ([Bug#8888](http://bugs.mysql.com/8888) [<http://bugs.mysql.com/8888>])
- Fixed option-parsing code for the embedded server to understand `K`, `M`, and `G` suffixes for the `net_buffer_length` and `max_allowed_packet` options. ([Bug#9472](http://bugs.mysql.com/9472) [<http://bugs.mysql.com/9472>])
- Fixed a crash when using `TIMESTAMP` columns with no minute or second parts in `GROUP BY` with the `new` system variable set to 1. ([Bug#9401](http://bugs.mysql.com/9401) [<http://bugs.mysql.com/9401>])
- If a `MyISAM` table on Windows had `INDEX DIRECTORY` or `DATA DIRECTORY` table options, `mysqldump` dumped the directory pathnames with single-backslash pathname separators. This would cause syntax errors when importing the dump file. `mysqldump` now changes ‘\’ to ‘/’ in the pathnames on Windows. ([Bug#6660](http://bugs.mysql.com/6660) [<http://bugs.mysql.com/6660>])
- Fixed a server crash caused by use of `NOW( )` is a subquery. ([Bug#8824](http://bugs.mysql.com/8824) [<http://bugs.mysql.com/8824>])

- Fixed problems with static variables to allow building on Fedora Core 3. ([Bug#6554](#) [<http://bugs.mysql.com/6554>])
- Some user variables were not being handled with “implicit” coercibility. ([Bug#9425](#) [<http://bugs.mysql.com/9425>])
- Setting the `max_error_count` system variable to 0 resulted in a setting of 1. ([Bug#9072](#) [<http://bugs.mysql.com/9072>])
- Fixed a collation coercibility problem that caused a union between binary and non-binary columns to fail. ([Bug#6519](#) [<http://bugs.mysql.com/6519>])
- Fixed a problem with the `tee` command in `mysql` that resulted in `mysql` crashing. ([Bug#8499](#) [<http://bugs.mysql.com/8499>])
- On Windows, create shared memory objects with the proper access rights to make them usable when the client and server are running under different accounts. ([Bug#8226](#) [<http://bugs.mysql.com/8226>])
- Bundled `zlib` in the source distribution was upgraded to 1.2.2. ([Bug#9118](#) [<http://bugs.mysql.com/9118>])
- Fixed server crash resulting from queries that combined `SELECT DISTINCT`, `SUM( )`, and `ROLLUP`. ([Bug#8615](#) [<http://bugs.mysql.com/8615>])
- Incorrect results were returned from queries that combined `SELECT DISTINCT`, `GROUP BY` , and `ROLLUP`. ([Bug#8616](#) [<http://bugs.mysql.com/8616>])
- Fixed a bug that under certain circumstances could allow a privilege escalation via database wildcards in `GRANT`. ([CVE-2004-0957](#) [<http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2004-0957>])
- Too many rows were returned from queries that combined `ROLLUP` and `LIMIT` if `SQL_CALC_FOUND_ROWS` was given. ([Bug#8617](#) [<http://bugs.mysql.com/8617>])
- `mysqldump` misinterpreted ‘`_`’ and ‘`%`’ characters in the names of tables to be dumped as wildcard characters. ([Bug#9123](#) [<http://bugs.mysql.com/9123>])
- Made the `relay_log_space_limit` system variable show up in the output of `SHOW VARIABLES`. ([Bug#7100](#) [<http://bugs.mysql.com/7100>])
- Use of `GROUP_CONCAT(x)` in a subquery, where `x` was an alias to a column in the outer query, resulted in a server crash. ([Bug#8656](#) [<http://bugs.mysql.com/8656>])
- The `CHARSET( )`, `COLLATION( )`, and `COERCIBILITY( )` functions sometimes returned `NULL`. `CHARSET( )` and `COLLATION( )` returned `NULL` when given any of these arguments that evaluated to `NULL`: A system function such as `DATABASE( )`; a column value; and a user variable. Now `CHARSET( )` and `COLLATION( )` return the system character set and collation; the column character set and collation; and `binary`. `COERCIBILITY(NULL)` now returns “ignorable” coercibility rather than `NULL`. ([Bug#9129](#) [<http://bugs.mysql.com/9129>])
- Expressions involving nested `CONCAT( )` calls and character set conversion of string constants could return an incorrect result. ([Bug#8785](#) [<http://bugs.mysql.com/8785>])
- The `MEMORY` storage engine did not properly increment an `AUTO_INCREMENT` column if there was a second composite index that included the column. ([Bug#8489](#) [<http://bugs.mysql.com/8489>])
- Fixed a bug in the filesort routine such that killing a filesort could cause an assertion failure. ([Bug#8799](#) [<http://bugs.mysql.com/8799>])

- `REPAIR TABLE` did not invalidate query results in the query cache that were generated from the table. (Bug#8480 [http://bugs.mysql.com/8480])
- If `max_join_size` was set, a query containing a subquery that exceeded the examined-rows limit could hang. (Bug#8726 [http://bugs.mysql.com/8726])
- Mixed-case database and table names in the grant tables were ignored for authentication if the `lower_case_table_names` system variable was set. `GRANT` will not create such privileges when `lower_case_table_names` is set, but it is possible to create them by direct manipulation of the grant tables, or that old grant records were present before setting the variable. (Bug#7989 [http://bugs.mysql.com/7989])
- `AES_DECRYPT(col_name, key)` could fail to return `NULL` for invalid values in `col_name`, if `col_name` was declared as `NOT NULL`. (Bug#8669 [http://bugs.mysql.com/8669])
- Ordering by unsigned expression (more complex than a column reference) was treating the value as signed, producing incorrectly sorted results. (Bug#7425 [http://bugs.mysql.com/7425])
- `HAVING` was treating unsigned columns as signed. (Bug#7425 [http://bugs.mysql.com/7425])
- Fixed a problem with boolean full-text searches on `utf8` columns where a double quote in the search string caused a server crash. (Bug#8351 [http://bugs.mysql.com/8351])
- `MIN(col_name)` and `MAX(col_name)` could fail to produce the correct result if `col_name` was contained in multiple indexes and the optimizer did not choose the first index that contained the column. (Bug#8893 [http://bugs.mysql.com/8893])
- Table creation for a `MyISAM` table failed if `DATA DIRECTORY` or `INDEX DIRECTORY` options were given that specified the pathname to the database directory where the table files would be created by default. (Bug#8707 [http://bugs.mysql.com/8707])
- Fixed a problem with `LIKE` pattern-matching for strings with the `cp1251_bin` binary collation. (Bug#8560 [http://bugs.mysql.com/8560])
- A join on two tables failed when each contained a `BIGINT UNSIGNED` column that were compared when their values exceeded  $2^{63} - 1$ . The match failed and the join returned no rows. (Bug#8562 [http://bugs.mysql.com/8562])
- For a query with both `GROUP BY` and `COUNT(DISTINCT)` clauses and a `FROM` clause with a subquery, `NULL` was returned for any `VARCHAR` column selected by the subquery. (Bug#8218 [http://bugs.mysql.com/8218])
- Fixed an optimizer bug that caused incorrectly ordered result from a query that used a `FULLTEXT` index to retrieve rows and there was another index that was usable for `ORDER BY`. For such a query, `EXPLAIN` showed `fulltext` join type, but regular (not `FULLTEXT`) index in the `Key` column. (Bug#6635 [http://bugs.mysql.com/6635])
- For a statement string that contained multiple slow queries, only the last one would be written to the slow query log. (Bug#8475 [http://bugs.mysql.com/8475])
- When the server was started with `--skip-name-resolve`, specifying hostname values that included netmasks in `GRANT` statements did not work. (Bug#8471 [http://bugs.mysql.com/8471])
- The `--set-character-set` option for `myisamchk` was changed to `--set-collation`. The value needed for specifying how to sort indexes is a collation name, not a character set name. (Bug#8349 [http://bugs.mysql.com/8349])
- Hostname matching didn't work if a netmask was specified for table-specific privileges. (Bug#3309 [http://bugs.mysql.com/3309])

- Binary data stored in `BLOB` or `BINARY` columns would be erroneously dumped if `mysqldump` was invoked with `--hex-blob` and `--skip-extended-insert` arguments. This happened if data contained characters larger than `0x7F` ([Bug#8830](#) [<http://bugs.mysql.com/8830>]).
- Corruption of `MyISAM` table indexes could occur with `TRUNCATE TABLE` if the table had already been opened. For example, this was possible if the table had been opened implicitly by selecting from a `MERGE` table that mapped to the `MyISAM` table. The server now issues an error message for `TRUNCATE TABLE` under these conditions. ([Bug#8306](#) [<http://bugs.mysql.com/8306>])
- Fixed handling of table-name matching in `mysqlhotcopy` to accommodate `DBD: :mysql 2.9003` and up (which implement identifier quoting). ([Bug#8136](#) [<http://bugs.mysql.com/8136>])
- In the `mysql_real_escape_string()` C API function, when a multi-byte character is encountered that is illegal in the current character set, escape only the first byte, not each byte. This avoids creating a valid character from an invalid one. ([Bug#8378](#) [<http://bugs.mysql.com/8378>])
- Fixed a problem with the `cp1250_czech_cs` collation that caused empty literal strings not to compare equal to empty character columns. ([Bug#8840](#) [<http://bugs.mysql.com/8840>])
- Fixed a problem in index cost calculation that caused a `USE INDEX` or `FORCE INDEX` hint not to be used properly for a `LEFT JOIN` across indexed `BLOB` columns. ([Bug#7520](#) [<http://bugs.mysql.com/7520>])
- The data type for `MAX(datetime_col)` was returned as `VARCHAR` rather than `DATETIME` if the query included a `GROUP BY` clause. ([Bug#5615](#) [<http://bugs.mysql.com/5615>])
- `FOUND_ROWS()` returned an incorrect value for preceding `SELECT` statements that used no table or view. ([Bug#6089](#) [<http://bugs.mysql.com/6089>])
- In string literals with an escape character (`'\`'`) followed by a multi-byte character that has a second byte of `'\`'`, the literal was not interpreted correctly. The next character now is escaped, not just the next byte. ([Bug#8303](#) [<http://bugs.mysql.com/8303>])
- InnoDB: Work around a problem in AIX 5.1 patched with ML7 security patch: InnoDB would refuse to open its `ibdata` files, complaining about an operating system error 0.
- InnoDB: Fixed a memory corruption bug if one created a table with a primary key that contained at least two column prefixes. An example: `CREATE TABLE t(a char(100), b tinyblob, PRIMARY KEY(a(5), b(10)))`.
- InnoDB: Do not try to space-pad `BLOB` columns containing `ucs2` characters. This avoids an assertion failure that was introduced when fixing [Bug#7350](#) [<http://bugs.mysql.com/7350>]. ([Bug#8771](#) [<http://bugs.mysql.com/8771>])
- InnoDB: Fixed a bug: MySQL 4.1.8 to 4.1.10 could complain that an InnoDB table created with MySQL 3.23.49 or earlier was in the new compact InnoDB table format of 5.0.3 or later, and InnoDB would refuse to use that table. There is nothing wrong with the table, it is `mysqld` that is in error. Workaround: wait that 4.1.11 is released before doing an upgrade, or dump the table and re-create it with any MySQL version `>= 3.23.50` before upgrading.
- InnoDB: Honor the `--tmpdir` startup option when creating temporary files. Previously, InnoDB temporary files were always created in the temporary directory of the operating system. On Netware, InnoDB will continue to ignore `--tmpdir`. ([Bug#5822](#) [<http://bugs.mysql.com/5822>])
- InnoDB: If MySQL wrote to its binlog, but for some reason `trx->update_undo` and `trx->insert_undo` were `NULL` in InnoDB, then `trx->commit_lsn` was garbage, and InnoDB could assert in the log flush of `trx_commit_complete_for_mysql()`. ([Bug#9277](#) [<http://bugs.mysql.com/9277>])

- **InnoDB**: If InnoDB cannot allocate memory, keep retrying for 60 seconds before we intentionally crash `mysqld`; maybe the memory shortage is just temporary.
- **InnoDB**: If one used `LOCK TABLES`, created an InnoDB temp table, and did a multiple-table update where a `MyISAM` table was the update table and the temp table was a read table, then InnoDB asserted in `row0sel.c` because `n_mysql_tables_in_use` was 0. Also, we remove the assertion altogether and just print an error to the `.err` log if this important consistency check fails. ([Bug#8677](http://bugs.mysql.com/8677) [<http://bugs.mysql.com/8677>])
- `mysqldump` now avoids writing `SET NAMES` to the dump output if the server is older than version 4.1 and would not understand that statement. ([Bug#7997](http://bugs.mysql.com/7997) [<http://bugs.mysql.com/7997>])
- Fixed a bug in `my_print_defaults` that made it ignore the `--defaults-extra-file` and `--defaults-file` options.
- Retrieving from a view defined as a `SELECT` that mixed `UNION ALL` and `UNION DISTINCT` resulted in a different result than retrieving from the original `SELECT`. ([Bug#6565](http://bugs.mysql.com/6565) [<http://bugs.mysql.com/6565>])
- Worked around a bug in support for NSS support in `glibc` when static linking is used and LDAP is one of the NSS sources. The workaround is to detect when the bug causes a segmentation fault and issue a diagnostic message with information about the problem. ([Bug#3037](http://bugs.mysql.com/3037) [<http://bugs.mysql.com/3037>], [Bug#4872](http://bugs.mysql.com/4872) [<http://bugs.mysql.com/4872>])
- If the `mysql` prompt was configured to display the default database name, and that database was dropped, `mysql` did not update the prompt. ([Bug#4802](http://bugs.mysql.com/4802) [<http://bugs.mysql.com/4802>])
- `perror` was printing a spurious extra line of output ("Error code ###: Unknown error ###" printed directly before the correct line with the error message). ([Bug#8517](http://bugs.mysql.com/8517) [<http://bugs.mysql.com/8517>])
- The `CHAR ( )` function was not ignoring `NULL` arguments, contrary to the documentation. ([Bug#6317](http://bugs.mysql.com/6317) [<http://bugs.mysql.com/6317>])
- Neither `SHOW ERRORS` nor `SHOW WARNINGS` were displaying Error-level messages. ([Bug#6572](http://bugs.mysql.com/6572) [<http://bugs.mysql.com/6572>])
- Creating a table using a name containing a character that is illegal in `character_set_client` resulted in the character being stripped from the name and no error. The character now is considered an error. ([Bug#8041](http://bugs.mysql.com/8041) [<http://bugs.mysql.com/8041>])
- Fixed a problem with the Cyrillic letters I and SHORT I being treated the same by the `utf8_general_ci` collation. ([Bug#8385](http://bugs.mysql.com/8385) [<http://bugs.mysql.com/8385>])
- The `MAX_CONNECTIONS_PER_HOUR` resource limit was not being reset hourly and thus imposed an absolute limit on number of connections per account until the server is restarted or the limits flushed. ([Bug#8350](http://bugs.mysql.com/8350) [<http://bugs.mysql.com/8350>])
- With a database was dropped with `lower_case_table_names=2`, tables in the database also were dropped but not being flushed properly from the table cache. If the database was re-created, the tables also would appear to have been re-created. ([Bug#8355](http://bugs.mysql.com/8355) [<http://bugs.mysql.com/8355>])
- Changed `mysql_server_end ( )` C API function to restore more variables to their initial state so that a subsequent call to `mysql_server_init ( )` would not cause a client program crash. ([Bug#7344](http://bugs.mysql.com/7344) [<http://bugs.mysql.com/7344>])
- Fixed a problem with accented letters improperly being treated as distinct with the `utf_general_ci` collation. ([Bug#7878](http://bugs.mysql.com/7878) [<http://bugs.mysql.com/7878>])
- `ENUM` and `SET` columns in privilege tables incorrectly had a case-sensitive collation, resulting in

failure of assignments of values that did not have the same lettercase as given in the column definitions. The collation was changed to be case insensitive. ([Bug#7617](#) [<http://bugs.mysql.com/7617>])

- An expression that tested a case-insensitive character column against string constants that differed in lettercase could fail because the constants were treated as having a binary collation. (For example, `WHERE city='London' AND city='london'` could fail.) ([Bug#7098](#) [<http://bugs.mysql.com/7098>], [Bug#8690](#) [<http://bugs.mysql.com/8690>])
- The output of the `STATUS (\s)` command in `mysql` had the values for the server and client character sets reversed. ([Bug#7571](#) [<http://bugs.mysql.com/7571>])
- If the slave was running with `--replicate-*-table` options which excluded one temporary table and included another, and the two tables were used in a single `DROP TEMPORARY TABLE IF EXISTS` statement, as the ones the master automatically writes to its binary log upon client's disconnection when client has not explicitly dropped these, the slave could forget to delete the included replicated temporary table. Only the slave needs to be upgraded. ([Bug#8055](#) [<http://bugs.mysql.com/8055>])
- Treat user variables as having `IMPLICIT` derivation (coercibility) to avoid “Illegal mix of collations” errors when replicating user variables. ([Bug#6676](#) [<http://bugs.mysql.com/6676>])
- When setting integer system variables to a negative value with `SET VARIABLES`, the value was treated as a positive value modulo  $2^{32}$ . ([Bug#6958](#) [<http://bugs.mysql.com/6958>])
- Fixed a bug in bundled `readline` library that caused segmentation fault in `mysql` when user entered Shift+Enter. ([Bug#5672](#) [<http://bugs.mysql.com/5672>])
- Fix conversion of strings `->` double to get higher accuracy for floating point values that are integers, like: `123.45E+02` ([Bug#7840](#) [<http://bugs.mysql.com/7840>]).
- Fixed a bug in `MATCH ... AGAINST` in natural language mode that could cause a server crash if the `FULLTEXT` index was not used in a join (`EXPLAIN` did not show `fulltext` join mode) and the search query matched no rows in the table ([Bug#8522](#) [<http://bugs.mysql.com/8522>]).
- Platform and architecture information in version information produced for `--version` option on Windows was always `Win95/Win98 (i32)`. More accurately determine platform as `Win32` or `Win64` for 32-bit or 64-bit Windows, and architecture as `ia32` for x86, `ia64` for Itanium, and `axp` for Alpha. ([Bug#4445](#) [<http://bugs.mysql.com/4445>])
- Fixed a rare race condition which could lead to `FLUSH TABLES WITH READ LOCK` hanging. ([Bug#8682](#) [<http://bugs.mysql.com/8682>])
- Fixed a bug in replication that caused the master to stamp generated statements (such as `SET` commands) with an `error_code` intended only for another statement. This could happen, for example, when a statements generates a duplicate key error on the master but must be replicated. ([Bug#8412](#) [<http://bugs.mysql.com/8412>])
- If multiple semicolon-separated statements were received in a single packet, they were written to the binary log as a single event rather than as separate per-statement events. For a server serving as a replication master, this caused replication to fail when the event was sent to slave servers. ([Bug#8436](#) [<http://bugs.mysql.com/8436>])

## D.1.13. Changes in release 4.1.10 (12 February 2005)

**Note:** The security improvements related to creation of table files and to user-defined functions were made after MySQL 4.1.10 was released and are present in MySQL 4.1.10a. We would like to thank Stefano Di Paola <[stefano.dipaola@wisec.it](mailto:stefano.dipaola@wisec.it)> for making us aware of these.



Functionality added or changed:

- Added back faster subquery execution from 4.1.8. This adds also back a bug from 4.1.8 in comparing `NULL` to the value of a subquery. See [Section A.8.4, “Open Issues in MySQL”](#).
- Security improvement: The server creates `.frm`, `.MYD`, `.MYI`, `.MRG`, `.ISD`, and `.ISM` table files only if a file with the same name does not already exist. Thanks to Stefano Di Paola <[stefano.dipaola@wisec.it](mailto:stefano.dipaola@wisec.it)> for finding and informing us about this issue. (CVE-2005-0711 [<http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2005-0711>])
- Security improvement: User-defined functions should have at least one symbol defined in addition to the `xxx` symbol that corresponds to the main `xxx()` function. These auxiliary symbols correspond to the `xxx_init()`, `xxx_deinit()`, `xxx_reset()`, `xxx_clear()`, and `xxx_add()` functions. `mysqld` by default no longer loads UDFs unless they have at least one auxiliary symbol defined in addition to the main symbol. The `--allow-suspicious-udfs` option controls whether UDFs that have only an `xxx` symbol can be loaded. By default, the option is off. `mysqld` also checks UDF filenames when it reads them from the `mysql.func` table and rejects those that contain directory pathname separator characters. (It already checked names as given in [CREATE FUNCTION](#) statements.) See [Section 19.2.4.1, “UDF Calling Sequences for Simple Functions”](#), [Section 19.2.4.2, “UDF Calling Sequences for Aggregate Functions”](#), and [Section 19.2.4.6, “User-Defined Function Security Precautions”](#). Thanks to Stefano Di Paola <[stefano.dipaola@wisec.it](mailto:stefano.dipaola@wisec.it)> for finding and informing us about this issue. (CVE-2005-0709 [<http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2005-0709>], CVE-2005-0710 [<http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2005-0710>])
- Setting the connection collation to a value different from the server collation followed by a [CREATE TABLE](#) statement that included a quoted default value resulted in a server crash. (Bug#8235 [<http://bugs.mysql.com/8235>])
- Thread stack size was increased from 192KB to 256KB on Linux/IA-64 (too small stack size was causing server crashes on some queries). (Bug#8391 [<http://bugs.mysql.com/8391>])
- From the Windows distribution, predefined accounts without passwords for remote users (“root@%”, “@%”) were removed (other distributions never had them).
- Added `mysql_library_init()` and `mysql_library_end()` as synonyms for the `mysql_server_init()` and `mysql_server_end()` C API functions. `mysql_library_init()` and `mysql_library_end()` are `#define` symbols, but the names more clearly indicate that they should be called when beginning and ending use of a MySQL C API library no matter whether the application uses `libmysqlclient` or `libmysqld`. (Bug#6149 [<http://bugs.mysql.com/6149>])
- The server now issues a warning when `lower_case_table_names=2` and the data directory is on a case-sensitive filesystem, just as when `lower_case_table_names=0` on a case-insensitive filesystem. (Bug#7887 [<http://bugs.mysql.com/7887>])
- The server now issues a warning to the error log when it encounters older tables that contain character columns that might be interpreted by newer servers to have a different column length. (Bug#6913 [<http://bugs.mysql.com/6913>]) See [Section 2.11.1, “Upgrading from MySQL 4.0 to 4.1”](#), for a discussion of this problem and what to do about it.
- InnoDB: When MySQL/InnoDB is compiled on Mac OS X 10.2 or earlier, detect the operating system version at run time and use the `fcntl()` file flush method on Mac OS X versions 10.3 and later. In Mac OS X, `fsync()` does not flush the write cache in the disk drive, but the special `fcntl()` does; however, the flush request is ignored by some external devices. Failure to flush the buffers may cause severe database corruption at power outages.
- InnoDB: A shared record lock (`LOCK_REC_NOT_GAP`) is now taken for a matching record in the

foreign key check because inserts can be allowed into gaps.

- InnoDB: Relaxed locking in `INSERT...SELECT`, single table `UPDATE...SELECT` and single table `DELETE...SELECT` clauses when `innodb_locks_unsafe_for_binlog` is used and isolation level of the transaction is not serializable. InnoDB uses consistent read in these cases for a selected table.

Bugs fixed:

- `FOUND_ROWS()` returned an incorrect value after a `SELECT SQL_CALC_FOUND_ROWS DISTINCT` statement that selected constants and included `GROUP BY` and `LIMIT` clauses. (Bug#7945 [<http://bugs.mysql.com/7945>])
- Fixed a bug in cardinality estimations for `HASH` indexes of `TEMPORARY` tables created using `MEMORY` storage engine. As a result queries that were using this index (as shown by `EXPLAIN`) could have returned incorrect results. (Bug#8371 [<http://bugs.mysql.com/8371>])
- Corrected a problem with references to `DUAL` where statements such as `SELECT 1 AS a FROM DUAL` would succeed but statements such as `SELECT 1 AS a FROM DUAL LIMIT 1` would fail. (Bug#8023 [<http://bugs.mysql.com/8023>])
- Fixed a server crash caused by `DELETE FROM tbl_name ... WHERE ... ORDER BY tbl_name.col_name` when the `ORDER BY` column was qualified with the table name. (Bug#8392 [<http://bugs.mysql.com/8392>])
- `mysqld` had problems finding its language files if the `basedir` value was specified as a very long pathname. (Bug#8015 [<http://bugs.mysql.com/8015>])
- Updates were being written to the binary log when there were `binlog-do-db` or `binlog-ignore-db` options even when there was no current database, contrary to Section 5.11.4, “The Binary Log”. (Bug#6749 [<http://bugs.mysql.com/6749>])
- Fixed conversion of floating-point values to character fields when the absolute value of the float was less than 1, and also fixed calculation of length for negative values. (Bug#7774 [<http://bugs.mysql.com/7774>])
- Column headers in query results retrieved from the query cache could be corrupted when a non-4.1 client was served a result originally generated for a 4.1 client. The query cache was not keeping track of which client/server protocol was being used. (Bug#6511 [<http://bugs.mysql.com/6511>])
- Fixed `LOAD INDEX` statement to actually load index in memory. (Bug#8452 [<http://bugs.mysql.com/8452>])
- If multiple prepared statements were executed without retrieving their results, executing one of them again would cause the client program to crash. (Bug#8330 [<http://bugs.mysql.com/8330>])
- Non-numeric values inserted into a `YEAR` column were being stored as 2000 rather than as 0000. (Bug#6067 [<http://bugs.mysql.com/6067>])
- Fixed a failure of multiple-table updates to replicate properly on slave servers when `-replicate-*-table` options had been specified. (Bug#7011 [<http://bugs.mysql.com/7011>])
- `mysql_stmt_close()` C API function was not clearing an error indicator when a previous prepare call failed, causing subsequent invocations of error-retrieving calls to indicate spurious error values. (Bug#7990 [<http://bugs.mysql.com/7990>])
- Fixed failure of `CREATE TABLE ... LIKE` Windows when the source or destination table was

located in a symlinked database directory. ([Bug#6607](http://bugs.mysql.com/6607) [<http://bugs.mysql.com/6607>])

- With `lower_case_table_names` set to 1, `mysqldump` on Windows could write the same table name in different lettercase for different SQL statements. Fixed so that consistent lettercase is used. ([Bug#5185](http://bugs.mysql.com/5185) [<http://bugs.mysql.com/5185>]) `HAVING` that referred to `RAND()` or a user-defined function in the `SELECT` part through an alias could cause a crash or wrong value. ([Bug#8216](http://bugs.mysql.com/8216) [<http://bugs.mysql.com/8216>])
- If one used `CONVERT_TZ()` function in `SELECT`, which in its turn was used in `CREATE TABLE` statements, then system time zone tables were added to list of tables joined in `SELECT` and thus erroneous result was produced. ([Bug#7899](http://bugs.mysql.com/7899) [<http://bugs.mysql.com/7899>])
- Fixed a bug in `CONV()` function returning unsigned `BIGINT` number (third argument is positive, and return value does not fit in 32 bits). ([Bug#7751](http://bugs.mysql.com/7751) [<http://bugs.mysql.com/7751>])
- Fixed a failure of the `IN()` operator to return correct result if all values in the list were constants and some of them were using substring functions, for example, `LEFT()`, `RIGHT()`, or `MID()`.
- Fixed problem with `SHOW INDEX` reporting `Sub_part` values in bytes rather than characters for columns with a multi-byte character set. ([Bug#7943](http://bugs.mysql.com/7943) [<http://bugs.mysql.com/7943>])
- Fixed a crash in `CONVERT_TZ()` function when its second or third argument was from a `const` table (see [Section 7.2.1, “Optimizing Queries with EXPLAIN”](#)). ([Bug#7705](http://bugs.mysql.com/7705) [<http://bugs.mysql.com/7705>])
- Correct a problem with `mysql_config`, which was failing to produce proper `zlib` option for linking under some circumstances. ([Bug#6273](http://bugs.mysql.com/6273) [<http://bugs.mysql.com/6273>])
- Fixed a problem with calculation of number of columns in row comparison against a subquery. ([Bug#8020](http://bugs.mysql.com/8020) [<http://bugs.mysql.com/8020>])
- Fixed erroneous output resulting from `SELECT DISTINCT` combined with a subquery and `GROUP BY`. ([Bug#7946](http://bugs.mysql.com/7946) [<http://bugs.mysql.com/7946>])
- Fixed server crash in comparing a nested row expression (for example `row(1, (2, 3))`) with a subquery. ([Bug#8022](http://bugs.mysql.com/8022) [<http://bugs.mysql.com/8022>])
- Fixed server crash resulting from certain correlated subqueries with forward references (referring to an alias defined later in the outer query). ([Bug#8025](http://bugs.mysql.com/8025) [<http://bugs.mysql.com/8025>])
- Fixed server crash resulting from re-execution of prepared statements containing subqueries. ([Bug#8125](http://bugs.mysql.com/8125) [<http://bugs.mysql.com/8125>])
- Removed a dependence of boolean full-text search on `--default-character-set` option. ([Bug#8159](http://bugs.mysql.com/8159) [<http://bugs.mysql.com/8159>])
- Fixed a crash in a boolean full-text search in certain joins. ([Bug#8234](http://bugs.mysql.com/8234) [<http://bugs.mysql.com/8234>])
- Fixed erroneous comparison where strings that began with `CHAR(31)` were considered equal to the empty string. ([Bug#8134](http://bugs.mysql.com/8134) [<http://bugs.mysql.com/8134>])
- Add description of `debug` command to `mysqladmin` help output. ([Bug#8207](http://bugs.mysql.com/8207) [<http://bugs.mysql.com/8207>])
- `percona.exe` was always returning “Unknown error” on Windows. See [Section 8.17, “percona — Explain Error Codes”](#). ([Bug#7390](http://bugs.mysql.com/7390) [<http://bugs.mysql.com/7390>])
- Modify `SET` statements produced by `mysqldump` to write quoted strings using single quotes rather than double quotes. This avoids problems if the dump file is reloaded while the `ANSI_QUOTES`

SQL mode is in effect. (Bug#8148 [<http://bugs.mysql.com/8148>])

- Fixed a bug where `ALTER TABLE` improperly would accept an index on a `TIMESTAMP` column that `CREATE TABLE` would reject. (Bug#7884 [<http://bugs.mysql.com/7884>])
- Fixed a bug in multiple-table `UPDATE` statements that could cause spurious `Table '#sql_....' is full` errors if the number of rows to update is big enough. (Bug#7788 [<http://bugs.mysql.com/7788>])
- Fixed a problem where `SHOW INDEX` on a `MERGE` table could crash a debugging version of the server. (Bug#7377 [<http://bugs.mysql.com/7377>])
- Fixed a problem where adding an `ORDER BY` clause for an indexed column would cause a `SELECT` to return an empty result. (Bug#7331 [<http://bugs.mysql.com/7331>])
- Fixed a problem where `ALTER TABLE` on a `TEMPORARY` table with a mixed-lettercase name could cause the table to disappear when `lower_case_table_names` was set to 2. (Bug#7261 [<http://bugs.mysql.com/7261>])
- Fixed a problem with key cache statistics being reported incorrectly by the server after receipt of a `SIGHUP` signal. (Bug#4285 [<http://bugs.mysql.com/4285>])
- Fixed a problem that caused `mysql_stmt_prepare()` to be very slow when used in client programs on Windows. (Bug#5787 [<http://bugs.mysql.com/5787>])
- For indexes, `SHOW CREATE TABLE` now displays the index type even if it is the default, for storage engines that support multiple index types. (Bug#7235 [<http://bugs.mysql.com/7235>])
- Fixed a bug where the use of `GROUP_CONCAT()` with `HAVING` caused a server crash. (Bug#7769 [<http://bugs.mysql.com/7769>])
- Fixed a bug where comparing the result of a subquery to a non-existent column caused a server crash on Windows. (Bug#7885 [<http://bugs.mysql.com/7885>])
- Fixed a bug which caused `TIMEDIFF()` function to return wrong results if one of its arguments had non-zero microsecond part (Bug#7586 [<http://bugs.mysql.com/7586>]).
- Fixed a bug which caused `TIMESTAMP` columns with display width specified to be not identical to `DATETIME` columns when server was run in `MAXDB` mode (Bug#7418 [<http://bugs.mysql.com/7418>]).
- Fixed a bug in `UNION` statements that resulted in the wrong number of the examined rows reported in the slow query log.
- Fixed a bug in a combination of `-not` and `trunc*` operators of full-text search. Using more than one truncated negative search term, was causing empty result set.
- InnoDB: Fixed a bug introduced in 4.1.9 to the Windows version if you used `innodb_file_per_table`. `mysqld` would stop and complain about Windows error number 87 in a file operation. (See the Bugs database or the 4.1.9 change notes about a workaround for that bug in 4.1.9). (Bug#8021 [<http://bugs.mysql.com/8021>])
- InnoDB: Corrected the handling of trailing spaces in the `ucs2` character set. (Bug#7350 [<http://bugs.mysql.com/7350>])
- InnoDB: Use native `tmpfile()` function on Netware. All InnoDB temporary files are created under `sys:\tmp`. Previously, InnoDB temporary files were never deleted on Netware.
- InnoDB: Fix a race condition that could cause the assertion `space->n_pending_flushes ==`

0 to fail in `fil0fil.c`, in `fil_space_free()`, in `DROP TABLE` or in `ALTER TABLE`.

- InnoDB: `ALTER TABLE ... ADD CONSTRAINT PRIMARY KEY ...` complained about bad foreign key definition. (Bug#7831 [<http://bugs.mysql.com/7831>])
- InnoDB: Fix a theoretical hang over the adaptive hash latch in InnoDB if one runs `INSERT ... SELECT ...` (binlog not enabled), or a multiple-table `UPDATE` or `DELETE`, and only the read tables are InnoDB type, the rest are MyISAM. (Bug#7879 [<http://bugs.mysql.com/7879>])
- Fixed a bug in `max_heap_table_size` handling, that resulted in `Table is full` error when the table was still smaller than the limit. (Bug#7791 [<http://bugs.mysql.com/7791>]).
- Fixed a symlink vulnerability in the `mysqlaccess` script. Reported by Javier Fernandez-Sanguino Pena and [Debian Security Audit Team](http://www.debian.org/security/audit) [<http://www.debian.org/security/audit>]. (CVE-2005-0004 [<http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2005-0004>])
- `mysqlbinlog` forgot to add backquotes around the collation of user variables (causing later parsing problems as `BINARY` is a reserved word). (Bug#7793 [<http://bugs.mysql.com/7793>])
- Ensured that `mysqldump --single-transaction` sets its transaction isolation level to `REPEATABLE READ` before proceeding (otherwise if the MySQL server was configured to run with a default isolation level lower than `REPEATABLE READ` it could give an inconsistent dump). (Bug#7850 [<http://bugs.mysql.com/7850>])
- Changed `mysql` client so that including `\p` as part of a prompt command uses the name of the shared memory connection when the connection is using shared memory. (Bug#7922 [<http://bugs.mysql.com/7922>])
- Fixed a problem in the server where executing a multi-statement query more than once with the query cache active could yield incorrect result sets. (Bug#7966 [<http://bugs.mysql.com/7966>])
- Fixed that a 4.1.10 slave can connect to a master < 3.23.50 without hanging (the reason for the hang is a bug in these quite old masters -- `SELECT @@unknown_var` hangs them -- which was fixed in MySQL 3.23.50). (Bug#7965 [<http://bugs.mysql.com/7965>])
- Fixed a bug where MySQL was allowing concurrent updates (inserts, deletes) to a table if binary logging is enabled. Changed to ensure that all updates are executed in a serialized fashion, because they are executed serialized when binlog is replayed. (Bug#7879 [<http://bugs.mysql.com/7879>])

## D.1.14. Changes in release 4.1.9 (11 January 2005)

Functionality added or changed:

- `mysqld_safe` no longer tests for the presence of the data directory when using a relatively located server binary. It just assumes the directory is there, and fails to start up if it is not. This allows the data directory location to be specified on the command line, and avoids running a server binary that was not intended. (Bug#7249 [<http://bugs.mysql.com/7249>])
- The naming scheme of the Windows installation packages has changed slightly:
  - The platform suffix was changed from `-win` to `-win32`
  - The product descriptions `-noinstall` and `-essential` have been moved in front of the version number

Examples: `mysql-essential-4.1.9-win32.msi`, `mysql-noinstall-4.1.9-win32.zip` See [Section 2.3, "Installing MySQL on Windows"](#).

- The Mac OS X 10.3 installation disk images now include a MySQL Preference Pane for the Mac OS X Control Panel that enables the user to start and stop the MySQL server via the GUI and activate and deactivate the automatic MySQL server startup on bootup.
- The `MySQL-shared-compat` Linux RPM now includes the 3.23 as well as the 4.0 `libmysql-client.so` shared libraries. (Bug#6342 [http://bugs.mysql.com/6342])
- `Seconds_Behind_Master` is `NULL` (which means “unknown”) if the slave SQL thread is not running, or if the slave I/O thread is not running or not connected to master. It is zero if the SQL thread has caught up with the I/O thread. It no longer grows indefinitely if the master is idle.
- InnoDB: Do not acquire an internal InnoDB table lock in `LOCK TABLES` if `AUTOCOMMIT=1`. This helps in porting old MyISAM applications to InnoDB. InnoDB table locks in that case caused deadlocks very easily.
- InnoDB: Print a more descriptive error and refuse to start InnoDB if the size of `ibdata` files is smaller than what is stored in the tablespace header; `innodb_force_recovery` overrides this.
- The MySQL server aborts immediately instead of simply issuing a warning if it is started with the `--log-bin` option but cannot initialize the binary log at startup (that is, an error occurs when writing to the binary log file or binary log index file).
- The binary log file and binary log index file now behave like MyISAM when there is a "disk full" or "quota exceeded" error. See Section A.4.3, “How MySQL Handles a Full Disk”.

Bugs fixed:

- Fixed problem where running `mysql_fix_privilege_tables` could result in grant table columns with too-short lengths if the server character set had been set to a multi-byte character set first. (Bug#7539 [http://bugs.mysql.com/7539])
- InnoDB: Fixed the **critical bug** if you enabled `innodb_file_per_table` in `my.cnf`. If you shut down `mysqld`, records could disappear from the secondary indexes of a table. Unfortunately, on Windows a new Bug#8021 [http://bugs.mysql.com/8021] was introduced. Windows users of `innodb_file_per_table` should put a line `innodb_flush_method=unbuffered` to their `my.cnf` or `my.ini` to work around 8021. (Bug#7496 [http://bugs.mysql.com/7496])
- InnoDB: Fixed a bug: 32-bit `mysqld` binaries built on HP-UX-11 did not work with InnoDB files greater than 2 GB in size. (Bug#6189 [http://bugs.mysql.com/6189])
- InnoDB: Return a sensible error code from `DISCARD TABLESPACE` if it fails because the table is referenced by a `FOREIGN KEY`.
- InnoDB: Fixed a bug: InnoDB failed to drop a table in the background drop queue if the table was referenced by a `FOREIGN KEY` constraint.
- InnoDB: Fixed a bug: if we dropped a table where an `INSERT` was waiting for a lock to check a `FOREIGN KEY` constraint, then an assertion would fail in `lock_reset_all_on_table()`.
- InnoDB: Fix a little bug: we looked at the physical size of a stored SQL `NULL` value from a wrong field in the index; this has probably caused no bugs visible to the user. It caused only some extra space to be used in some rare cases.
- InnoDB: Use the `fcntl()` file flush method on Mac OS X versions 10.3 and up. Apple had disabled `fsync()` in Mac OS X for internal disk drives, which caused corruption at power outages.
- `mysqladmin password` now checks whether the server has `--old-passwords` turned on or

predates 4.1 and uses the old-format password if so. ([Bug#7451](http://bugs.mysql.com/7451) [<http://bugs.mysql.com/7451>])

- Added a `--default-character-set` option to `mysqladmin` to avoid problems when the default character set is not `latin1`. ([Bug#7524](http://bugs.mysql.com/7524) [<http://bugs.mysql.com/7524>])
- Fix a problem with truncation of `FLOAT` values. ([Bug#7361](http://bugs.mysql.com/7361) [<http://bugs.mysql.com/7361>])
- Fixed a bug in `PROCEDURE ANALYSE ( )`, which did not quote some `ENUM` values properly. ([Bug#2813](http://bugs.mysql.com/2813) [<http://bugs.mysql.com/2813>])
- Fixed a bug that caused incorrect results for complex datetime expressions containing casts of datetime values to `TIME` or `DATE` values. ([Bug#6914](http://bugs.mysql.com/6914) [<http://bugs.mysql.com/6914>])
- Include compression library flags in the output from `mysql_config --lib_r`. ([Bug#7021](http://bugs.mysql.com/7021) [<http://bugs.mysql.com/7021>])
- Corrected a problem with `mysql_config` not producing all relevant flags from `CFLAGS`. ([Bug#6964](http://bugs.mysql.com/6964) [<http://bugs.mysql.com/6964>])
- Corrected a problem with `mysqld_safe` not properly capturing output from `ps`. ([Bug#5878](http://bugs.mysql.com/5878) [<http://bugs.mysql.com/5878>])
- Fixed a bug that caused a linking failure when linking both the MySQL client library and IMAP library. ([Bug#7428](http://bugs.mysql.com/7428) [<http://bugs.mysql.com/7428>])
- Fixed table corruption bug when using `INSERT DELAYED` with prepared statements.
- Fixed a bug that caused microseconds to be gobbled from the string result of the `STR_TO_DATE` function, if there is some other specifier in the format string following `%f`. ([Bug#7458](http://bugs.mysql.com/7458) [<http://bugs.mysql.com/7458>])
- Made the MySQL server accept executing `SHOW CREATE DATABASE` even if the connection has an open transaction or locked tables. Refusing it made `mysqldump --single-transaction` sometimes fail to print a complete `CREATE DATABASE` statement for some dumped databases. ([Bug#7358](http://bugs.mysql.com/7358) [<http://bugs.mysql.com/7358>])
- Fixed that, when encountering a “disk full” or “quota exceeded” write error, `MyISAM` sometimes didn't sleep and retry the write, thus resulting in a corrupted table. ([Bug#7714](http://bugs.mysql.com/7714) [<http://bugs.mysql.com/7714>])
- Fixed that `--expire-log-days` was not honored if using only transactions. ([Bug#7236](http://bugs.mysql.com/7236) [<http://bugs.mysql.com/7236>])
- Fixed that a slave could crash after replicating many `ANALYZE TABLE`, `OPTIMIZE TABLE`, or `REPAIR TABLE` statements from the master. ([Bug#6461](http://bugs.mysql.com/6461) [<http://bugs.mysql.com/6461>], [Bug#7658](http://bugs.mysql.com/7658) [<http://bugs.mysql.com/7658>])

## D.1.15. Changes in release 4.1.8 (14 December 2004)

**Note:** Due to a `libtool`-related bug in the source distribution, the creation of shared `libmysqlclient` libraries was not possible (the resulting files were missing the `.so` file name extension). The file `ltmain.sh` was updated to fix this problem and the resulting source distribution was released as `mysql-4.1.8a.tar.gz`. This modification did not affect the binary packages. ([Bug#7401](http://bugs.mysql.com/7401) [<http://bugs.mysql.com/7401>])

Functionality added or changed:

- Automatic character set conversion formerly was done for operations that mix a column and a string such as assigning a string to a column, when this was possible without loss of information. Automatic conversion for operations that mix columns and strings has been expanded to cover many functions (such as `CONCAT()`) and assignment operators. This reduces the frequency of `Illegal mix of collations` errors.
- For `ALTER DATABASE`, the database name now can be omitted to apply the change to the default database.
- Added `WITH CONSISTENT SNAPSHOT` clause to `START TRANSACTION` to begin a transaction with a consistent read.
- Added `--order-by-primary` to `mysqldump`, to sort each table's data in a dump file. This may be useful when dumping a `MyISAM` table which will be loaded into an InnoDB table. Dumping a `MyISAM` table with this option is considerably slower than without.
- InnoDB: Do not periodically write `SHOW INNODB STATUS` information to a temporary file unless the configuration option `innodb_status_file=1` is set.
- InnoDB: Commit after every 10,000 copied rows when executing `ALTER TABLE`. This makes it much faster to recover from an aborted `ALTER TABLE` or `OPTIMIZE TABLE`.
- `FULLTEXT` index block size is changed to be 1024 instead of 2048.
- Added `--disable-log-bin` option to `mysqlbinlog`. Using this option you can disable binary logging for the statements produced by `mysqlbinlog`. That is, `mysqlbinlog --disable-log-bin <file_name> | mysql` won't write any statements to the MySQL server binary log.
- The `--master-data` option for `mysqldump` now takes an optional argument of 1 or 2 to produce a non-commented or commented `CHANGE MASTER TO` statement. The default is 1 for backward compatibility.
- `mysqldump --single-transaction --master-data` now is able to take an online (non-blocking) dump of InnoDB and report the corresponding binary log coordinates. This makes a backup suitable for point-in-time recovery, roll-forward or replication slave creation. See [Section 8.12, “mysqldump — A Database Backup Program”](#).
- Added `--lock-all-tables` to `mysqldump` to lock all tables by acquiring a global read lock.
- Added `--hex-lob` option to `mysqldump` for dumping binary string columns using hexadecimal notation.
- Added `mysql_hex_string()` C API function that hex-encodes a string.
- In the normal log MySQL now prints the log position for `Binlog Dump` requests.
- Added `[mysql_cluster]` section to `my.cnf` file for configuration settings specific to MySQL Cluster. `ndb-connectstring` variable moved here.
- A connection doing a rollback now displays "Rolling back" in the `State` column of `SHOW PROCESSLIST`.
- `mysqlbinlog` now prints an informative commented line (thread id, timestamp, server id, and so forth) before each `LOAD DATA INFILE`, like it does for other queries; unless `--short-form` is used.
- The statements `CREATE TABLE`, `TRUNCATE TABLE`, `DROP DATABASE`, and `CREATE DATABASE` cause an implicit commit.



Bugs fixed:

- A multiple-table `DELETE` could cause MySQL to crash when using `InnoDB` tables. ([Bug#5837](#) [<http://bugs.mysql.com/5837>], [Bug#6378](#) [<http://bugs.mysql.com/6378>])
- Some data definition statements (`CREATE TABLE` where the table was not a temporary table, `TRUNCATE TABLE`, `DROP DATABASE`, and `CREATE DATABASE`) were not being written to the binary log after a `ROLLBACK`. This also caused problems with replication. ([Bug#6883](#) [<http://bugs.mysql.com/6883>])
- Fixed incorrect referencing to column by name from subquery to outer query in case of using temporary table by outer query and placing subquery in the `WHERE` clause. ([Bug#7079](#) [<http://bugs.mysql.com/7079>])
- Fixed a bug in authentication code that allowed a malicious user to crash the server with specially crafted packets (using a modified client library). ([Bug#7187](#) [<http://bugs.mysql.com/7187>])
- Fixed a crashing bug in a string function `LEFT`, when this function is part of the expression which is used as `GROUP BY` field. ([Bug#7101](#) [<http://bugs.mysql.com/7101>])
- Fixed bug which caused MySQL to require privileges on system time zone description tables for implicit access to them (that is, if one set `time_zone` variable or used `CONVERT_TZ( )` function) in case when some table-level or column-level privileges existed. ([Bug#6765](#) [<http://bugs.mysql.com/6765>])
- `mysql_stmt_data_seek(stmt, 0)` now rewinds a counter and enables buffered rows to be re-fetched on the client side. ([Bug#6996](#) [<http://bugs.mysql.com/6996>])
- Fixed an insufficient privilege check in `SHOW CREATE TABLE` command. ([Bug#7043](#) [<http://bugs.mysql.com/7043>])
- Fixed a rare memory corruption (that resulted in a crash) in `MATCH ... AGAINST` on columns that use multi-byte character sets. ([Bug#6269](#) [<http://bugs.mysql.com/6269>])
- Fixed `NULL` processing in `ALL/SOME` subqueries. ([Bug#6247](#) [<http://bugs.mysql.com/6247>])
- Fixed execution of complex queries with subqueries. ([Bug#6406](#) [<http://bugs.mysql.com/6406>], [Bug#6841](#) [<http://bugs.mysql.com/6841>])
- Fixed initialization of some internal structures for first execution. ([Bug#6517](#) [<http://bugs.mysql.com/6517>])
- Backported a fix for the full-text interface from MySQL 5.0. ([Bug#6523](#) [<http://bugs.mysql.com/6523>])
- Fixed `NULL` value handling in case of empty results in subqueries. ([Bug#6806](#) [<http://bugs.mysql.com/6806>])
- Prevent adding `CREATE TABLE ... SELECT` query to the binary log when the insertion of new records partially failed. ([Bug#6682](#) [<http://bugs.mysql.com/6682>])
- `INSERT ... SELECT` no longer reports spurious "column truncated" warnings ([Bug#6284](#) [<http://bugs.mysql.com/6284>])
- Fixed a bug that could cause "Record has changed since last read in table" error message in some queries on `HEAP` tables that contain only one row. ([Bug#6748](#) [<http://bugs.mysql.com/6748>])
- `mysqld_safe` was in many cases ignoring any `--no-defaults`, `--defaults-file`, or `-`

`-defaults-extra-file` arguments. Those arguments are now honored, and this may change what options are passed to `mysqld` in some installations.

- The server was interpreting `CHAR BINARY` and `VARCHAR BINARY` columns from 4.0 tables as having the `BINARY` and `VARBINARY` data types. Now they are interpreted as `CHAR` and `VARCHAR` columns that have the binary collation of the column's character set. (This is the same way that `CHAR BINARY` and `VARCHAR BINARY` are handled for new tables created in 4.1.)
- Fixed spurious "duplicate key" error from `REPLACE` or `INSERT ... ON DUPLICATE KEY UPDATE` statements performing multiple-row insert in the table that had unique and full-text indexes. ([Bug#6784](http://bugs.mysql.com/6784) [<http://bugs.mysql.com/6784>])
- Fixed a bug in execution of subqueries in `SET` and `DO` statements which caused wrong results to be returned from subsequent queries. ([Bug#6462](http://bugs.mysql.com/6462) [<http://bugs.mysql.com/6462>])
- Fixed a bug that allowed server to accept datetime values with wrong year part. The server now also performs same checks for datetime values passed through `MYSQL_TIME` structures as for datetime values passed as strings. ([Bug#6266](http://bugs.mysql.com/6266) [<http://bugs.mysql.com/6266>])
- Fixed a bug with `INSERT` for a table with `FULLTEXT` indexes. Under rare circumstances, this could result in a corrupted table if words of different lengths may be considered equal. This is possible in some collations, for example, in `utf8_general_ci` or `latin1_german2_ci`. ([Bug#6265](http://bugs.mysql.com/6265) [<http://bugs.mysql.com/6265>])
- InnoDB: Do not intentionally crash `mysqld` if the buffer pool is exhausted by the lock table; return error 1206 instead. Do not intentionally crash `mysqld` if we cannot allocate the memory for the InnoDB buffer pool. ([Bug#6817](http://bugs.mysql.com/6817) [<http://bugs.mysql.com/6817>]) ([Bug#6827](http://bugs.mysql.com/6827) [<http://bugs.mysql.com/6827>])
- InnoDB: Let InnoDB's `FOREIGN KEY` parser to remove the `latin1` character `0xA0` from the end of an unquoted identifier. The EMS MySQL Manager in `ALTER TABLE` adds that character after a table name, which caused error 121 when we tried to add a new constraint.
- InnoDB: Refuse to open new-style tables created with MySQL 5.0.3 or later. ([Bug#7089](http://bugs.mysql.com/7089) [<http://bugs.mysql.com/7089>])
- InnoDB: Do not call `rewind()` when displaying `SHOW INNODB STATUS` information on `stderr`.
- InnoDB: Made the foreign key parser better aware of quotes. ([Bug#6340](http://bugs.mysql.com/6340) [<http://bugs.mysql.com/6340>])
- InnoDB: If one used `INSERT IGNORE` to insert several rows at a time, and the first inserts were ignored because of a duplicate key collision, then InnoDB in a replication slave assigned `AUTO_INCREMENT` values 1 bigger than in the master. This broke the MySQL replication. ([Bug#6287](http://bugs.mysql.com/6287) [<http://bugs.mysql.com/6287>])
- InnoDB: Fixed a bug: InnoDB ignored in `innodb_data_file_path` the `max` specification in `:autoextend:max:2000M`. This bug was introduced in 4.1.1.
- InnoDB: Fixed a bug: `innodb_locks_unsafe_for_binlog` still uses next-key locking ([Bug#6747](http://bugs.mysql.com/6747) [<http://bugs.mysql.com/6747>]). InnoDB used next-key locking when record matched completely to search tuple. This unnecessary next-key locking is now removed when `innodb_locks_unsafe_for_binlog` option is used.
- InnoDB: Fix two hangs: `FOREIGN KEY` constraints treated table and database names as case-insensitive. `RENAME TABLE t TO T` would hang in an endless loop if `t` had a foreign key constraint defined on it. Fix also a hang over the dictionary mutex that would occur if one tried in `ALTER TABLE` or `RENAME TABLE` to create a foreign key constraint name that collided with another

existing name. ([Bug#3478](http://bugs.mysql.com/3478) [<http://bugs.mysql.com/3478>])

- If `STMT_ATTR_UPDATE_MAX_LENGTH` is set for a prepared statement, `mysql_stmt_store_result()` updates `field->max_length` for numeric columns as well. ([Bug#6096](http://bugs.mysql.com/6096) [<http://bugs.mysql.com/6096>])
- Prepared statements now handle `ZEROFILL` when converting `integer` to `string`.
- Fixed crash when a call to `mysql_stmt_store_result()` occurred without a preceding call to `mysql_stmt_bind_result()`.
- Fixed crash in prepared statements when using `SELECT * FROM t1 NATURAL JOIN t2...`
- Fixed crash in prepared statements when using `SELECT ... PROCEDURE`.
- Fixed crash in prepared statements when using subqueries.
- `GROUP_CONCAT(...ORDER BY)` when used with prepared statements gave wrong sorting order.
- `CREATE TABLE created_table` didn't signal when table was created. This could cause a `DROP TABLE created_table` in another thread to wait "forever".
- Server warnings now are reset when you execute a prepared statement.
- Improved performance of identifier comparisons (if many tables or columns are specified).
- `OPTIMIZE TABLE`, `REPAIR TABLE`, and `ANALYZE TABLE` are now replicated without any error code in the binary log. ([Bug#5551](http://bugs.mysql.com/5551) [<http://bugs.mysql.com/5551>])
- `LOAD DATA INFILE` now works with option `replicate-rewrite-db`. ([Bug#6353](http://bugs.mysql.com/6353) [<http://bugs.mysql.com/6353>])
- Fixed a bug which caused a crash when only the slave I/O thread was stopped and started. ([Bug#6148](http://bugs.mysql.com/6148) [<http://bugs.mysql.com/6148>])
- Changed semantics of `CREATE/ALTER/DROP DATABASE` statements so that replication of `CREATE DATABASE` is possible when using `--binlog-do-db` and `--binlog-ignore-db`. ([Bug#6391](http://bugs.mysql.com/6391) [<http://bugs.mysql.com/6391>])
- If a connection had an open transaction but had done no updates to transactional tables (for example if had just done a `SELECT FOR UPDATE` then executed a non-transactional update, that update automatically committed the transaction (thus releasing InnoDB's row-level locks etc). ([Bug#5714](http://bugs.mysql.com/5714) [<http://bugs.mysql.com/5714>])
- If a connection was interrupted by a network error and did a rollback, the network error code got stored into the `BEGIN` and `ROLLBACK` binary log events; that caused superfluous slave stops. ([Bug#6522](http://bugs.mysql.com/6522) [<http://bugs.mysql.com/6522>])
- A sequence of `BEGIN` (or `SET AUTOCOMMIT=0`), `FLUSH TABLES WITH READ LOCK`, transactional update, `COMMIT`, `FLUSH TABLES WITH READ LOCK` could hang the connection forever and possibly the MySQL server itself. This happened for example when running the `innobackup` script several times. ([Bug#6732](http://bugs.mysql.com/6732) [<http://bugs.mysql.com/6732>])
- `mysqlbinlog` did not print `SET PSEUDO_THREAD_ID` statements in front of `LOAD DATA INFILE` statements inserting into temporary tables, thus causing potential problems when rolling forward these statements after restoring a backup. ([Bug#6671](http://bugs.mysql.com/6671) [<http://bugs.mysql.com/6671>])

## D.1.16. Changes in release 4.1.7 (23 October 2004: Produc-

## tion)

Functionality added or changed:

- `MOD( )` no longer rounds arguments with a fractional part to integers. Now it returns exact remainder after division. (Bug#6138 [<http://bugs.mysql.com/6138>])
- InnoDB: Added a startup option and settable system variable `innodb_table_locks` for making `LOCK TABLE` acquire also InnoDB locks. The default value is 1, which means that `LOCK TABLES` causes also InnoDB internally to take a table lock. In applications using `AUTOCOMMIT=1` and `LOCK TABLES`, InnoDB's internal table locks that were added in 4.0.20 and 4.1.2 can cause deadlocks. You can set `innodb_table_locks=0` in `my.cnf` to remove that problem. (Bug#3299 [<http://bugs.mysql.com/3299>], Bug#5998 [<http://bugs.mysql.com/5998>])
- See Section 14.2.17, “Restrictions on InnoDB Tables”. (Bug#3299 [<http://bugs.mysql.com/3299>], Bug#5998 [<http://bugs.mysql.com/5998>]) InnoDB: `SHOW TABLE STATUS` now shows the creation time of the table for InnoDB. Note that this timestamp might not be the correct time because, for example, `ALTER TABLE` changes this timestamp.
- InnoDB: If `innodb_thread_concurrency` would be exceeded, let a thread sleep 10 ms before entering the FIFO queue; previously, the value was 50 ms.

Bugs fixed:

- Fixed a bug with `FOUND_ROWS( )` used together with `LIMIT` clause in prepared statements. (Bug#6088 [<http://bugs.mysql.com/6088>])
- Fixed a bug with `NATURAL JOIN` in prepared statements. (Bug#6046 [<http://bugs.mysql.com/6046>]).
- Fixed a bug in join of tables from different databases having columns with identical names (prepared statements). (Bug#6050 [<http://bugs.mysql.com/6050>])
- Now implicit access to system time zone description tables (which happens when you set the `time_zone` variable or use `CONVERT_TZ( )` function) does not require any privileges. (Bug#6116 [<http://bugs.mysql.com/6116>])
- Fixed a bug which caused the server to crash when the deprecated `libmysqlclient` function `mysql_create_db( )` was called. (Bug#6081 [<http://bugs.mysql.com/6081>])
- Fixed `REVOKE ALL PRIVILEGES, GRANT OPTION FROM user` so that all privileges are revoked correctly. (Bug#5831 [<http://bugs.mysql.com/5831>]). This corrects a case that the fix in 4.1.6 could miss.
- Fixed crash when selecting from a `HEAP` table with `key_column IS NOT NULL`. This could also cause a crash if not all index parts were used. (Bug#6082 [<http://bugs.mysql.com/6082>])
- Fixed a bug that could cause `MyISAM` index corruption when key values start with character codes below `BLANK`. This was caused by the new key sort order in 4.1. (Bug#6151 [<http://bugs.mysql.com/6151>])
- InnoDB: Fixed a bug in `LOAD DATA INFILE...REPLACE` printing duplicate key error when executing the same load query several times. (Bug#5835 [<http://bugs.mysql.com/5835>])
- Fixed a bug in the prepared statements protocol when wrong metadata was sent for `SELECT` statements not returning a result set (such as `SELECT ... INTO OUTFILE`). (Bug#6059

[<http://bugs.mysql.com/6059>])

- Fixed bug which allowed one to circumvent missing UPDATE privilege if one had INSERT and SELECT privileges for table with primary key. (Bug#6173 [<http://bugs.mysql.com/6173>])
- Fixed a bug in `libmysqlclient` with wrong conversion of negative time values to strings. (Bug#6049 [<http://bugs.mysql.com/6049>]).
- Fixed a bug in `libmysqlclient` with wrong conversion of zero date values (0000-00-00) to strings. (Bug#6058 [<http://bugs.mysql.com/6058>])
- Fixed a bug that caused the server to crash on attempt to prepare a statement with `RAND(?)`. (Bug#5985 [<http://bugs.mysql.com/5985>])
- Fixed a bug with handling of `DATE`, `TIME`, and `DATETIME` columns in the binary protocol. The problem is compiler-specific and could have been observed on HP-UX, AIX, Solaris9, when compiling with native compiler. (Bug#6025 [<http://bugs.mysql.com/6025>])
- Fixed a bug with handling of `TINYINT` columns in the binary protocol. The problem is specific to platforms where the C compiler has the `char` data type unsigned by default. (Bug#6024 [<http://bugs.mysql.com/6024>])
- InnoDB: Fixed problem introduced in MySQL 4.0.21 where a connection starting a transaction, doing updates, then `FLUSH TABLES WITH READ LOCK`, then `COMMIT`, would cause replication slaves to stop (complaining about error 1223). Bug surfaced when using the InnoDB `innobackup` script. (Bug#5949 [<http://bugs.mysql.com/5949>])
- InnoDB: Release the dictionary latch during a long cascaded `FOREIGN KEY` operation, so that we do not starve other users doing `CREATE TABLE` or other DDL operation. This caused a notorious 'Long semaphore wait' message to be printed to the `.err` log. (Bug#5961 [<http://bugs.mysql.com/5961>])

## D.1.17. Changes in release 4.1.6 (10 October 2004)

Functionality added or changed:

- Added option `--sigint-ignore` to the `mysql` command line client to make it ignore `SIGINT` signals (typically the result of the user pressing Control-C).
- InnoDB: Added the startup option and settable global variable `innodb_max_purge_lag` for delaying `INSERT`, `UPDATE` and `DELETE` operations when the purge operations are lagging. The default value of this parameter is zero, meaning that there are no delays. See [Section 14.2.13](#), “Implementation of Multi-Versioning”.
- InnoDB: The `innodb_autoextend_increment` startup option that was introduced in release 4.1.5 was made a settable global variable. (Bug#5736 [<http://bugs.mysql.com/5736>])
- InnoDB: If `DROP TABLE` is invoked on an InnoDB table for which the `.ibd` file is missing, print to error log that the table was removed from the InnoDB data dictionary, and allow MySQL to delete the `.frm` file. Maybe `DROP TABLE` should issue a warning in this case.
- `TIMESTAMP` columns now can store `NULL` values. To create such a column, you must explicitly specify the `NULL` attribute in the column specification. (Unlike all other data types, `TIMESTAMP` columns are `NOT NULL` by default.)
- Now if `ALTER TABLE` converts one `AUTO_INCREMENT` column to another `AUTO_INCREMENT`

column it preserves zero values (this includes the case that we don't change such column at all).

- Now if `ALTER TABLE` converts some column to `TIMESTAMP NOT NULL` column it converts `NULL` values to current timestamp value (One can still get old behavior by setting system `TIMESTAMP` variable to zero).
- On Windows, the MySQL configuration files included in the package now use `.ini` instead of `.cnf` as the file name suffix.

Bugs fixed:

- Fixed a bug that caused the server to crash on attempt to execute a prepared statement with a sub-query inside a boolean expression. (Bug#5987 [http://bugs.mysql.com/5987])
- Fixed a bug that caused the server to sometimes choose non-optimal execution plan for a prepared statement executed with changed placeholder values. (Bug#6042 [http://bugs.mysql.com/6042])
- InnoDB: Make the check for excessive semaphore waits tolerate glitches in the system clock (do not crash the server if the system time is adjusted while InnoDB is under load.). (Bug#5898 [http://bugs.mysql.com/5898])
- InnoDB: Fixed a bug in the InnoDB `FOREIGN KEY` parser that prevented `ALTER TABLE` of tables containing '#' in their names. (Bug#5856 [http://bugs.mysql.com/5856])
- InnoDB: Fixed a bug that prevented `ALTER TABLE t DISCARD TABLESPACE` from working. (Bug#5851 [http://bugs.mysql.com/5851])
- InnoDB: `SHOW CREATE TABLE` now obeys the `SET SQL_MODE=ANSI` and `SET SQL_QUOTE_SHOW_CREATE=0` settings. (Bug#5292 [http://bugs.mysql.com/5292])
- InnoDB: Fixed a bug that caused `CREATE TEMPORARY TABLE ... ENGINE=InnoDB` to terminate `mysqld` when running in `innodb_file_per_table` mode. Per-table tablespaces for temporary tables from now on are created in the temporary directory of `mysqld`. (Bug#5137 [http://bugs.mysql.com/5137])
- InnoDB: Fixed some (not all) UTF-8 bugs in column prefix indexes. (Bug#5975 [http://bugs.mysql.com/5975])
- InnoDB: If one updated a column so that its size changed, or updated it to an externally stored (`TEXT` or `BLOB`) value, then ANOTHER externally stored column would show up as 512 bytes of good data + 20 bytes of garbage in a consistent read that fetched the old version of the row. (Bug#5960 [http://bugs.mysql.com/5960])
- InnoDB: Change error code to `HA_ERR_ROW_IS_REFERENCED` if we cannot `DROP` a parent table referenced by a `FOREIGN KEY` constraint; this error number is less misleading than the previous number `HA_ERR_CANNOT_ADD_FOREIGN`, but misleading still. (Bug#6202 [http://bugs.mysql.com/6202])
- Fixed `REVOKE ALL PRIVILEGES, GRANT OPTION FROM user` so that all privileges are revoked correctly. (Bug#5831 [http://bugs.mysql.com/5831])
- Fixed a bug that caused the server to crash when character set conversion was implicitly used in prepared mode; for example, as in `'abc' LIKE CONVERT('abc' as utf8)`. (Bug#5688 [http://bugs.mysql.com/5688])
- The `mysql_change_user()` C API function now frees all prepared statements associated with the connection. (Bug#5315 [http://bugs.mysql.com/5315])

- Fixed a bug when inserting `NULL` into an `AUTO_INCREMENT` column failed, when using prepared statements. (Bug#5510 [<http://bugs.mysql.com/5510>])
- Fixed slave SQL thread so that the `SET COLLATION_SERVER. . .` statements it replicates don't advance its position (so that if it gets interrupted before the actual update query, it later redoes the `SET`). (Bug#5705 [<http://bugs.mysql.com/5705>])
- Fixed that if the slave SQL thread found a syntax error in a query (which should be rare, as the master parsed it successfully), it stops. (Bug#5711 [<http://bugs.mysql.com/5711>])
- Fixed that if a write to a `MyISAM` table fails because of a full disk or an exceeded disk quota, it prints a message to the error log every 10 minutes, and waits until disk space becomes available. (Bug#3248 [<http://bugs.mysql.com/3248>])
- Now MySQL does not prefer columns, which are mentioned in select list but are renamed, over columns from other tables participating in `FROM` clause when it resolves `GROUP BY` clause (for example, `SELECT t1.a AS c FROM t1, t2 ORDER BY a` produces an error if both `t1` and `t2` tables contain `a` column). (Bug#4302 [<http://bugs.mysql.com/4302>])
- Behavior of `ALTER TABLE` converting column containing `NULL` values to `AUTO_INCREMENT` column is no longer affected by `NO_AUTO_VALUE_ON_ZERO` mode. (Bug#5915 [<http://bugs.mysql.com/5915>]).

## D.1.18. Changes in release 4.1.5 (16 September 2004)

Functionality added or changed:

- **InnoDB**: Added configuration option `innodb_autoextend_increment` for setting the size in megabytes by which **InnoDB** tablespaces are extended when they become full. The default value is 8, corresponding to the fixed increment of 8MB in previous versions of MySQL.
- **InnoDB**: The new Windows installation wizard of MySQL makes **InnoDB** as the MySQL default table type on Windows, unless explicitly specified otherwise. Note that it places the `my.ini` file in the installation directory of the MySQL server. See [Section 2.3.4.14](#), “The Location of the `my.ini` File”.

Bugs fixed:

- Fixed a bug which caused the server to crash on attempt to execute a prepared statement with `BETWEEN ? AND ?` and a datetime column. (Bug#5748 [<http://bugs.mysql.com/5748>])
- Fixed name resolving of external fields of subqueries if subquery placed in select list of query with grouping. (Bug#5326 [<http://bugs.mysql.com/5326>])
- Fixed detection of using same table for updating and selecting in multi-update queries. (Bug#5455 [<http://bugs.mysql.com/5455>])
- The values of the `max_sort_length`, `sql_mode`, and `group_concat_max_len` system variables now are stored in the query cache with other query information to avoid returning an incorrect result from the query cache. (Bug#5394 [<http://bugs.mysql.com/5394>]) (Bug#5515 [<http://bugs.mysql.com/5515>])
- Fixed syntax analyzer with `sql_mode=IGNORE_SPACE`. It happened to take phrases like `default .07` as `identifier.identifier`. (Bug#5318 [<http://bugs.mysql.com/5318>])

- Fixed illegal internal field length of user variables of integer type. This showed up when creating a table as `SELECT @var_name`. (Bug#4788 [<http://bugs.mysql.com/4788>])
- Fixed a buffer overflow in prepared statements API (libmysqlclient) when a statement containing thousands of placeholders was executed. (Bug#5194 [<http://bugs.mysql.com/5194>])
- Fixed a bug in the server when after reaching a certain limit of prepared statements per connection (97), statement ids began to overlap, so occasionally wrong statements were chosen for execution. (Bug#5399 [<http://bugs.mysql.com/5399>])
- Fixed a bug in prepared statements when `LIKE` used with arguments in different character sets crashed server on first execute. (Bug#4368 [<http://bugs.mysql.com/4368>])
- Fixed a bug in prepared statements when providing '0000-00-00' date to a parameter lead to server crash. (Bug#4231 [<http://bugs.mysql.com/4231>], Bug#4562 [<http://bugs.mysql.com/4562>])
- Fixed a bug in `OPTIMIZE TABLE` that could cause table corruption on `FULLTEXT` indexes. (Bug#5327 [<http://bugs.mysql.com/5327>])
- InnoDB: Fixed a bug that InnoDB only allowed a maximum of 1000 connections inside InnoDB at the same time. A higher number could cause an assertion failure in `sync0arr.c`, line 384. Now we allow 1000, 10000, or 50000, depending on the buffer pool size. (Bug#5414 [<http://bugs.mysql.com/5414>])

## D.1.19. Changes in release 4.1.4 (26 August 2004: Gamma)

**Note:** To fix a compile problem on systems that do not have `automake` 1.7 installed, an updated 4.1.4a source tarball has been published. In addition to resolving this `automake` dependency (Bug#5319 [<http://bugs.mysql.com/5319>]), it also fixes some reported `libedit` compile errors when using a non-`gcc` compiler (Bug#5353 [<http://bugs.mysql.com/5353>]).

Functionality added or changed:

- Added the `CSV` storage engine.
- Made internal representation of `TIMESTAMP` values in `InnoDB` in 4.1 to be the same as in 4.0. This difference resulted in incorrect datetime values in `TIMESTAMP` columns in `InnoDB` tables after an upgrade from 4.0 to 4.1. (Bug#4492 [<http://bugs.mysql.com/4492>]) **Warning: extra steps during upgrade required!** Unfortunately this means that if you are upgrading from 4.1.x, where  $x \leq 3$ , to 4.1.4 you should use `mysqldump` for saving and then restoring your `InnoDB` tables with `TIMESTAMP` columns.
- The `mysqld-opt` Windows server was renamed to `mysqld`. This completes the Windows server renaming begun in MySQL 4.1.2. See Section 2.3.8, “Selecting a MySQL Server type”.
- Added Latin language collations for the `ucs2` and `utf8` Unicode character sets. These are called `ucs2_roman_ci` and `utf8_roman_ci`.
- Corrected the name of the Mac OS X StartupItem script (it must match the name of the subdirectory, which was renamed to `MySQLCOM` in MySQL 4.1.2). Thanks to Bryan McCormack for reporting this.
- Added `--start-datetime`, `--stop-datetime`, `--start-position`, and `--stop-position` options to `mysqlbinlog`. These make point-in-time recovery easier.
- Killing a `CHECK TABLE` statement does not result in the table being marked as “corrupted” any



more; the table remains as if `CHECK TABLE` had not even started. See [Section 13.5.5.3](#), “`KILL Syntax`”.

- Made the MySQL server ignore `SIGHUP` and `SIGQUIT` on Mac OS X 10.3. This is needed because under this OS, the MySQL server receives lots of these signals (reported as [Bug#2030](#) [<http://bugs.mysql.com/2030>]).
- Support of usage of column aliases qualified by table name or alias in `ORDER BY` and `GROUP BY` was dropped. For example the following query `SELECT a AS b FROM t1 ORDER BY t1.b` is not allowed. One should use `SELECT a AS b FROM t1 ORDER BY t1.a` or `SELECT a AS b FROM t1 ORDER BY b` instead. This was non-standard (since aliases are defined on query level not on table level) and caused problems with some queries.

Bugs fixed:

- Fixed a bug that caused libmysql to crash when attempting to fetch a value of `MEDIUMINT` column. ([Bug#5126](#) [<http://bugs.mysql.com/5126>])
- Fixed a bug that caused the MySQL server to crash when attempting to execute a prepared statement with `SELECT ... INTO @var` for a second time. ([Bug#5034](#) [<http://bugs.mysql.com/5034>])
- Fixed execution of optimized `IN` subqueries that use compound indexes. ([Bug#4435](#) [<http://bugs.mysql.com/4435>])
- Prohibited resolving of table fields in inner queries if fields do not take part in grouping for queries with grouping (inside aggregate function arguments, all table fields are still allowed). ([Bug#4814](#) [<http://bugs.mysql.com/4814>])
- Fixed a crash after `SLAVE STOP` if the IO thread was in a special state. ([Bug#4629](#) [<http://bugs.mysql.com/4629>])
- Fixed an old bug in concurrent accesses to `MERGE` tables (even one `MERGE` table and `MyISAM` tables), that could have resulted in a crash or hang of the server. ([Bug#2408](#) [<http://bugs.mysql.com/2408>], [CVE-2004-0837](#) [<http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2004-0837>])
- Fixed a bug that caused server crash on attempt to execute for a second time a prepared statement with `NOT` in `WHERE` or `ON` clauses. ([Bug#4912](#) [<http://bugs.mysql.com/4912>])
- `MATCH ... AGAINST` now works in a subquery. ([Bug#4769](#) [<http://bugs.mysql.com/4769>])
- Fixed a bug that omitted the `.err` extension of the error log file (`--log-error`) when the host-name contained a domain name. The domain name is now replaced by the extension. ([Bug#4997](#) [<http://bugs.mysql.com/4997>])
- Fixed a crash in `myisamchk`. ([Bug#4901](#) [<http://bugs.mysql.com/4901>])
- Fixed a bug which caused server crash if one used the `CONVERT_TZ( )` function with time zone described in database as parameter and this time zone was not used before. ([Bug#4508](#) [<http://bugs.mysql.com/4508>])
- Support for `%T`, `%r`, `%V`, `%v` and `%X`, `%x` format specifiers was added to `STR_TO_DATE( )` function. ([Bug#4756](#) [<http://bugs.mysql.com/4756>])
- Fixed a bug (hang) in `NATURAL JOIN` where joined table had no common column. ([Bug#4807](#) [<http://bugs.mysql.com/4807>])

- Fixed a crash caused by `UNHEX(NULL)`. (Bug#4441 [<http://bugs.mysql.com/4441>])
- `mysql_fix_privilege_tables` didn't correctly handle the argument of its `--password=password_val` option. (Bug#4240 [<http://bugs.mysql.com/4240>], Bug#4543 [<http://bugs.mysql.com/4543>])
- Fixed that `mysqlbinlog --read-from-remote-server` sometimes couldn't accept 2 binary logs on command line. (Bug#4507 [<http://bugs.mysql.com/4507>])
- Fixed that `mysqlbinlog --position --read-from-remote-server` had wrong # at lines. (Bug#4506 [<http://bugs.mysql.com/4506>])
- If `CREATE TEMPORARY TABLE t SELECT` failed while loading the data, the temporary table was not dropped. (Bug#4551 [<http://bugs.mysql.com/4551>])
- Fixed that when a multiple-table `DROP TABLE` failed to drop a table on the master server, the error code was not written to the binary log. (Bug#4553 [<http://bugs.mysql.com/4553>])
- When the slave SQL thread was replicating a `LOAD DATA INFILE` statement, it didn't show the statement in the output of `SHOW PROCESSLIST`. (Bug#4326 [<http://bugs.mysql.com/4326>])
- Fixed an assertion failure when reading the grant tables (Bug#4407 [<http://bugs.mysql.com/4407>])
- Fixed that `CREATE TABLE ... TYPE=HEAP ... AS SELECT...` caused replication slave to stop. (Bug#4971 [<http://bugs.mysql.com/4971>])
- Fixed that `mysql_options(...,MYSQL_OPT_LOCAL_INFILE,...)` failed to disable `LOAD DATA LOCAL INFILE`. (Bug#5038 [<http://bugs.mysql.com/5038>])
- Fixed that `disable-local-infile` option had no effect if client read it from a configuration file using `mysql_options(...,MYSQL_READ_DEFAULT,...)`. (Bug#5073 [<http://bugs.mysql.com/5073>])
- Fixed that `SET GLOBAL SYNC_BINLOG` did not work on some platforms (Mac OS X). (Bug#5064 [<http://bugs.mysql.com/5064>])
- Fixed that `mysql-test-run` failed on the `rpl_trunc_binlog` test if running test from the installed (the target of 'make install') directory. (Bug#5050 [<http://bugs.mysql.com/5050>])
- Fixed that `mysql-test-run` failed on the `grant_cache` test when run as Unix user 'root'. (Bug#4678 [<http://bugs.mysql.com/4678>])
- Fixed an unlikely deadlock which could happen when using `KILL`. (Bug#4810 [<http://bugs.mysql.com/4810>])
- Fixed a crash when one connection got `KILLED` while it was doing `START SLAVE`. (Bug#4827 [<http://bugs.mysql.com/4827>])
- Made `FLUSH TABLES WITH READ LOCK` block `COMMIT` if server is running with binary logging; this ensures that the binary log position is trustable when doing a full backup of tables and the binary log. (Bug#4953 [<http://bugs.mysql.com/4953>])
- Fixed that the counter of an `auto_increment` column was not reset by `TRUNCATE TABLE` if the table was a temporary table. (Bug#5033 [<http://bugs.mysql.com/5033>])
- Fixed bug which caused error to be reported when column from `ORDER BY` clause was present in two tables participating in `SELECT` even if the second instance of column in select list was renamed. (Bug#4302 [<http://bugs.mysql.com/4302>])

## D.1.20. Changes in release 4.1.3 (28 June 2004: Beta)

**Note:** The initial release of MySQL 4.1.3 for Windows accidentally was not compiled with support for the Spatial Extensions (OpenGIS). This was fixed by rebuilding from the same 4.1 code snapshot with the missing option and releasing those packages as version 4.1.3a.

To enable compiling the newly released PHP 5 against MySQL 4.1.3 on Windows, the Windows packages had to be rebuilt once more to add a few missing symbols to the MySQL client library. These packages were released as MySQL 4.1.3b.

Functionality added or changed:

- **Warning: Incompatible change:** C API change: `mysql_shutdown()` now requires a second argument. This is a source-level incompatibility that affects how you compile client programs; it does not affect the ability of compiled clients to communicate with older servers. See [Section 17.2.3.63](#), “`mysql_shutdown()`”.
- **Warning: Incompatible change:** The `timezone` system variable has been removed and replaced by `system_time_zone`. See [Section 5.2.2](#), “System Variables”.
- Support for per-connection time zones was added. Now you can set the current time zone for a connection by setting the `@@time_zone` system variable to a value such as `'+10:00'` or `'Europe/Moscow'` (where `'Europe/Moscow'` is the name of one of the time zones described in the system tables). Functions like `CURRENT_TIMESTAMP`, `UNIX_TIMESTAMP`, and so forth honor this time zone. Values of `TIMESTAMP` type are also interpreted as values in this time zone. So now our `TIMESTAMP` type behaves similar to Oracle's `TIMESTAMP WITH LOCAL TIME ZONE`. That is, values stored in such a column are normalized toward UTC and converted back to the current connection time zone when they are retrieved from such a column. To set up the tables that store time zone information, see [Section 2.10](#), “Post-Installation Setup and Testing”.
- Basic time zone conversion function `CONVERT_TZ()` was added. It assumes that its first argument is a datetime value in the time zone specified by its second argument and returns the equivalent datetime value in the time zone specified by its third argument.
- Added the `ARCHIVE` storage engine.
- Added SQL syntax for prepared statements. See [Section 13.7](#), “SQL Syntax for Prepared Statements”.
- Language-specific collations were added for the `ucs2` and `utf8` Unicode character sets: Icelandic, Latvian, Romanian, Slovenian, Polish, Estonian, Swedish, Turkish, Czech, Danish, Lithuanian, Slovak, Spanish, Traditional Spanish.
- `CHECK TABLE` now can be killed. It then marks the table as corrupted. See [Section 13.5.5.3](#), “KILL Syntax”.
- `OPTIMIZE TABLE` for InnoDB tables is now mapped to `ALTER TABLE` instead of `ANALYZE TABLE`.
- `sync_frm` is now a settable global variable (not only a startup option).
- Added the `sync_binlog=N` global variable and startup option, which makes the MySQL server synchronize its binary log to disk (`flush_sync_binlog()`) after every *N*th write to the binary log.
- Changed the slave SQL thread to print fewer useless error messages (no more message duplication; no more messages when an error is skipped (because of `slave-skip-errors`)).
- `DROP DATABASE IF EXISTS`, `DROP TABLE IF EXISTS`, single-table `DELETE` and single-

table `UPDATE` are now written to the binary log even if they changed nothing on the master (for example, even if the `DELETE` matched no row). The old behavior sometimes caused bad surprises in replication setups.

- Replication and `mysqlbinlog` now have better support for the case that the session character set and collation variables are changed within a given session. See [Section 6.7, “Replication Features and Known Problems”](#).
- Added `--innodb-safe-binlog` server option, which adds consistency guarantees between the content of `InnoDB` tables and the binary log. See [Section 5.11.4, “The Binary Log”](#).
- `LIKE` now supports the use of a prepared statement parameter or delimited constant expression as the argument to `ESCAPE` ([Bug#4200](#) [<http://bugs.mysql.com/4200>]).

#### Bugs fixed:

- Fixed `CREATE DATABASE IF NOT EXISTS` for Win32 which caused an error if database existed. ([Bug#4378](#) [<http://bugs.mysql.com/4378>])
- Added missing `root` account to Windows version of `mysqld`. ([Bug#4242](#) [<http://bugs.mysql.com/4242>])
- Fixed bug in prepared `EXPLAIN` statement which led to server crash. ([Bug#4271](#) [<http://bugs.mysql.com/4271>])
- Fixed a bug of using parameters in some prepared statements via SQL syntax. ([Bug#4280](#) [<http://bugs.mysql.com/4280>])
- Fixed a bug in `MERGE` tables created with `INSERT_METHOD=LAST`, that were not able to report a key number that caused “Duplicate entry” error for `UNIQUE` key in `INSERT`. As a result, error message was not precise enough (error 1022 instead of error 1062) and `INSERT ... ON DUPLICATE KEY UPDATE` did not work. ([Bug#4008](#) [<http://bugs.mysql.com/4008>])
- Fixed a bug in `DELETE` from a table with `FULLTEXT` indexes which under rare circumstances could result in a corrupted table, if words of different lengths may be considered equal (which is possible in some collations, for example, in `utf8_general_ci` or `latin1_german2_ci`). ([Bug#3808](#) [<http://bugs.mysql.com/3808>])
- Fixed too-early unlocking of tables if we have subquery in `HAVING` clause. ([Bug#3984](#) [<http://bugs.mysql.com/3984>])
- Fixed a bug in `mysqldump` when it didn't return an error if the output device was filled ([Bug#1851](#) [<http://bugs.mysql.com/1851>])
- Fixed a bug in client-side conversion of string column to `MYSQL_TIME` application buffer (prepared statements API). ([Bug#4030](#) [<http://bugs.mysql.com/4030>])
- Fixed a bug with server crash on attempt to execute a non-prepared statement. ([Bug#4236](#) [<http://bugs.mysql.com/4236>])
- Fixed a bug with server crash on attempt to prepare a statement with character set introducer. ([Bug#4105](#) [<http://bugs.mysql.com/4105>])
- Fixed bug which caused different number of warnings to be generated when bad datetime as string or as number was inserted into `DATETIME` or `TIMESTAMP` column. ([Bug#2336](#) [<http://bugs.mysql.com/2336>])

- Fixed some byte order bugs with prepared statements on machines with high-byte-first. (Bug#4173 [http://bugs.mysql.com/4173])
- Fixed unlikely bug in the range optimizer when using many `IN( )` queries on different key parts. (Bug#4157 [http://bugs.mysql.com/4157])
- Fixed problem with `NULL` and derived tables. (Bug#4097 [http://bugs.mysql.com/4097])
- Fixed wrong `UNION` results if display length of fields for numeric types was set less than real length of values in them. (Bug#4067 [http://bugs.mysql.com/4067])
- Fixed a bug in `mysql_stmt_close( )`, which hung up when attempting to close statement after failed `mysql_stmt_fetch( )`. (Bug#4079 [http://bugs.mysql.com/4079])
- Fixed bug of re-execution optimized `COUNT( * )`, `MAX( )` and `MIN( )` functions in prepared statements. (Bug#2687 [http://bugs.mysql.com/2687])
- Fixed a bug with `COUNT(DISTINCT)` performance degradation in cases like `COUNT(DISTINCT a TEXT, b CHAR(1))` (no index used). (Bug#3904 [http://bugs.mysql.com/3904])
- Fixed a bug in `MATCH ... AGAINST(... IN BOOLEAN MODE)` that under rare circumstances could cause wrong results if in the data's collation one byte could match many (like in `utf8_general_ci` or `latin1_german2_ci`). (Bug#3964 [http://bugs.mysql.com/3964])
- Fixed a bug in prepared statements protocol, when microseconds part of `MYSQL_TYPE_TIME/MYSQL_TYPE_DATETIME` columns was not sent to the client. (Bug#4026 [http://bugs.mysql.com/4026])
- Fixed a bug that using `--with-charset` with `configure` didn't affect the MySQL client library. (Bug#3990 [http://bugs.mysql.com/3990])
- Fixed a bug in authentication code that allowed a malicious user to bypass password verification with specially crafted packets, using a modified client library. (CVE-2004-0627 [http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2004-0627], CVE-2004-0628 [http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2004-0628])
- Fixed bug with wrong result of `CONCAT(?, col_name)` in prepared statements. (Bug#3796 [http://bugs.mysql.com/3796])
- Fixed potential memory overrun in `mysql_real_connect( )` (which required a compromised DNS server and certain operating systems). (Bug#4017 [http://bugs.mysql.com/4017], CVE-2004-0836 [http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2004-0836])
- During the installation process of the server RPM on Linux, `mysqld` was run as the `root` system user, and if you had `--log-bin=file_name`, where the file was located somewhere outside of the data directory, it created binary log files owned by `root` in this directory that remained owned by `root` after the installation. This is now fixed by starting `mysqld` as the `mysql` system user instead. (Bug#4038 [http://bugs.mysql.com/4038])
- Made `DROP DATABASE` honor the value of `lower_case_table_names`. (Bug#4066 [http://bugs.mysql.com/4066])
- The slave SQL thread refused to replicate `INSERT ... SELECT` if it examined more than 4 billion rows. (Bug#3871 [http://bugs.mysql.com/3871])
- `mysqlbinlog` didn't escape the string content of user variables, and did not deal well when these variables were in non-ASCII character sets; this is now fixed by always printing the string content of user variables in hexadecimal. The character set and collation of the string is now also printed. (Bug#3875 [http://bugs.mysql.com/3875])

- Fixed incorrect destruction of expression which led to crash of server on complex `AND/OR` expressions if query was ignored (either by a replication server because of `--replicate-*-table` rules, or by any MySQL server because of a syntax error). ([Bug#3969](http://bugs.mysql.com/3969) [http://bugs.mysql.com/3969], [Bug#4494](http://bugs.mysql.com/4494) [http://bugs.mysql.com/4494])

## D.1.21. Changes in release 4.1.2 (28 May 2004)

Functionality added or changed:

- `CHAR BYTE` is an alias for the `BINARY` data type. (Previously, it was an alias for `CHAR BINARY`.)
- **Warning: Incompatible change:** Handling of the `FLOAT` and `DOUBLE` floating-point data types is more strict to follow standard SQL. For example, a data type of `FLOAT(3,1)` stores a maximum value of 99.9. Previously, the server allowed larger numbers to be stored. That is, it stored a value such as 100.0 as 100.0. Now the server clips 100.0 to the maximum allowable value of 99.9. If you have tables that were created before MySQL 4.1.2 and that contain floating-point data not strictly legal for the column type, you should alter the data types of those columns. For example:

```
ALTER TABLE tbl_name MODIFY col_name FLOAT(4,1);
```

- **Warning: Incompatible change:** String comparison now works according to the SQL standard. Because we have that `'a' = 'a '` then from it must follow that `'a' > 'a\t'`. (The latter was not the case before MySQL 4.1.2.) To implement it, we had to change how storage engines compare strings internally. As a side effect, if you have a table where a `CHAR` or `VARCHAR` column in some row has a value with the last character less than `ASCII(32)`, you have to repair this table. `CHECK TABLES` tells you if this problem exists. ([Bug#3152](http://bugs.mysql.com/3152) [http://bugs.mysql.com/3152])
- Added support for `DEFAULT CURRENT_TIMESTAMP` and for `ON UPDATE CURRENT_TIMESTAMP` specifications for `TIMESTAMP` columns. Now you can explicitly say that a `TIMESTAMP` column should be set automatically to the current timestamp for `INSERT` and/or `UPDATE` statements, or even prevent the column from updating automatically. Only one column with such an auto-set feature per table is supported. `TIMESTAMP` columns created with earlier versions of MySQL behave as before. Behavior of `TIMESTAMP` columns that were created without explicit specification of default/on as earlier depends on its position in table: If it is the first `TIMESTAMP` column, it be treated as having been specified as `TIMESTAMP DEFAULT CURRENT_TIMESTAMP ON UPDATE CURRENT_TIMESTAMP`. In other cases, it would be treated as a `TIMESTAMP DEFAULT 0` column. `NOW` is supported as an alias for `CURRENT_TIMESTAMP`. **Warning: Incompatible change:** Unlike in previous versions, explicit specification of default values for `TIMESTAMP` column is never ignored and turns off the auto-set feature (unless you have `CURRENT_TIMESTAMP` as the default).
- **Warning: Incompatible change:** Renamed prepared statements C API functions:

Old Name	New Name
<code>mysql_bind_param()</code>	<code>mysql_stmt_bind_param()</code>
<code>mysql_bind_result()</code>	<code>mysql_stmt_bind_result()</code>
<code>mysql_prepare()</code>	<code>mysql_stmt_prepare()</code>
<code>mysql_execute()</code>	<code>mysql_stmt_execute()</code>
<code>mysql_fetch()</code>	<code>mysql_stmt_fetch()</code>
<code>mysql_fetch_column()</code>	<code>mysql_stmt_fetch_column()</code>
<code>mysql_param_count()</code>	<code>mysql_stmt_param_count()</code>
<code>mysql_param_result()</code>	<code>mysql_stmt_param_metadata()</code>

<code>mysql_get_metadata()</code>	<code>mysql_stmt_result_metadata()</code>
<code>mysql_send_long_data()</code>	<code>mysql_stmt_send_long_data()</code>

Now all functions that operate with a `MYSQL_STMT` structure begin with the prefix `mysql_stmt_`.

- **Warning: Incompatible change:** The signature of the `mysql_stmt_prepare()` function was changed to `int mysql_stmt_prepare(MYSQL_STMT *stmt, const char *query, unsigned long length)`. To create a `MYSQL_STMT` handle, you should use the `mysql_stmt_init()` function, not `mysql_stmt_prepare()`.
- **Warning: Incompatible change:** The `Type` output column for `SHOW TABLE STATUS` now is labeled `Engine`.
- Added the `EXAMPLE` storage engine.
- The `mysqld` Windows server was renamed to `mysqld-debug`. See [Section 2.3.8, “Selecting a MySQL Server type”](#).
- Added `Handler_discover` status variable.
- Added support for character set conversion and `MYSQL_TYPE_BLOB` type code in prepared statement protocol.
- Added explanation of hidden `SELECT` of `UNION` in output of `EXPLAIN SELECT` statement.
- `mysql` command-line client now supports multiple `-e` options. ([Bug#591](#) [<http://bugs.mysql.com/591>])
- New `myisam_data_pointer_size` system variable. See [Section 5.2.2, “System Variables”](#).
- The `--log-warnings` server option now is enabled by default. Disable with `--skip-log-warnings`.
- The `--defaults-file=file_name` option now requires that the filename must exist (safety fix). ([Bug#3413](#) [<http://bugs.mysql.com/3413>])
- `mysqld_multi` now creates the log in the directory named by `datadir` (from the `[mysqld]` section in `my.cnf` or compiled in), not in `/tmp`. Thanks to Christian Hammers from Debian Security Team for reporting this. ([CVE-2004-0388](#) [<http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2004-0388>])
- `SHOW GRANTS` with no `FOR` clause or with `FOR CURRENT_USER()` shows the privileges for the current session.
- The improved character set support introduced in MySQL 4.1.0 for the `MyISAM` and `HEAP` storage engines is now available for `InnoDB` as well.
- A name of “Primary” no longer can be specified as an index name. (That name is reserved for the `PRIMARY KEY` if the table has one.) ([Bug#856](#) [<http://bugs.mysql.com/856>])
- MySQL now issues a warning when a `SET` or `ENUM` column with duplicate values in the list is created. ([Bug#1427](#) [<http://bugs.mysql.com/1427>])
- Now `SQL_SELECT_LIMIT` variable has no influence on subqueries. ([Bug#2600](#))

[<http://bugs.mysql.com/2600>)]

- `UNHEX ( )` function implemented. See [Section 12.3, “String Functions”](#).
- The `mysql` command-line client no longer stores in the history file multiple copies of identical queries that are run consecutively.
- Multi-line statements in the `mysql` command-line client now are stored in the history file as a single line.
- `UUID ( )` function implemented. Note that it does not work with replication yet. See [Section 12.9.4, “Miscellaneous Functions”](#).
- Prepared statements with all types of subqueries fixed.
- MySQL now supports up to 64 indexes per table.
- `MyISAM` tables now support keys up to 1000 bytes long.
- `MyISAM` and `InnoDB` tables now support index prefix lengths up to 1000 bytes long.
- If you try to create a key with a key part that is too long, and it is safe to auto-truncate it to a smaller length, MySQL now does so. A warning is generated, rather than an error.
- The `ft_boolean_syntax` variable now can be changed while the server is running. See [Section 5.2.2, “System Variables”](#).
- `REVOKE ALL PRIVILEGES, GRANT FROM user_list` is changed to a more consistent `REVOKE ALL PRIVILEGES, GRANT OPTION FROM user_list`. ([Bug#2642](#) [<http://bugs.mysql.com/2642>])
- Internal string-to-number conversion now supports only SQL:2003 compatible syntax for numbers. In particular, `'0x10'+0` does not work anymore. (Actually, it worked only on some systems before, such as Linux. It did not work on others, such as FreeBSD or Solaris. Making these queries OS-independent was the goal of this change.) Use `CONV ( )` to convert hexadecimal numbers to decimal. Example: `CONV(MID( '0x10' , 3 ) , 16 , 10)+0`.
- `mysqlhotcopy` now works on NetWare.
- `ALTER TABLE DROP PRIMARY KEY` no longer drops the first `UNIQUE` index if there is no primary index. ([Bug#2361](#) [<http://bugs.mysql.com/2361>])
- Added `latin1_spanish_ci` (Modern Spanish) collation for the `latin1` character set.
- Added the `ENGINE` table option as a synonym for the `TYPE` option for `CREATE TABLE` and `ALTER TABLE`.
- Added the `--default-storage-engine` server option as a synonym for `--default-table-type`.
- Added the `storage_engine` system variable as a synonym for `table_type`.
- Added `init_connect` and `init_slave` system variables. The values should be SQL statements to be executed when each client connects or each time a slave's SQL thread starts, respectively.
- C API enhancement: `SERVER_QUERY_NO_INDEX_USED` and `SERV-ER_QUERY_NO_GOOD_INDEX_USED` flags are now set in the `server_status` field of the `MYSQL` structure. It is these flags that make the query to be logged as slow if `mysqld` was started



with `--log-slow-queries --log-queries-not-using-indexes`.

- For replication of `MEMORY` (`HEAP`) tables: Made the master automatically write a `DELETE FROM` statement to its binary log when a `MEMORY` table is opened for the first time since master's startup. This is for the case where the slave has replicated a non-empty `MEMORY` table, then the master is shut down and restarted: the table is now empty on master; the `DELETE FROM` empties it on slave too. Note that even with this fix, between the master's restart and the first use of the table on master, the slave still has out-of-date data in the table. But if you use the `init-file` option to populate the `MEMORY` table on the master at startup, it ensures that the failing time interval is zero. (Bug#2477 [<http://bugs.mysql.com/2477>])
- When a session having open temporary tables terminates, the statement automatically written to the binary log is now `DROP TEMPORARY TABLE IF EXISTS` instead of `DROP TEMPORARY TABLE`, for more robustness.
- The MySQL server now returns an error if `SET SQL_LOG_BIN` or `SET SQL_LOG_UPDATE` is issued by a user without the `SUPER` privilege (in previous versions it just silently ignored the statement in this case).
- Changed that when the MySQL server has binary logging disabled (that is, no `--log-bin` option was used), then no transaction binary log cache is allocated for connections. This should save `bin-log-cache-size` bytes of memory (32KB by default) for every connection.
- Added `Binlog_cache_use` and `Binlog_cache_disk_use` status variables that count the number of transactions that used transaction binary log and that had to flush this temporary binary log to disk instead of using only the in-memory buffer. They can be used for tuning the `bin-log-cache-size` system variable.
- Added option `--replicate-same-server-id`.
- The Mac OS X Startup Item has been moved from the directory `/Library/StartupItems/MySQL` to `/Library/StartupItems/MySQLCOM` to avoid a file name collision with the MySQL Startup Item installed with Mac OS X Server. See [Section 2.13.2](#), "Mac OS X Notes".
- Added option `--to-last-log` to `mysqlbinlog`, for use in conjunction with `--read-from-remote-server`.
- The `FLOAT` and `DECIMAL` types now obey (precision,scale) settings. (Bug#10897 [<http://bugs.mysql.com/10897>])
- Added the `mysql_set_local_infile_handler()` and `mysql_set_local_infile_default()` C API functions.

Bugs fixed:

- Fixed check of `EXPLAIN` of `UNION`. (Bug#3639 [<http://bugs.mysql.com/3639>])
- Fixed a bug in a query that used `DISTINCT` and `ORDER BY` by column's real name, while the column had an alias, specified in `SELECT` clause. (Bug#3681 [<http://bugs.mysql.com/3681>])
- `mysqld` could crash when a table was altered and used at the same time. This was a 4.1.2-specific bug. (Bug#3643 [<http://bugs.mysql.com/3643>]).
- Fixed bug when using impossible `WHERE` with `PROCEDURE ANALYSE()`. (Bug#2238 [<http://bugs.mysql.com/2238>]).

- Fixed security problem in new authentication where password was not checked for changed `GRANT` accounts until `FLUSH PRIVILEGES` was executed. (Bug#3404 [http://bugs.mysql.com/3404])
- Fixed crash of `GROUP_CONCAT()` on expression with `ORDER BY` and external `ORDER BY` in a query. (Bug#3752 [http://bugs.mysql.com/3752])
- Fixed a bug in `ALL/SOME` subqueries in case of optimization (key field present in subquery). (Bug#3646 [http://bugs.mysql.com/3646])
- Fixed a bug in `SHOW GRANTS` and `EXPLAIN SELECT` character set conversion. (Bug#3403 [http://bugs.mysql.com/3403])
- Prepare statements parameter do not cause error message as fields used in select list but not included in `ORDER BY` list.
- `UNION` statements did not consult `SQL_SELECT_LIMIT` value when set. This is now fixed properly, which means that this limit is applied to the top level query, unless `LIMIT` for entire `UNION` is used.
- Fixed a bug in multiple-table `UPDATE` statements that resulted in an error when one of the tables was not updated but was used in the nested query, contained therein.
- Fixed `mysql_stmt_send_long_data()` behavior on second execution of prepared statement and in case when long data had zero length. (Bug#1664 [http://bugs.mysql.com/1664])
- Fixed crash on second execution of prepared statement with `UNION`. (Bug#3577 [http://bugs.mysql.com/3577])
- Fixed incorrect results of aggregate functions in subquery with empty result set. (Bug#3505 [http://bugs.mysql.com/3505])
- You can now call `mysql_stmt_attr_set(..., STMT_ATTR_UPDATE_MAX_LENGTH)` to tell the client library to update `MYSQL_FIELD->max_length` when doing `mysql_stmt_store_result()`. (Bug#1647 [http://bugs.mysql.com/1647]).
- Added support for unsigned integer types to prepared statement API (Bug#3035 [http://bugs.mysql.com/3035]).
- Fixed crash in prepared statements when subquery in the `FROM` clause with parameter used. (Bug#3020 [http://bugs.mysql.com/3020])
- Fixed unknown error when negative value bind to unsigned. (Bug#3223 [http://bugs.mysql.com/3223])
- Fixed aggregate function in prepared statements. (Bug#3360 [http://bugs.mysql.com/3360])
- Incorrect error message when wrong table used in multiple-table `DELETE` statement in prepared statements. (Bug#3411 [http://bugs.mysql.com/3411])
- Requiring `UPDATE` privilege for tables which are not updated in multiple-table `UPDATE` statement in prepared statements.
- Fixed prepared statement support for `INSERT`, `REPLACE`, `CREATE`, `DELETE`, `SELECT`, `DO`, `SET` and `SHOW`. All other commands are prohibited via prepared statement interface. (Bug#3398 [http://bugs.mysql.com/3398], Bug#3406 [http://bugs.mysql.com/3406], Bug#2811 [http://bugs.mysql.com/2811])
- Fixed a lot of bugs in `GROUP_CONCAT()`. (Bug#2695 [http://bugs.mysql.com/2695], Bug#3381 [http://bugs.mysql.com/3381], Bug#3319 [http://bugs.mysql.com/3319])

- Added optimization that allows for prepared statements using a large number of tables or tables with a large number of columns to be re-executed significantly faster. (Bug#2050 [<http://bugs.mysql.com/2050>])
- Fixed bug that caused execution of prepared statements to fail then table that this statement were using left table cache. This bug showed up as if this prepared statement used random garbage as column names or as server crashes. (Bug#3307 [<http://bugs.mysql.com/3307>])
- Fixed a problem resulting from setting the `character_set_results` variable to `NULL`. (Bug#3296 [<http://bugs.mysql.com/3296>])
- Fixed query cache statistics.
- Fixed bug in `ANALYZE TABLE` on a `BDB` table inside a transaction that hangs server thread. (Bug#2342 [<http://bugs.mysql.com/2342>])
- Fixed a symlink vulnerability in the `mysqlbug` script. (Bug#3284 [<http://bugs.mysql.com/3284>], CVE-2004-0381 [<http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2004-0381>])
- Fixed a bug in parallel repair (`myisamchk -p, myisam_repair_threads`); sometimes the repair process failed to repair a table. (Bug#1334 [<http://bugs.mysql.com/1334>])
- A query that uses both `UNION [DISTINCT]` and `UNION ALL` now works correctly. (Bug#1428 [<http://bugs.mysql.com/1428>])
- Table default character set affects `LONGBLOB` columns. (Bug#2821 [<http://bugs.mysql.com/2821>])
- `CONCAT_WS ( )` makes the server die in case of illegal mix of collations. (Bug#3087 [<http://bugs.mysql.com/3087>])
- UTF8 charset breaks joins with mixed column/string constant. (Bug#2959 [<http://bugs.mysql.com/2959>])
- Fixed `DROP DATABASE` to report number of tables deleted.
- Fixed bug in privilege checking of `ALTER TABLE RENAME`. (Bug#3270 [<http://bugs.mysql.com/3270>], CVE-2004-0835 [<http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2004-0835>])
- Fixed memory leak in the client library when statement handle was freed on closed connection (call to `mysql_stmt_close` after `mysql_close`). (Bug#3073 [<http://bugs.mysql.com/3073>])
- Fixed server segmentation faults when processing malformed prepared statements. (Bug#2795 [<http://bugs.mysql.com/2795>], Bug#2274 [<http://bugs.mysql.com/2274>])
- Fixed using subqueries with `OR` and `AND` functions. (Bug#2838 [<http://bugs.mysql.com/2838>])
- Fixed comparison of tables/database names with `--lower_case_table_names` option. (Bug#2880 [<http://bugs.mysql.com/2880>])
- Removed try to check `NULL` if index built on column where `NULL` is impossible in `IN` subquery optimization. (Bug#2393 [<http://bugs.mysql.com/2393>])
- Fixed incorrect parsing of subqueries in the `FROM` clause. (Bug#2421 [<http://bugs.mysql.com/2421>])
- Fixed processing of `RAND ( )` in subqueries with static tables. (Bug#2645 [<http://bugs.mysql.com/2645>])
- Fixed bug with quoting of table names in `mysqldump` for various values of `sql_mode` of server.

[Bug#2591](http://bugs.mysql.com/2591) [http://bugs.mysql.com/2591])

- Fixed bug with storing values that are out of range for `DOUBLE` and `FLOAT` columns. ([Bug#2082](http://bugs.mysql.com/2082) [http://bugs.mysql.com/2082])
- Fixed bug with compiling `--with-pstack` with binutils 2.13.90. ([Bug#1661](http://bugs.mysql.com/1661) [http://bugs.mysql.com/1661])
- Fixed a bug in the `GRANT` system. When a password was assigned to an account at the global level and then privileges were granted at the database level (without specifying any password), the existing password was replaced temporarily in memory until the next `FLUSH PRIVILEGES` operation or the server was restarted. ([Bug#2953](http://bugs.mysql.com/2953) [http://bugs.mysql.com/2953])
- Fixed a bug in full-text search on multi-byte character set (such as UTF8) that appeared when a search word was shorter than a matching word from the index (for example, searching for “Uppsala” when table data contain “Uppsa\*la”). ([Bug#3011](http://bugs.mysql.com/3011) [http://bugs.mysql.com/3011])
- Fixed a bug that made `Max_used_connections` to be less than the actual maximum number of connections in use simultaneously.
- Fixed calculation of `Index_length` in `HEAP` table status for `BTREE` indexes. ([Bug#2719](http://bugs.mysql.com/2719) [http://bugs.mysql.com/2719])
- Fixed `mysql_stmt_affected_rows()` call to always return number of rows affected by given statement. ([Bug#2247](http://bugs.mysql.com/2247) [http://bugs.mysql.com/2247])
- Fixed crash in `MATCH ... AGAINST()` on a phrase search operator with a missing closing double quote. ([Bug#2708](http://bugs.mysql.com/2708) [http://bugs.mysql.com/2708])
- Fixed output of `mysqldump --tab`. ([Bug#2705](http://bugs.mysql.com/2705) [http://bugs.mysql.com/2705])
- Fix for a bug in `UNION` operations that prevented proper handling of `NULL` columns. This happened only if a column in the first `SELECT` node was `NOT NULL`. ([Bug#2508](http://bugs.mysql.com/2508) [http://bugs.mysql.com/2508])
- Fix for a bug in `UNION` operations with `InnoDB` storage engine, when some columns from one table were used in one `SELECT` statement and some were used in another `SELECT` statement. ([Bug#2552](http://bugs.mysql.com/2552) [http://bugs.mysql.com/2552])
- Fixed a few years old bug in the range optimizer that caused a segmentation fault on some very rare queries. ([Bug#2698](http://bugs.mysql.com/2698) [http://bugs.mysql.com/2698])
- Fixed bug with `SHOW CREATE TABLE ...` which didn't properly double quotes. ([Bug#2593](http://bugs.mysql.com/2593) [http://bugs.mysql.com/2593])
- Queries with subqueries in `FROM` clause locks all tables at once for now. This also fixed bugs in `EXPLAIN` of subqueries in `FROM` output. ([Bug#2120](http://bugs.mysql.com/2120) [http://bugs.mysql.com/2120])
- Fixed bug with `mysqldump` not quoting “tricky” names correctly. ([Bug#2592](http://bugs.mysql.com/2592) [http://bugs.mysql.com/2592])
- Fix for a bug that prevented table / column privileges from being loaded on startup. ([Bug#2546](http://bugs.mysql.com/2546) [http://bugs.mysql.com/2546])
- Fixed bug in replication with `CREATE TABLE ... LIKE ...` that resulted in a statement not being written to the binary log. ([Bug#2557](http://bugs.mysql.com/2557) [http://bugs.mysql.com/2557])
- Fixed memory leak in `INSERT ... ON DUPLICATE KEY UPDATE ...`. ([Bug#2438](http://bugs.mysql.com/2438) [http://bugs.mysql.com/2438])

- Fixed bug in the parser, making the syntax `CONVERT(expr, type)` legal again.
- Fixed parsing of short-form IP addresses in `INET_ATON()`. (Bug#2310 [<http://bugs.mysql.com/2310>])
- Fixed a bug in `CREATE ... SELECT` that sometimes caused a string column with a multi-byte character set (such as `utf8`) to have insufficient length to hold the data.
- Fixed a rare table corruption on adding data (`INSERT`, `REPLACE`, `UPDATE`, etc. but not `DELETE`) to a `FULLTEXT` index. (Bug#2417 [<http://bugs.mysql.com/2417>])
- Compile the `MySQL-client` RPM package against `libreadline` instead of `libedit`. (Bug#2289 [<http://bugs.mysql.com/2289>])
- Fix for a crashing bug that was caused by not setting `vio_timeout()` virtual function for all protocols. This bug occurred on Windows. (Bug#2025 [<http://bugs.mysql.com/2025>])
- Fix for a bug that caused `mysql` client program to erroneously cache the value of the current database. (Bug#2025 [<http://bugs.mysql.com/2025>])
- Fix for a bug that caused client/server communication to be broken when `mysql_set_server_option()` or `mysql_get_server_option()` were invoked. (Bug#2207 [<http://bugs.mysql.com/2207>])
- Fix for a bug that caused wrong results when `CAST()` was applied on `NULL` to signed or unsigned integer column. (Bug#2219 [<http://bugs.mysql.com/2219>])
- Fix for a crashing bug that occurred in the `mysql` client program when database name was longer than expected. (Bug#2221 [<http://bugs.mysql.com/2221>])
- Fixed a bug in `CHECK TABLE` that sometimes resulted in a spurious error `Found key at page ... that points to record outside datafile` for a table with a `FULLTEXT` index. (Bug#2190 [<http://bugs.mysql.com/2190>])
- Fixed bug in `GRANT` with table-level privilege handling. (Bug#2178 [<http://bugs.mysql.com/2178>])
- Fixed bug in `ORDER BY` on a small column. (Bug#2147 [<http://bugs.mysql.com/2147>])
- Fixed a bug with the `INTERVAL()` function when 8 or more comparison arguments are provided. (Bug#1561 [<http://bugs.mysql.com/1561>])
- Packaging: Fixed a bug in the Mac OS PKG `postinstall` script (`mysql_install_db` was called with an obsolete argument).
- Packaging: Added missing file `mysql_create_system_tables` to the server RPM package. This bug was fixed for the 4.1.1 RPMs by updating the MySQL-server RPM from `MySQL-server-4.1.1-0` to `MySQL-server-4.1.1-1`. The other RPMs were not affected by this change.
- Fixed a bug in `myisamchk` and `CHECK TABLE` that sometimes resulted in a spurious error `Found key at page ... that points to record outside datafile` for a table with a `FULLTEXT` index. (Bug#1977 [<http://bugs.mysql.com/1977>])
- Fixed a hang in full-text indexing of strings in multi-byte (all besides `utf8`) charsets. (Bug#2065 [<http://bugs.mysql.com/2065>])
- Fixed a crash in full-text indexing of UTF8 data. (Bug#2033 [<http://bugs.mysql.com/2033>])
- Replication: a rare race condition in the slave SQL thread that could lead to an incorrect complaint that the relay log is corrupted. (Bug#2011 [<http://bugs.mysql.com/2011>])

- Replication: If a client connects to a slave server and issues an administrative statement for a table (for example, `OPTIMIZE TABLE` or `REPAIR TABLE`), this could sometimes stop the slave SQL thread. This does not lead to any corruption, but you must use `START SLAVE` to get replication going again. (Bug#1858 [<http://bugs.mysql.com/1858>])
- Replication: in the slave SQL thread, a multiple-table `UPDATE` could produce an incorrect complaint that some record was not found in one table, if the `UPDATE` was preceded by a `INSERT . . . SELECT`. (Bug#1701 [<http://bugs.mysql.com/1701>])
- Replication: sometimes the master gets a non-fatal error during the execution of a statement but finally the statements succeeds (for example, a write to a `MyISAM` table first receives "no space left on device" but is able to finally complete, see Section A.4.3, "How MySQL Handles a Full Disk"); the bug was that the master forgot to reset the error code to 0 after success, so the error code got into its binary log, thus making the slave giving false alarms like "did not get the same error as on master". (Bug#2083 [<http://bugs.mysql.com/2083>])
- Removed a misleading "check permissions on master.info" from a replication error message, because the cause of the problem could be different from permissions. (Bug#2121 [<http://bugs.mysql.com/2121>])
- Fixed a crash when the replication slave was unable to create the first relay log. (Bug#2145 [<http://bugs.mysql.com/2145>])
- `ALTER DATABASE` caused the client to hang if the database did not exist. (Bug#2333 [<http://bugs.mysql.com/2333>])
- Multiple-table `DELETE` statements were never replicated by the slave if there were any `--replicate-*-table` options. (Bug#2527 [<http://bugs.mysql.com/2527>])
- Fixed bug in `ALTER TABLE RENAME`, when rename to the table with the same name in another database silently dropped destination table if it existed. (Bug#2628 [<http://bugs.mysql.com/2628>])
- The MySQL server did not report any error if a statement (submitted through `mysql_real_query()` or `mysql_stmt_prepare()`) was terminated by garbage characters. This can happen if you pass a wrong `length` parameter to these functions. The result was that the garbage characters were written into the binary log. (Bug#2703 [<http://bugs.mysql.com/2703>])
- Fixed bug in client library that caused `mysql_stmt_fetch` and `mysql_stmt_store_result()` to hang if they were called without prior call of `mysql_stmt_execute()`. Now they give an error instead. (Bug#2248 [<http://bugs.mysql.com/2248>])
- Made clearer the error message that one gets when an update is refused because of the `--read-only` option. (Bug#2757 [<http://bugs.mysql.com/2757>])
- Fixed that `--replicate-wild-*-table` rules apply to `ALTER DATABASE` when the table pattern is `%`, as is the case for `CREATE DATABASE` and `DROP DATABASE`. (Bug#3000 [<http://bugs.mysql.com/3000>])
- Fixed that when a `Rotate` event is found by the slave SQL thread in the middle of a transaction, the value of `Relay_Log_Pos` in `SHOW SLAVE STATUS` remains correct. (Bug#3017 [<http://bugs.mysql.com/3017>])
- Corrected the master's binary log position that `InnoDB` reports when it is doing a crash recovery on a slave server. (Bug#3015 [<http://bugs.mysql.com/3015>])
- Changed the column `Seconds_Behind_Master` in `SHOW SLAVE STATUS` to never show a value of -1. (Bug#2826 [<http://bugs.mysql.com/2826>])

- Changed that when a `DROP TEMPORARY TABLE` statement is automatically written to the binary log when a session ends, the statement is recorded with an error code of value zero (this ensures that killing a `SELECT` on the master does not result in a superfluous error on the slave). (Bug#3063 [<http://bugs.mysql.com/3063>])
- Changed that when a thread handling `INSERT DELAYED` (also known as a `delayed_insert` thread) is killed, its statements are recorded with an error code of value zero (killing such a thread does not endanger replication, so we thus avoid a superfluous error on the slave). (Bug#3081 [<http://bugs.mysql.com/3081>])
- Fixed deadlock when two `START SLAVE` commands were run at the same time. (Bug#2921 [<http://bugs.mysql.com/2921>])
- Fixed that a statement never triggers a superfluous error on the slave, if it must be excluded given the `--replicate-*` options. The bug was that if the statement had been killed on the master, the slave would stop. (Bug#2983 [<http://bugs.mysql.com/2983>])
- The `--local-load` option of `mysqlbinlog` now requires an argument.
- Fixed a segmentation fault when running `LOAD DATA FROM MASTER` after `RESET SLAVE`. (Bug#2922 [<http://bugs.mysql.com/2922>])
- `mysqlbinlog --read-from-remote-server` read all binary logs following the one that was requested. It now stops at the end of the requested file, the same as it does when reading a local binary log. There is an option `--to-last-log` to get the old behavior. (Bug#3204 [<http://bugs.mysql.com/3204>])
- Fixed `mysqlbinlog --read-from-remote-server` to print the exact positions of events in the "at #" lines. (Bug#3214 [<http://bugs.mysql.com/3214>])
- Fixed a rare error condition that caused the slave SQL thread spuriously to print the message `Bin-log has bad magic number` and stop when it was not necessary to do so. (Bug#3401 [<http://bugs.mysql.com/3401>])
- Fixed the `Exec_master_log_pos` column and its disk image in the `relay-log.info` file to be correct if the master had version 3.23. (The value was too big by six bytes.) This bug does not exist in MySQL 5.0. (Bug#3400 [<http://bugs.mysql.com/3400>])
- Fixed `mysqlbinlog` not to forget to print a `USE` statement under rare circumstances where the binary log contained a `LOAD DATA INFILE` statement. (Bug#3415 [<http://bugs.mysql.com/3415>])
- Fixed a memory corruption when replicating a `LOAD DATA INFILE` when the master had version 3.23. Some smaller problems remain in this setup, See [Section 6.7, "Replication Features and Known Problems"](#). (Bug#3422 [<http://bugs.mysql.com/3422>])
- Multiple-table `DELETE` statements were always replicated by the slave if there were some `--replicate-*--ignore-table` options and no `--replicate-*--do-table` options. (Bug#3461 [<http://bugs.mysql.com/3461>])
- Fixed a crash of the MySQL slave server when it was built with `--with-debug` and replicating itself. (Bug#3568 [<http://bugs.mysql.com/3568>])
- Fixed that in some replication error messages, a very long query caused the rest of the message to be invisible (truncated), by putting the query last in the message. (Bug#3357 [<http://bugs.mysql.com/3357>])
- Fixed a bug in `REPAIR TABLE` that resulted sometimes in a corrupted table, if the table contained `FULLTEXT` indexes and many words of different lengths that are considered equal (which is possible in certain collations, such as `latin1_german2_ci` or `utf8_general_ci`). (Bug#3835

[<http://bugs.mysql.com/3835>])

- Fixed a crash of `mysqld` that was started with binary logging disabled, but with a non-zero value for the `expire_logs_days` system variable. (Bug#3807 [<http://bugs.mysql.com/3807>])
- If `server-id` was not set using startup options but with `SET GLOBAL`, the replication slave still complained that it was not set. (Bug#3829 [<http://bugs.mysql.com/3829>])

## D.1.22. Changes in release 4.1.1 (01 December 2003)

This release includes all fixes in MySQL 4.0.16 and most of the fixes in MySQL 4.0.17.

Functionality added or changed:

- **Warning: Incompatible change:** Renamed the C API `mysql_prepare_result()` function to `mysql_get_metadata()` because the old name was confusing.
- **Warning: Incompatible change:** Client authentication now is based on 41-byte passwords in the `user` table, not 45-byte passwords as in 4.1.0. Any 45-byte passwords created for 4.1.0 must be re-set after running the `mysql_fix_privilege_tables` script.
- **Important note:** If you upgrade to `InnoDB-4.1.1` or higher, you cannot downgrade to a version lower than 4.1.1 any more! That is because earlier versions of `InnoDB` are not aware of multiple tablespaces.
- Added `secure_auth` global server system variable and `--secure-auth` server option that disallow authentication for accounts that have old (pre-4.1.1) passwords.
- Added `--secure-auth` option to `mysql` command-line client. If this option is set, the client refuses to send passwords in old (pre-4.1.1) format.
- `EXPLAIN` now supports an `EXTENDED` option. When given, `EXPLAIN` generates extra information that may be viewed with the `SHOW WARNINGS` statement.
- Table aliases are not case sensitive if `lower_case_table_names` is non-zero.
- The `--old-protocol` option for `mysqld` is no longer supported and has been removed.
- Renamed `bdb_version` system variable to `version_bdb`.
- `MyISAM` tables now use a better checksum algorithm (if checksum is enabled with `CREATE TABLE ... CHECKSUM=1`). Old tables will appear to have incorrect checksum, and should be repaired.
- New `CHECKSUM TABLE` statement for reporting table checksum values.
- Added `character_set_client`, `character_set_connection`, `character_set_database`, `character_set_results`, `character_set_server`, `character_set_system`, `collation_connection`, `collation_database`, and `collation_server` system variables to provide information about character sets and collations.
- It is now possible to create multiple key caches, assign table indexes to particular caches, and to preload indexes into caches. See [Section 13.5.5.1, “CACHE INDEX Syntax”](#). See [Section 13.5.5.4, “LOAD INDEX INTO CACHE Syntax”](#). Structured system variables are introduced as a means of grouping related key cache parameters. See [Section 5.2.3.1, “Structured System Variables”](#).
- Added `preload_buffer_size` system variable.



- New `COERCIBILITY()` function to return the collation coercibility of a string.
- The `--quote-names` option for `mysqldump` now is enabled by default.
- `mysqldump` now includes a statement in the dump output to set `FOREIGN_KEY_CHECKS` to 0 to avoid problems with tables having to be reloaded in a particular order when the dump is reloaded. The existing `FOREIGN_KEY_CHECKS` value is saved and restored.
- You can revoke all privileges from a user with `REVOKE ALL PRIVILEGES, GRANT FROM user_list`.
- Added `IGNORE` option for `DELETE` statement.
- MySQL source distributions now also include the MySQL Internals Manual `internals.texi`.
- Added `mysql_set_server_option()` C API client function to allow multiple statement handling in the server to be enabled or disabled.
- The `mysql_next_result()` C API function now returns `-1` if there are no more result sets.
- Renamed `CLIENT_MULTI_QUERIES` connect option flag to `CLIENT_MULTI_STATEMENTS`. To allow for a transition period, the old option continues to be recognized for a while.
- Require `DEFAULT` before table and database default character set. This enables us to use `ALTER TABLE tbl_name ... CHARACTER SET=...` to change the character set for all `CHAR`, `VARCHAR`, and `TEXT` columns in a table.
- Added `MATCH ... AGAINST( ... WITH QUERY EXPANSION)` and the `ft_query_expansion_limit` system variable.
- Removed unused `ft_max_word_len_for_sort` system variable.
- Removed unused `ft_max_word_len_for_sort` variable from `myisamchk`.
- Full-text search now supports multi-byte character sets and the Unicode `utf8` character set. (The Unicode `ucs2` character set is not yet supported.)
- Phrase search in `MATCH ... AGAINST ( ... IN BOOLEAN MODE)` no longer matches partial words.
- Added aggregate function `BIT_XOR()` for bitwise XOR operations.
- Replication over SSL now works.
- The `START SLAVE` statement now supports an `UNTIL` clause for specifying that the slave SQL thread should be started but run only until it reaches a given position in the master's binary logs or in the slave's relay logs.
- Produce warnings even for single-row `INSERT` statements, not just for multiple-row `INSERT` statements. Previously, it was necessary to set `SQL_WARNINGS=1` to generate warnings for single-row statements.
- Added `delimiter (\d)` command to the `mysql` command-line client for changing the statement delimiter (terminator). The default delimiter is semicolon.
- `CHAR`, `VARCHAR`, and `TEXT` columns now have lengths measured in characters rather than in bytes. The character size depends on the column's character set. This means, for example, that a `CHAR(n)` column for a multi-byte character set takes more storage than before. Similarly, index values on such columns are measured in characters, not bytes.

- `LIMIT` no longer accepts negative arguments (they used to be treated as very big positive numbers before).
- The `DATABASE()` function now returns `NULL` rather than the empty string if there is no database selected.
- Added `--sql-mode=NO_AUTO_VALUE_ON_ZERO` option to suppress the usual behavior of generating the next sequence number when zero is stored in an `AUTO_INCREMENT` column. With this mode enabled, zero is stored as zero; only storing `NULL` generates a sequence number.
- Added `DROP USER 'user_name'@'host_name'` statement to drop an account that has no privileges.
- The interface to aggregate user-defined functions has changed a bit. You must now declare a `xxx_clear()` function for each aggregate function `XXX()`. `xxx_clear()` is used instead of `xxx_reset()`.
- Added new `ADDTIME()`, `DATE()`, `DATEDIFF()`, `LAST_DAY()`, `MAKEDATE()`, `MAKETIME()`, `MICROSECOND()`, `SUBTIME()`, `TIME()`, `TIMEDIFF()`, `TIMESTAMP()`, `UTC_DATE()`, `UTC_TIME()`, `UTC_TIMESTAMP()`, and `WEEKOFYEAR()` functions.
- Added new syntax for `ADDDATE()` and `SUBDATE()`. The second argument now may be a number representing the number of days to be added to or subtracted from the first date argument.
- Added new `type` values `DAY_MICROSECOND`, `HOURL_MICROSECOND`, `MINUTE_MICROSECOND`, `SECOND_MICROSECOND`, and `MICROSECOND` for `DATE_ADD()`, `DATE_SUB()`, and `EXTRACT()`.
- Added new `%f` microseconds format specifier for `DATE_FORMAT()` and `TIME_FORMAT()`.
- All queries in which at least one `SELECT` does not use indexes properly now are written to the slow query log when long log format is used.
- It is now possible to create a `MERGE` table from `MyISAM` tables in different databases. Formerly, all the `MyISAM` tables had to be in the same database, and the `MERGE` table had to be created in that database as well.
- Added new `COMPRESS()`, `UNCOMPRESS()`, and `UNCOMPRESSED_LENGTH()` functions.
- When using `SET sql_mode='mode'` for a complex mode (such as `ANSI`), we now update the `sql_mode` variable to include all the individual options implied by the complex mode.
- Added the OLAP (On-Line Analytical Processing) function `ROLLUP`, which provides summary rows for each `GROUP BY` level.
- Added `SQLSTATE` codes for all server errors.
- Added `mysql_sqlstate()` and `mysql_stmt_sqlstate()` C API client functions that return the `SQLSTATE` error code for the last error.
- `TIME` columns with hour values greater than 24 were returned incorrectly to the client.
- `ANALYZE TABLE`, `OPTIMIZE TABLE`, `REPAIR TABLE`, and `FLUSH` statements are now stored in the binary log and thus replicated to slaves. This logging does not occur if the optional `NO_WRITE_TO_BINLOG` keyword (or its alias `LOCAL`) is given. Exceptions are that `FLUSH LOGS`, `FLUSH MASTER`, `FLUSH SLAVE`, and `FLUSH TABLES WITH READ LOCK` are not logged in any case. For a syntax example, see [Section 13.5.5.2, “FLUSH Syntax”](#).
- New global system variable `relay_log_purge` to enable or disable automatic relay log purging.

- `LOAD DATA` now produces warnings that can be fetched with `SHOW WARNINGS`.
- Added support for syntax `CREATE TABLE table2 (LIKE table1)` that creates an empty table `table2` with a definition that is exactly the same as `table1`, including any indexes.
- `CREATE TABLE tbl_name (...) TYPE=storage_engine` now generates a warning if the named storage engine is not available. The table is still created as a `MyISAM` table, as before.
- Most subqueries are now much faster than before.
- Added `PURGE BINARY LOGS` as an alias for `PURGE MASTER LOGS`.
- Disabled the `PURGE LOGS` statement that was added in version 4.1.0. The statement now should be issued as `PURGE MASTER LOGS` or `PURGE BINARY LOGS`.
- Added `SHOW BDB LOGS` as an alias for `SHOW LOGS`.
- Added `SHOW MASTER LOGS` as an alias for `SHOW BINARY LOGS`. (In 4.1.0, `SHOW MASTER LOGS` was renamed to `SHOW BINARY LOGS`. Now you can use either one.)
- Added `Slave_IO_State` and `Seconds_Behind_Master` columns to the output of `SHOW SLAVE STATUS`. `Slave_IO_State` indicates the state of the slave I/O thread, and `Seconds_Behind_Master` indicates the number of seconds by which the slave is late compared to the master.
- The `--lower-case-table-names=1` server option now also makes aliases case insensitive. ([Bug#534](http://bugs.mysql.com/534) [<http://bugs.mysql.com/534>])
- Changed that the relay log is flushed to disk by the slave I/O thread every time it reads a relay log event. This reduces the risk of losing some part of the relay log in case of brutal crash.

Bugs fixed:

- Fixed `mysql` parser not to erroneously interpret ‘;’ character within `/* ... */` comment as statement terminator.
- Fixed merging types and length of result set columns for `UNION` operations. The types and lengths now are determined taking into account values for all `SELECT` statements in the `UNION`, not just the first `SELECT`.
- Fixed a bug in privilege handling that caused connections from certain IP addresses to be assigned incorrect database-level privileges. A connection could be assigned the database privileges of the previous successful authentication from one of those IP addresses, even if the IP address username and database name were different. ([Bug#1636](http://bugs.mysql.com/1636) [<http://bugs.mysql.com/1636>])
- Error-handling functions were not called properly when an error resulted from `[CREATE | REPLACE | INSERT] ... SELECT` statements.
- `HASH`, `BTREE`, `RTREE`, `ERRORS`, and `WARNINGS` no longer are reserved words. ([Bug#724](http://bugs.mysql.com/724) [<http://bugs.mysql.com/724>])
- Fix for bug in `ROLLUP` when all tables were `const` tables. ([Bug#714](http://bugs.mysql.com/714) [<http://bugs.mysql.com/714>])
- Fixed a bug in `UNION` that prohibited `NULL` values from being inserted into result set columns where the first `SELECT` of the `UNION` retrieved `NOT NULL` columns. The type and `max_length` of the result column is now defined based on all `UNION` parts.

- Fixed name resolution of columns of reduced subqueries in unions. ([Bug#745](#) [<http://bugs.mysql.com/745>])
- Fixed memory overrun in subqueries in select list with `WHERE` clause bigger than outer query `WHERE` clause. ([Bug#726](#) [<http://bugs.mysql.com/726>])
- Fixed a bug that caused `MyISAM` tables with `FULLTEXT` indexes created in 4.0.x to be unreadable in 4.1.x.
- Fixed a data loss bug in `REPAIR TABLE . . . USE_FRM` when used with tables that contained `TIMESTAMP` columns and were created in 4.0.x.
- Fixed reduced subquery processing in `ORDER BY/GROUP BY` clauses. ([Bug#442](#) [<http://bugs.mysql.com/442>])
- Fixed name resolution of outer columns of subquery in `INSERT/REPLACE` statements. ([Bug#446](#) [<http://bugs.mysql.com/446>])
- Fixed bug in marking columns of reduced subqueries. ([Bug#679](#) [<http://bugs.mysql.com/679>])
- Fixed a bug that made `CREATE FULLTEXT INDEX` syntax illegal.
- Fixed a crash when a `SELECT` that required a temporary table (marked by `Using temporary` in `EXPLAIN` output) was used as a derived table in `EXPLAIN` command. ([Bug#251](#) [<http://bugs.mysql.com/251>])
- Fixed a rare table corruption bug in `DELETE` from a big table with a `new` (created by MySQL-4.1) full-text index.
- `LAST_INSERT_ID( )` now returns 0 if the last `INSERT` statement didn't insert any rows.
- Fixed missing last character in function output. ([Bug#447](#) [<http://bugs.mysql.com/447>])
- Fixed a rare replication bug when a transaction spanned two or more relay logs, and the slave was stopped while executing the part of the transaction that was in the second or later relay log. Then replication would resume at the beginning of the second or later relay log, which was incorrect. (It should resume at `BEGIN`, in the first relay log.) ([Bug#53](#) [<http://bugs.mysql.com/53>])
- `CONNECTION_ID( )` now is properly replicated. ([Bug#177](#) [<http://bugs.mysql.com/177>])
- The new `PASSWORD( )` function in 4.1 is now properly replicated. ([Bug#344](#) [<http://bugs.mysql.com/344>])
- Fixed a bug with double freed memory.
- Fixed a crashing bug in `UNION` operations that involved temporary tables.
- Fixed a bug that under certain circumstances could allow a privilege escalation via database wild-cards in `GRANT`. ([Bug#3924](#) [<http://bugs.mysql.com/3924>], [CVE-2004-0957](#) [<http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2004-0957>])
- Fixed a crashing bug in `DERIVED TABLES` when `EXPLAIN` is used on a `DERIVED TABLES` with a join.
- Fixed a crashing bug in `DELETE` with `ORDER BY` and `LIMIT` caused by an uninitialized array of reference pointers.
- Fixed a bug in the `USER( )` function caused by an error in the size of the allocated string.

- Fixed a crashing bug when attempting to create a table containing a spatial (GIS) column with a storage engine that does not support spatial types.
- Fixed a crashing bug in `UNION` caused by the empty select list and a non-existent column being used in some of the individual `SELECT` statements.
- Fixed a replication bug with a 3.23 master and a 4.0 slave: The slave lost the replicated temporary tables if `FLUSH LOGS` was issued on the master. (Bug#254 [http://bugs.mysql.com/254])
- Fixed a security bug: A server compiled without SSL support still allowed connections by users who had the `REQUIRE SSL` option specified for their accounts.
- When an undefined user variable was used in a updating query on the master (such as `INSERT INTO t VALUES (@a)`, where `@a` had never been set by this connection before), the slave could replicate the query incorrectly if a previous transaction on the master used a user variable of the same name. (Bug#1331 [http://bugs.mysql.com/1331])
- Fixed bug with prepared statements: Using the `?` prepared statement parameter as the argument to certain functions or statement clauses caused a server crash when `mysql_prepare()` was invoked. (Bug#1500 [http://bugs.mysql.com/1500])
- Fixed bug with prepared statements: after call to `mysql_prepare` placeholders became allowed in all consequent statements, even if they are not prepared (Bug#1946 [http://bugs.mysql.com/1946])
- `SLAVE START` (which is a deprecated syntax, `START SLAVE` should be used instead) could crash the slave. (Bug#2516 [http://bugs.mysql.com/2516])
- Fixed bug in `ALTER TABLE RENAME`, when rename to the table with the same name in another database silently dropped destination table if it existed. (Bug#2628 [http://bugs.mysql.com/2628])
- When not specify hostname in `SET PASSWORD FOR user` it's now defaulted to `%` instead of the current host.

## D.1.23. Changes in release 4.1.0 (03 April 2003: Alpha)

Functionality added or changed:

- **Warning: Incompatible change:** `TIMESTAMP` is now returned as a string of type `'YYYY-MM-DD HH:MM:SS'` and different timestamp lengths are not supported.  
  
This change was necessary for SQL standards compliance. In a future version, a further change will be made (backward compatible with this change), allowing the timestamp length to indicate the desired number of digits of fractions of a second.
- Renamed `SHOW MASTER LOGS` statement to `SHOW BINARY LOGS`.
- Allow `DEFAULT(col_name)` in expressions; it produces the column's default value.
- Added `--compatible` option to `mysqldump` for producing output that is compatible with other database systems or with older MySQL servers.
- The `--opt` option for `mysqldump` now is enabled by default, as are all the options implied by `-opt`.
- New `CHARSET()` and `COLLATION()` functions to return the character set and collation of a string.
- Allow index type to be specified explicitly for some storage engines via `USING type_name` syn-

tax in index definition.

- New function `IS_USED_LOCK()` for determining the connection identifier of the client that holds a given advisory lock.
- New more secure client authentication based on 45-byte passwords in the `user` table. (CVE-2000-0981 [<http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2000-0981>])
- Added `old-password` command to `mysqladmin` for changing password but storing it using the old password-hashing format.
- New `CRC32()` function to compute cyclic redundancy check value.
- On Windows, we are now using shared memory to communicate between server and client when they are running on the same machine and you are connecting to `localhost`.
- `REPAIR TABLE` of `MyISAM` tables now uses less temporary disk space when sorting char columns.
- `DATE/DATETIME` checking is now a bit stricter to support the ability to automatically distinguish between date, datetime, and time with microseconds. For example, dates of type `YYYYMMDD HHMM` are no longer supported; you must either have separators between each `DATE/TIME` part or not at all.
- Server side help for all MySQL functions. One can now type `help week` in the `mysql` client and get help for the `week()` function.
- Added new `mysql_get_server_version()` C API client function.
- Fixed bug in `libmysqlclient` that fetched column defaults.
- Fixed bug in `mysql` command-line client in interpreting quotes within comments. (Bug#539 [<http://bugs.mysql.com/539>])
- Added `record_in_range()` method to `MERGE` tables to be able to choose the right index when there are many to choose from.
- Replication now works with `RAND()` and user variables `@var`.
- Allow one to change mode for `ANSI_QUOTES` on the fly.
- `EXPLAIN SELECT` now can be killed. See [Section 13.5.5.3, “KILL Syntax”](#).
- `REPAIR TABLE` and `OPTIMIZE TABLE` now can be killed. See [Section 13.5.5.3, “KILL Syntax”](#).
- Allow empty index lists to be specified for `USE INDEX`, `IGNORE INDEX`, and `FORCE INDEX`.
- `DROP TEMPORARY TABLE` now drops only temporary tables and doesn't end transactions.
- Added support for `UNION` in derived tables.
- New faster client/server protocol that supports prepared statements, bound parameters, and bound result columns, binary transfer of data, warnings.
- Added database and real table name (in case of alias) to the `MYSQL_FIELD` structure.
- Multi-line queries: You can now issue several queries at once and then read the results in one go.
- In `CREATE TABLE foo (a INT not null primary key)` the `PRIMARY` word is now optional.

- In `CREATE TABLE` the attribute `SERIAL` is now an alias for `BIGINT UNSIGNED NOT NULL AUTO_INCREMENT UNIQUE`.
- `SELECT ... FROM DUAL` is an alias for `SELECT ...` (To be compatible with some other database systems).
- If one creates a too long `CHAR/VARCHAR` it's now automatically changed to `TEXT` or `BLOB`; One get a warning in this case.
- One can specify the different `BLOB/TEXT` types with the syntax `BLOB(length)` and `TEXT(length)`. MySQL automatically changes it to one of the internal `BLOB/TEXT` types.
- `CHAR BYTE` is an alias for the `CHAR BINARY` data type.
- `VARCHARACTER` is an alias for `VARCHAR`.
- New operators `integer MOD integer` and `integer DIV integer`. `DIV` is now a reserved word.
- `SERIAL DEFAULT VALUE` added as an alias for `AUTO_INCREMENT`.
- `TRUE` and `FALSE` added as alias for 1 and 0, respectively.
- Aliases are now forced in derived tables, as per standard SQL.
- Fixed `SELECT .. LIMIT 0` to return proper row count for `SQL_CALC_FOUND_ROWS`.
- One can specify many temporary directories to be used in a round-robin fashion with: `--tmpdir=dirname1:dirname2:dirname3`.
- Subqueries: `SELECT * from t1 where t1.a=(SELECT t2.b FROM t2)`.
- Derived tables:

```
SELECT a.col1, b.col2
FROM (SELECT MAX(col1) AS col1 FROM root_table) a,
other_table b
WHERE a.col1=b.col1;
```

- Character sets to be defined per column, table and database.
- Unicode (UTF8) support.
- New `CONVERT(... USING ...)` syntax for converting string values between character sets.
- `BTREE` index on `MEMORY (HEAP)` tables.
- Faster embedded server (new internal communication protocol).
- One can add a comment per column in `CREATE TABLE`.
- `SHOW FULL COLUMNS FROM tbl_name` shows column comments.
- `ALTER DATABASE`.
- Support for GIS (Geometrical data). See [Chapter 16, Spatial Extensions](#).
- `SHOW [COUNT(*)] WARNINGS` shows warnings from the last command.
- One can specify a data type for a column in `CREATE TABLE ... SELECT` by defining the

column in the `CREATE` part.

```
CREATE TABLE foo (a TINYINT NOT NULL) SELECT b+1 AS a FROM bar;
```

- `expr SOUNDS LIKE expr` same as `SOUNDEX(expr)=SOUNDEX(expr)`.
- Added new `VARIANCE(expr)` function returns the variance of `expr`
- One can create a table from the existing table using `CREATE [TEMPORARY] TABLE [IF NOT EXISTS] table (LIKE table)`. The table can be either normal or temporary.
- New options `--reconnect` and `--skip-reconnect` for the `mysql` client, to reconnect automatically or not if the connection is lost.
- `START SLAVE (STOP SLAVE)` no longer returns an error if the slave is started (stopped); it returns a warning instead.
- `SLAVE START` and `SLAVE STOP` are no longer accepted by the query parser; use `START SLAVE` and `STOP SLAVE` instead.

## D.2. Changes in release 4.0.x (Recent; still supported)

Version 4.0 of the MySQL server includes many enhancements and new features:

- The `InnoDB` storage engine is now included in the standard binaries, adding transactions, row-level locking, and foreign keys. See [Section 14.2, “The InnoDB Storage Engine”](#).
- A query cache, offering vastly increased performance for many applications. By caching complete result sets, later identical queries can return instantly. See [Section 5.13, “The MySQL Query Cache”](#).
- Improved full-text indexing with boolean mode, truncation, and phrase searching. See [Section 12.7, “Full-Text Search Functions”](#).
- Enhanced `MERGE` tables, now supporting `INSERT` statements and `AUTO_INCREMENT`. See [Section 14.3, “The MERGE Storage Engine”](#).
- `UNION` syntax in `SELECT`. See [Section 13.2.7.2, “UNION Syntax”](#).
- Multiple-table `DELETE` statements. See [Section 13.2.1, “DELETE Syntax”](#).
- `libmysqld`, the embedded server library. See [Section 17.1, “libmysqld, the Embedded MySQL Server Library”](#).
- Additional `GRANT` privilege options for even tighter control and security. See [Section 13.5.1.2, “GRANT Syntax”](#).
- Management of user resources in the `GRANT` system, particularly useful for ISPs and other hosting providers. See [Section 5.8.4, “Limiting Account Resources”](#).
- Dynamic server variables, allowing configuration changes to be made without having to stop and restart the server. See [Section 13.5.3, “SET Syntax”](#).
- Improved replication code and features. See [Chapter 6, Replication](#).
- Numerous new functions and options.



- Changes to existing code for enhanced performance and reliability.

For a full list of changes, please refer to the changelog sections for each individual 4.0.x release.

## D.2.1. Changes in release 4.0.28 (Not yet released)

This is a bugfix release for the previous production release family.

This section documents all changes and bug fixes that have been applied since the last official MySQL release. If you would like to receive more fine-grained and personalized *update alerts* about fixes that are relevant to the version and features you use, please consider subscribing to *MySQL Network* (a commercial MySQL offering). For more details please see <http://www.mysql.com/network/advisors.html>.

Functionality added or changed:

- The `mysqldumpslow` script has been moved from client RPM packages to server RPM packages. This corrects a problem where `mysqldumpslow` could not be used with a client-only RPM install, because it depends on `my_print_defaults` which is in the server RPM. (Bug#20216 [<http://bugs.mysql.com/20216>])

Bugs fixed:

- `mysqlhotcopy` did not copy `RAID` directories with names that contained non-decimal hex digits. (It copied only directories containing the characters 0 through 9 and ignored those containing a through f.) (Bug#18777 [<http://bugs.mysql.com/18777>])
- Using `SELECT` and a table join while running a concurrent `INSERT` operation would join incorrect rows. (Bug#14400 [<http://bugs.mysql.com/14400>])
- A query with a `WHERE` clause containing `column = ELT(int_value_1, value_list)` OR `column = ELT(int_value_2, value_list)` could return unexpected results. (Bug#12728 [<http://bugs.mysql.com/12728>])

## D.2.2. Changes in release 4.0.27 (06 May 2006)

This is a security fix release and bugfix release for the MySQL 4.0 production release family.

This release includes the patches for recently reported security vulnerabilities in the MySQL client-server protocol. We would like to thank Stefano Di Paola <[stefano.dipaola@wisec.it](mailto:stefano.dipaola@wisec.it)> for finding and reporting these to us.

Functionality added or changed:

- The `MySQL-server` RPM now explicitly assigns the `mysql` system user to the `mysql` user group during the postinstallation process. This corrects an issue with upgrading the server on some Linux distributions whereby a previously existing `mysql` user was not changed to the `mysql` group, resulting in wrong groups for files created following the installation. (Bug#12823 [<http://bugs.mysql.com/12823>])
- Better detection of connection timeout for replication servers on Windows allows elimination of extraneous `Lost connection` errors in the error log. (Bug#5588 [<http://bugs.mysql.com/5588>])

Bugs fixed:

- **Security fix:** A malicious client, using specially crafted invalid login or `COM_TABLE_DUMP` packets was able to read uninitialized memory, which potentially, though unlikely in MySQL, could have led to an information disclosure. (CVE-2006-1516 [<http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2006-1516>], CVE-2006-1517 [<http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2006-1517>]) Thanks to Stefano Di Paola <[stefano.dipaola@wisec.it](mailto:stefano.dipaola@wisec.it)> for finding and reporting this bug.
- `MySQL-shared-compat-4.0.26-0.i386.rpm` incorrectly depend on `glibc` 2.3 and cannot not be installed on a `glibc` 2.2 system. For MySQL 4.0, we recommend using the older `MySQL-shared-compat-4.0.25-0.i386.rpm` package. (Bug#16539 [<http://bugs.mysql.com/16539>])
- Running `myisampack` followed by `myisamchk` with the `--unpack` option would corrupt the `auto_increment` key. (Bug#12633 [<http://bugs.mysql.com/12633>])
- When `myisamchk` needed to rebuild a table, `AUTO_INCREMENT` information was lost. (Bug#10405 [<http://bugs.mysql.com/10405>])
- Avoid trying to include `<asm/atomic.h>` when it doesn't work in C++ code. (Bug#13621 [<http://bugs.mysql.com/13621>])
- `BIT_COUNT()` could return an incorrect value for right table columns in a `LEFT JOIN`. (Bug#13044 [<http://bugs.mysql.com/13044>])
- MySQL would not compile on Linux distributions that use the tinfo library. (Bug#18912 [<http://bugs.mysql.com/18912>])
- An `UPDATE` statement which tried to update a column with a name beginning with an asterisk would cause the server to crash. This was because the server would wrongly expand the `*` character to the list of all table columns, causing the list of columns to become longer than the list of values. Now the server performs this expansion only if the `*` character is followed by a space. (Bug#15610 [<http://bugs.mysql.com/15610>])
- An `INSERT ... SELECT` statement between tables in a `MERGE` set can return errors when statement involves insert into child table from merge table or vice-versa. (Bug#5390 [<http://bugs.mysql.com/5390>])
- Fixed problems with static variables to allow building on Fedora Core 3. (Bug#6554 [<http://bugs.mysql.com/6554>])
- A `LIMIT`-related optimization failed to take into account that `MyISAM` table indexes can be disabled, causing Error 124 when it tried to use such an index. (Bug#14616 [<http://bugs.mysql.com/14616>])
- For a table that had been opened with `HANDLER OPEN`, issuing `OPTIMIZE TABLE`, `ALTER TABLE`, or `REPAIR TABLE` caused a server crash. (Bug#14397 [<http://bugs.mysql.com/14397>])
- Queries of the form `(SELECT ...) ORDER BY ...` were being treated as a `UNION`. This improperly resulted in only distinct values being returned (because `UNION` by default eliminates duplicate results). Also, references to column aliases in `ORDER BY` clauses following parenthesized `SELECT` statements were not resolved properly. (Bug#7672 [<http://bugs.mysql.com/7672>])
- `SELECT DISTINCT` with a `GROUP BY` clause caused a server crash. (Bug#13855 [<http://bugs.mysql.com/13855>])
- `SHOW CREATE TABLE` did not display any `FOREIGN KEY` clauses if a temporary file could not

be created. Now `SHOW CREATE TABLE` displays an error message in an SQL comment if this occurs. (Bug#13002 [<http://bugs.mysql.com/13002>])

- MySQL programs in binary distributions for Solaris 8/9/10 x86 systems would not run on Pentium III machines. (Bug#6772 [<http://bugs.mysql.com/6772>])
- Queries against a `MERGE` table that has a composite index could produce incorrect results. (Bug#9112 [<http://bugs.mysql.com/9112>])
- The counters for the `Key_read_requests`, `Key_reads`, `Key_write_requests`, and `Key_writes` status variables were changed from `unsigned long` to `unsigned longlong` to accommodate larger values before the variables roll over and restart from 0. (Bug#12920 [<http://bugs.mysql.com/12920>])
- A concurrency problem for `CREATE ... SELECT` could cause a server crash. (Bug#12845 [<http://bugs.mysql.com/12845>])
- On HP-UX 11.x (PA-RISC), the `-L` option caused `mysqlimport` to crash. (Bug#12958 [<http://bugs.mysql.com/12958>])
- The server crashed when one thread resized the query cache while another thread was using it. (Bug#12848 [<http://bugs.mysql.com/12848>])

### D.2.3. Changes in release 4.0.26 (08 September 2005)

Functionality added or changed:

- Added the `mysql_get_client_version` C API function to the embedded server library. (It was present in the regular client library but inadvertently omitted from the embedded library.) (Bug#10266 [<http://bugs.mysql.com/10266>])

Bugs fixed:

- An optimizer estimate of zero rows for a non-empty `InnoDB` table used in a left or right join could cause incomplete rollback for the table. (Bug#12779 [<http://bugs.mysql.com/12779>])
- Query cache is switched off if a thread (connection) has tables locked. This prevents invalid results where the locking thread inserts values between a second thread connecting and selecting from the table. (Bug#12385 [<http://bugs.mysql.com/12385>])
- For PKG installs on Mac OS X, the preinstallation and postinstallation scripts were being run only for new installations and not for upgrade installations, resulting in an incomplete installation process. (Bug#11380 [<http://bugs.mysql.com/11380>])
- On Windows, applications that used the embedded server made it not possible to remove certain files in the data directory, even after the embedded server had been shut down. This occurred because a file descriptor was being held open. (Bug#12177 [<http://bugs.mysql.com/12177>])
- Creation of the `mysql` group account failed during the RPM installation. (Bug#12348 [<http://bugs.mysql.com/12348>])
- Attempting to repair a table having a fulltext index on a column containing words whose length exceeded 21 characters and where `myisam_repair_threads` was greater than 1 would crash the server. (Bug#11684 [<http://bugs.mysql.com/11684>])

- When two threads compete for the same table, a deadlock could occur if one thread has also a lock on another table through `LOCK TABLES` and the thread is attempting to remove the table in some manner and the other thread want locks on both tables. ([Bug#10600](http://bugs.mysql.com/10600) [<http://bugs.mysql.com/10600>])

## D.2.4. Changes in release 4.0.25 (05 July 2005)

Functionality added or changed:

- Security improvement: Applied a patch to fix a UDF library-loading vulnerability that could result in a buffer overflow and code execution. (<http://www.appsecinc.com/resources/alerts/mysql/2005-002.html>)
- Added `--with-big-tables` compilation option to `configure`. (Previously it was necessary to pass `-DBIG_TABLES` to the compiler manually in order to enable large table support.) See [Section 2.9.2, “Typical configure Options”](#), for details.

Bugs fixed:

- On Mac OS X, `libmysqlclient_r.a` now is built with `--fno-common` to make it possible to link a shared two-level namespace library against `libmysqlclient_r.a`. ([Bug#10638](http://bugs.mysql.com/10638) [<http://bugs.mysql.com/10638>])
- An error in the implementation of the MyISAM compression algorithm caused `myisampack` to fail with very large sets of data (total size of all the records in a single column needed to be  $\geq 3$  GB in order to trigger this issue). ([Bug#8321](http://bugs.mysql.com/8321) [<http://bugs.mysql.com/8321>])
- A problem with the `my_global.h` file caused compilation of MySQL to fail on single-processor Linux systems running 2.6 kernels. ([Bug#10364](http://bugs.mysql.com/10364) [<http://bugs.mysql.com/10364>])
- Fixed a portability problem testing for `crypt()` support that caused compilation problems when using OpenSSL/yaSSL on HP-UX and Mac OS X. ([Bug#10675](http://bugs.mysql.com/10675) [<http://bugs.mysql.com/10675>], [Bug#11150](http://bugs.mysql.com/11150) [<http://bugs.mysql.com/11150>])
- MyISAM table corruption could occur with `ANALYZE TABLE` if a write lock was acquired with `LOCK TABLES` and then an `INSERT` or `DELETE` was done prior to analyzing the table. ([Bug#10901](http://bugs.mysql.com/10901) [<http://bugs.mysql.com/10901>])
- Fixed a server crash resulting from `CREATE TABLE ... SELECT` that selected from a table being altered by `ALTER TABLE`. ([Bug#10224](http://bugs.mysql.com/10224) [<http://bugs.mysql.com/10224>])
- InnoDB: In `DROP DATABASE`, check for all referencing tables from other databases before dropping any tables. ([Bug#10335](http://bugs.mysql.com/10335) [<http://bugs.mysql.com/10335>])
- Fixed a problem with incorrect constant propagation resulting in incorrect evaluation of `AND/OR` queries. ([Bug#10095](http://bugs.mysql.com/10095) [<http://bugs.mysql.com/10095>])
- Fixed wrong buffer usage for auto-increment key with blob part that caused `CHECK TABLE` to report that the table was wrong. ([Bug#10045](http://bugs.mysql.com/10045) [<http://bugs.mysql.com/10045>])
- No error was raised for `BOOLEAN` full-text searches for storage engines that do not support full-text. ([Bug#7709](http://bugs.mysql.com/7709) [<http://bugs.mysql.com/7709>])
- The test in `configure` to see whether `CXX` specified `gcc` failed if `gcc` was specified as a full pathname. ([Bug#9690](http://bugs.mysql.com/9690) [<http://bugs.mysql.com/9690>])

- In the `mysql_real_escape_string()` C API function, when a multi-byte character is encountered that is illegal in the current character set, escape only the first byte, not each byte. This avoids creating a valid character from an invalid one. (Bug#9864 [http://bugs.mysql.com/9864]; this is a backport of Bug#8378 [http://bugs.mysql.com/8378] from MySQL 4.1.11 to 4.0.25)
- Fixed a deadlock resulting from use of `FLUSH TABLES WITH READ LOCK` while an `INSERT DELAYED` statement is in progress. (Bug#7823 [http://bugs.mysql.com/7823])
- Fixed a segmentation fault in `mysqlcheck` that occurred when the last table checked in `--auto-repair` mode returned an error (such as the table being a `MERGE` table). (Bug#9492 [http://bugs.mysql.com/9492])
- Fixed faulty display of `TIMESTAMP` columns retrieved as `col_name+0` while the `new` system variable is set to 1. (Bug#8894 [http://bugs.mysql.com/8894])
- Queries containing `CURRENT_USER()` incorrectly were registered in the query cache. (Bug#9796 [http://bugs.mysql.com/9796])
- An `UPDATE` that updated only some of the columns in a multiple-column index could result in a loop. (Bug#8942 [http://bugs.mysql.com/8942])
- `REPAIR TABLE` did not invalidate query results in the query cache that were generated from the table. (Bug#8480 [http://bugs.mysql.com/8480])
- Fixed a bug that caused concurrent inserts to be allowed into the tables in the `SELECT ... UNION ...` part of `INSERT ... SELECT ... UNION ...`. This could result in the incorrect order of queries in the binary log. (Bug#9922 [http://bugs.mysql.com/9922])
- Fixed a bug that under certain circumstances could allow a privilege escalation via database wildcards in `GRANT`. (Bug#3924 [http://bugs.mysql.com/3924], CVE-2004-0957 [http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2004-0957])
- `<=>` was not properly comparing `NULL` values in the `WHERE` clause of outer joins. (Bug#8711 [http://bugs.mysql.com/8711])
- InnoDB: Fixed a bug: MySQL-4.0.23 and 4.0.24 could complain that an InnoDB table created with MySQL-3.23.49 or earlier was in the new compact InnoDB table format of 5.0.3 or later, and InnoDB would refuse to use that table. (The same bug exists in 4.1.8 - 4.1.10.) There is nothing wrong with the table, it is `mysqld` that is in error. Workaround: wait that 4.0.25 or 4.1.11 is released before doing an upgrade, or dump the table and re-create it with any MySQL version  $\geq$  3.23.50 before upgrading to 4.0.23 or 4.0.24.

## D.2.5. Changes in release 4.0.24 (04 March 2005)

Functionality added or changed:

- Security improvement: The server creates `.frm`, `.MYD`, `.MYI`, `.MRG`, `.ISD`, and `.ISM` table files only if a file with the same name does not already exist. Thanks to Stefano Di Paola <stefano.dipaola@wisec.it> for finding and informing us about this issue. (CVE-2005-0711 [http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2005-0711])
- Security improvement: User-defined functions should have at least one symbol defined in addition to the `xxx` symbol that corresponds to the main `xxx()` function. These auxiliary symbols correspond to the `xxx_init()`, `xxx_deinit()`, `xxx_reset()`, `xxx_clear()`, and `xxx_add()` functions. `mysqld` by default no longer loads UDFs unless they have at least one auxiliary symbol defined in addition to the main symbol. The `--allow-suspicious-udfs` option controls

whether UDFs that have only an `xxx` symbol can be loaded. By default, the option is off. `mysqld` also checks UDF filenames when it reads them from the `mysql.func` table and rejects those that contain directory pathname separator characters. (It already checked names as given in `CREATE FUNCTION` statements.) See [Section 19.2.4.1, “UDF Calling Sequences for Simple Functions”](#), [Section 19.2.4.2, “UDF Calling Sequences for Aggregate Functions”](#), and [Section 19.2.4.6, “User-Defined Function Security Precautions”](#). Thanks to Stefano Di Paola <[stefano.dipaola@wisec.it](mailto:stefano.dipaola@wisec.it)> for finding and informing us about this issue. ([CVE-2005-0709](#) [<http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2005-0709>], [CVE-2005-0710](#) [<http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2005-0710>])

- **InnoDB:** Added configuration option and settable global variable `innodb_autoextend_increment` for setting the size in megabytes by which InnoDB tablespaces are extended when they become full. The default value is 8, corresponding to the fixed increment of 8MB in previous versions of MySQL.
- **InnoDB:** Do not acquire an internal InnoDB table lock in `LOCK TABLES` if `AUTOCOMMIT=1`. This helps in porting old MyISAM applications to InnoDB. InnoDB table locks in that case caused deadlocks very easily.

Bugs fixed:

- `AES_DECRYPT(col_name,key)` could fail to return `NULL` for invalid values in `col_name`, if `col_name` was declared as `NOT NULL`. ([Bug#8669](#) [<http://bugs.mysql.com/8669>])
- `FOUND_ROWS()` returned an incorrect value after a `SELECT SQL_CALC_FOUND_ROWS DISTINCT` statement that selected constants and included `GROUP BY` and `LIMIT` clauses. ([Bug#7945](#) [<http://bugs.mysql.com/7945>])
- Index cardinality was not being updated properly for `TEMPORARY` tables under some circumstances, such as `CREATE TABLE ... SELECT` followed by `ANALYZE TABLE`. ([Bug#7519](#) [<http://bugs.mysql.com/7519>])
- Fixed a server crash caused by `DELETE FROM tbl_name ... WHERE ... ORDER BY tbl_name.col_name` when the `ORDER BY` column was qualified with the table name. ([Bug#8392](#) [<http://bugs.mysql.com/8392>])
- Fixed a bug in `MATCH ... AGAINST` in natural language mode that could cause a server crash if the `FULLTEXT` index was not used in a join (`EXPLAIN` did not show `fulltext` join mode) and the search query matched no rows in the table ([Bug#8522](#) [<http://bugs.mysql.com/8522>]).
- Platform and architecture information in version information produced for `--version` option on Windows was always `Win95/Win98 (i32)`. More accurately determine platform as `Win32` or `Win64` for 32-bit or 64-bit Windows, and architecture as `ia32` for x86, `ia64` for Itanium, and `axp` for Alpha. ([Bug#4445](#) [<http://bugs.mysql.com/4445>])
- Fixed an optimization problem that allowed a negative number to be stored in a `DOUBLE UNSIGNED` column when it was assigned a value from a signed `DOUBLE` column. ([Bug#7700](#) [<http://bugs.mysql.com/7700>])
- Fixed a failure of multiple-table updates to replicate properly on slave servers when `-replicate-*-table` options had been specified. ([Bug#7011](#) [<http://bugs.mysql.com/7011>])
- Renamed `set_bit()` and `clear_bit()` functions in source code to avoid a conflict with functions of the same names in Linux kernel header files. ([Bug#7971](#) [<http://bugs.mysql.com/7971>])
- Part of the information being used to cache access-permission lookups was not always reinitialized properly, particularly for connections from localhost on Windows. The result was connection fail-

- ures that appeared to occur randomly. ([Bug#5569](http://bugs.mysql.com/5569) [<http://bugs.mysql.com/5569>])
- Corrected a problem with the `QUOTE ( )` function returning bad results. ([Bug#8248](http://bugs.mysql.com/8248) [<http://bugs.mysql.com/8248>])
  - Fixed a problem where `INSERT INTO ...SELECT` failed when the source and target table were the same. ([Bug#6034](http://bugs.mysql.com/6034) [<http://bugs.mysql.com/6034>])
  - Fixed a problem where RPM installation on Linux as a non-privileged user would result in incomplete installation. ([Bug#7347](http://bugs.mysql.com/7347) [<http://bugs.mysql.com/7347>])
  - Change thread stack size used for building Linux RPM distributions to avoid warnings about stack size during server startup. ([Bug#6226](http://bugs.mysql.com/6226) [<http://bugs.mysql.com/6226>])
  - Fixed a symlink vulnerability in the `mysqlaccess` script. Reported by Javier Fernandez-Sanguino Pena and [Debian Security Audit Team](http://www.debian.org/security/audit) [<http://www.debian.org/security/audit>]. ([CVE-2005-0004](http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2005-0004) [<http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2005-0004>])
  - Fixed support for C API function `mysql_list_fields( )`, which was accidentally broken in 4.0.22 ([Bug#6761](http://bugs.mysql.com/6761) [<http://bugs.mysql.com/6761>])
  - Make `query_cache_wlock_invalidate` system variable visible in `SHOW VARIABLES` output. ([Bug#7594](http://bugs.mysql.com/7594) [<http://bugs.mysql.com/7594>])
  - Fixed a bug which caused `FROM_UNIXTIME( )` function to return `NULL` for zero argument instead of the Epoch. ([Bug#7515](http://bugs.mysql.com/7515) [<http://bugs.mysql.com/7515>])
  - Now in datetime values two digit year is interpreted as year in 20th or 21st century even with zero month and day. ([Bug#7297](http://bugs.mysql.com/7297) [<http://bugs.mysql.com/7297>])
  - Fixed a bug in `QUOTE` function when used in conjunction with some other string functions. This lead to severe buffer overflow and server crashing. ([Bug#7495](http://bugs.mysql.com/7495) [<http://bugs.mysql.com/7495>])
  - InnoDB: Work around a problem in AIX 5.1 patched with ML7 security patch: InnoDB would refuse to open its `ibdata` files, complaining about an operating system error 0.
  - InnoDB: Fixed a memory corruption bug if one created a table with a primary key that contained at least two column prefixes. An example: `CREATE TABLE t(a char(100), b tinyblob, PRIMARY KEY(a(5), b(10)))`.
  - InnoDB: Use native `tmpfile( )` function on Netware. All InnoDB temporary files are created under `sys:\tmp`. Previously, InnoDB temporary files were never deleted on Netware.
  - InnoDB: Honor the `--tmpdir` startup option when creating temporary files. Previously, InnoDB temporary files were always created in the temporary directory of the operating system. On Netware, InnoDB will continue to ignore `--tmpdir`. ([Bug#5822](http://bugs.mysql.com/5822) [<http://bugs.mysql.com/5822>])
  - InnoDB: Fix a theoretical hang over the adaptive hash latch in InnoDB if one runs `INSERT ... SELECT ...` (binlog not enabled), or a multiple-table `UPDATE` or `DELETE`, and only the read tables are InnoDB type, the rest are MyISAM; this also fixes [Bug#7879](http://bugs.mysql.com/7879) [<http://bugs.mysql.com/7879>] for InnoDB type tables. ([Bug#7879](http://bugs.mysql.com/7879) [<http://bugs.mysql.com/7879>])
  - InnoDB: Fixed a bug: 32-bit `mysqld` binaries built on HP-UX-11 did not work with InnoDB files greater than 2 GB in size. ([Bug#6189](http://bugs.mysql.com/6189) [<http://bugs.mysql.com/6189>])
  - InnoDB: Fixed a bug: InnoDB failed to drop a table in the background drop queue if the table was referenced by a foreign key constraint.
  - InnoDB: Fixed a bug: if we dropped a table where an `INSERT` was waiting for a lock to check a

`FOREIGN KEY` constraint, then an assertion would fail in `lock_reset_all_on_table()`, since that operation assumes no waiting locks on the table or its records.

- Fixed that, when encountering a “disk full” or “quota exceeded” write error, `MyISAM` sometimes didn't sleep and retry the write, thus resulting in a corrupted table. ([Bug#7714](http://bugs.mysql.com/7714) [<http://bugs.mysql.com/7714>])
- Fixed that a slave could crash after replicating many `ANALYZE TABLE`, `OPTIMIZE TABLE`, or `REPAIR TABLE` statements from the master. ([Bug#6461](http://bugs.mysql.com/6461) [<http://bugs.mysql.com/6461>], [Bug#7658](http://bugs.mysql.com/7658) [<http://bugs.mysql.com/7658>])
- Fixed a bug where MySQL was allowing concurrent updates (inserts, deletes) to a table if binary logging is enabled. Changed to ensure that all updates are executed in a serialized fashion, because they are executed serialized when binlog is replayed. ([Bug#7879](http://bugs.mysql.com/7879) [<http://bugs.mysql.com/7879>])
- Fixed a bug in replication that caused the master to stamp generated statements (such as `SET` commands) with an `error_code` intended only for another statement. This could happen, for example, when a statements generates a duplicate key error on the master but must be replicated. ([Bug#8412](http://bugs.mysql.com/8412) [<http://bugs.mysql.com/8412>])
- Documented problem with using `mysqldump` in 4.0.x to dump `TIMESTAMP(2)` and `TIMESTAMP(4)` data types. ([Bug#6530](http://bugs.mysql.com/6530) [<http://bugs.mysql.com/6530>])

## D.2.6. Changes in release 4.0.23 (18 December 2004)

**Note:** Due to a `libtool`-related bug in the source distribution, the creation of shared `libmysqlclient` libraries was not possible (the resulting files were missing the `.so` file name extension). The file `ltmain.sh` was updated to fix this problem and the resulting source distribution was released as `mysql-4.0.23a.tar.gz`. This modification did not affect the binary packages. ([Bug#7401](http://bugs.mysql.com/7401) [<http://bugs.mysql.com/7401>])

Functionality added or changed:

- Added `--hex-blob` option to `mysqldump` for dumping binary string columns using hexadecimal notation.
- Added `mysql_hex_string()` C API function that hex-encodes a string.
- InnoDB: Do not periodically write `SHOW INNODB STATUS` information to a temporary file unless the configuration option `innodb_status_file=1` is set.
- InnoDB: Made the foreign key parser better aware of quotes. ([Bug#6340](http://bugs.mysql.com/6340) [<http://bugs.mysql.com/6340>])
- `mysqlbinlog` now prints an informative commented line (thread id, timestamp, server id, and so forth) before each `LOAD DATA INFILE`, like it does for other queries; unless `--short-form` is used.

Bugs fixed:

- A multiple-table `DELETE` could cause MySQL to crash when using InnoDB tables. ([Bug#5837](http://bugs.mysql.com/5837) [<http://bugs.mysql.com/5837>], [Bug#6378](http://bugs.mysql.com/6378) [<http://bugs.mysql.com/6378>])
- Corrected accounts in the `mysql.user` table in Windows distributions that had been created with a `Host` value of `build` rather than `%`. ([Bug#6000](http://bugs.mysql.com/6000) [<http://bugs.mysql.com/6000>])



- Prevent adding `CREATE TABLE ... SELECT` query to the binary log when the insertion of new records partially failed. (Bug#6682 [<http://bugs.mysql.com/6682>])
- Fixed bug which caused `FROM_UNIXTIME()` function to return wrong result if the argument was too big. (Bug#6439 [<http://bugs.mysql.com/6439>])
- Fixed bug which caused MySQL server to store wrong values in `TIMESTAMP` columns and give wrong results for `UNIX_TIMESTAMP()` function if it was run in time zone with leap seconds. (Bug#6387 [<http://bugs.mysql.com/6387>])
- InnoDB: Fixed a bug in `LOAD DATA INFILE...REPLACE` printing duplicate key error when executing the same load query several times. (Bug#5835 [<http://bugs.mysql.com/5835>])
- InnoDB: Refuse to open new-style tables created with MySQL 5.0.3 or later. (Bug#7089 [<http://bugs.mysql.com/7089>])
- InnoDB: Do not call `rewind()` when displaying `SHOW INNODB STATUS` information on `stderr`.
- InnoDB: If one used `INSERT IGNORE` to insert several rows at a time, and the first inserts were ignored because of a duplicate key collision, then InnoDB in a replication slave assigned `AUTO_INCREMENT` values 1 bigger than in the master. This broke the MySQL replication. (Bug#6287 [<http://bugs.mysql.com/6287>])
- InnoDB: Fix two hangs: `FOREIGN KEY` constraints treated table and database names as case-insensitive. `RENAME TABLE t TO T` would hang in an endless loop if `t` had a foreign key constraint defined on it. Fix also a hang over the dictionary mutex that would occur if one tried in `ALTER TABLE` or `RENAME TABLE` to create a foreign key constraint name that collided with another existing name. (Bug#3478 [<http://bugs.mysql.com/3478>])
- InnoDB: Treat character `0xA0` as space in InnoDB's `FOREIGN KEY` parser if MySQL treats it as space in the default charset. EMS MySQL Manager inserts character `0xA0` after the table name in an `ALTER`, which confused InnoDB's parser.
- Fixed a bug which caused a crash when only the slave I/O thread was stopped and restarted. (Bug#6148 [<http://bugs.mysql.com/6148>])
- If a connection had an open transaction but had done no updates to transactional tables (for example if had just done a `SELECT FOR UPDATE` then executed a non-transactional update, that update automatically committed the transaction (thus releasing InnoDB's row-level locks etc). (Bug#5714 [<http://bugs.mysql.com/5714>])
- If a connection was interrupted by a network error and did a rollback, the network error code got stored into the `BEGIN` and `ROLLBACK` binary log events; that caused superfluous slave stops. (Bug#6522 [<http://bugs.mysql.com/6522>])
- A sequence of `BEGIN` (or `SET AUTOCOMMIT=0`), `FLUSH TABLES WITH READ LOCK`, transactional update, `COMMIT`, `FLUSH TABLES WITH READ LOCK` could hang the connection forever and possibly the MySQL server itself. This happened for example when running the `innobackup` script several times. (Bug#6732 [<http://bugs.mysql.com/6732>])

## D.2.7. Changes in release 4.0.22 (27 October 2004)

Functionality added or changed:

- The `--with-openssl` option for `configure` now accepts a path prefix as an argument. -

`--with-openssl-includes` and `--with-openssl-libs` are still supported, but are needed only to override the default values. (Bug#5494 [http://bugs.mysql.com/5494])

- Added new `--without-man` option to `configure` to suppress building/installing the manual pages. (Bug#5379 [http://bugs.mysql.com/5379])
- InnoDB: New `mysqld` option `--innodb-table-locks` and session variable `innodb_table_locks` (on by default). In applications using `AUTOCOMMIT=1` and MySQL's `LOCK TABLES` command, InnoDB's internal table locks that were added in 4.0.20 can cause deadlocks. You can set `innodb_table_locks=0` in `my.cnf` to remove that problem. See Section 14.2.17, “Restrictions on InnoDB Tables”. (Bug#3299 [http://bugs.mysql.com/3299], Bug#5998 [http://bugs.mysql.com/5998])
- InnoDB: Added the startup option and settable global variable `innodb_max_purge_lag` for delaying `INSERT`, `UPDATE` and `DELETE` operations when the purge operations are lagging. The default value of this parameter is zero, meaning that there are no delays. See Section 14.2.13, “Implementation of Multi-Versioning”.
- InnoDB: Change error code to `HA_ERR_ROW_IS_REFERENCED` if we cannot `DROP` a parent table because it is referenced by a `FOREIGN KEY` constraint.

#### Bugs fixed:

- Fixed bug in server which caused connection stall when one of deprecated `libmysqlclient` functions `mysql_create_db()`, `mysql_rm_db()` were called and were going to return error. (Bug#6081 [http://bugs.mysql.com/6081])
- Fixed returning wrong query result from query cache if a temporary table was hiding a real table after putting results to query cache. (Bug#6084 [http://bugs.mysql.com/6084])
- Fixed `ENABLE KEYS`, which failed if `tmpdir` ran out of space. Now, a full repair is done in this case. (Bug#5625 [http://bugs.mysql.com/5625])
- Fixed an improper error message when trying to drop a table which is referenced by a `FOREIGN KEY` constraint. (Bug#5784 [http://bugs.mysql.com/5784])
- Fixed a bug that allowed `FLUSH TABLE(S)` to close `HANDLER` tables. `HANDLER` tables are now reopened after a `FLUSH TABLE(S)` the next time they are used. However, they lose their file position if this happens. (Bug#4286 [http://bugs.mysql.com/4286])
- Fixed a bug that allowed `HANDLER` tables with the same alias to be opened multiple times. `HANDLER` aliases must now be unique, even though it is syntactically correct in versions below 4.1 to qualify them with their base table's database name (for example, `test_db.handler_tbl` now conflicts with `another_db.handler_tbl`). (Bug#4335 [http://bugs.mysql.com/4335])
- Fixed crash when using MySQL 4.0 with privilege tables from MySQL 5.0.
- `mysqlimport` now reads input files locally from the client host only if the `--local` option is given. Previously, it assumed incorrectly in some cases that files were local even without `--local`. (Bug#5829 [http://bugs.mysql.com/5829])
- InnoDB: Make the check for excessive semaphore waits to tolerate glitches in the system clock (do not crash the server if the system time is adjusted while InnoDB is under load.). (Bug#5898 [http://bugs.mysql.com/5898])
- InnoDB: Fixed a bug in the InnoDB `FOREIGN KEY` parser that prevented `ALTER TABLE` of tables containing '#' in their names. (Bug#5856 [http://bugs.mysql.com/5856])

- InnoDB: Fixed problem introduced in 4.0.21 where a connection starting a transaction, doing updates, then `FLUSH TABLES WITH READ LOCK`, then `COMMIT`, would cause replication slaves to stop (complaining about error 1223). Bug surfaced when using the InnoDB `innobackup` script. (Bug#5949 [<http://bugs.mysql.com/5949>])
- InnoDB: If one updated a column so that its size changed, or updated it to an externally stored (`TEXT` or `BLOB`) value, then ANOTHER externally stored column would show up as 512 bytes of good data + 20 bytes of garbage in a consistent read that fetched the old version of the row. (Bug#5960 [<http://bugs.mysql.com/5960>])
- InnoDB: Release the dictionary latch during a long cascaded `FOREIGN KEY` operation, so that we do not starve other users doing `CREATE TABLE` or other DDL operations. This caused a notorious 'Long semaphore wait' message to be printed to the `.err` log. (Bug#5961 [<http://bugs.mysql.com/5961>])
- InnoDB: Let InnoDB remember row locking type (X or S) inside `LOCK TABLES`, also over plain consistent read `SELECTs`.
- InnoDB: Fixed a bug introduced in 4.0.21. An assertion failed if one used `mysqldump` with the option `-l` or `--opt`, or if one used `LOCK TABLES ... LOCAL`. (Workaround in 4.0.21: use `--quick` and `--single-transaction`. (Bug#5538 [<http://bugs.mysql.com/5538>])
- InnoDB: Having a column prefix index in the primary key, and the same column fully in a secondary key could cause an assertion failure in `row_build_row_ref()`. (Bug#5180 [<http://bugs.mysql.com/5180>])
- Fixed a bug which resulted in an erroneously calculated number of examined rows in `UNIONS`. This value is printed in the slow query log. (Bug#5879 [<http://bugs.mysql.com/5879>])
- Fixed bug with crash of server on some values of `read_rnd_buffer_size` (Bug#5492 [<http://bugs.mysql.com/5492>])
- Fixed bug which caused truncation of values read from or into `TIMESTAMP` fields if `--new` mode was enabled. (Bug#4131 [<http://bugs.mysql.com/4131>])
- `mysqladmin` now returns a status of 0 even when the server denies access; such an error means the server is running. (Bug#3120 [<http://bugs.mysql.com/3120>])
- Fixed that if the slave SQL thread found a syntax error in a query (which should be rare, as the master parsed it successfully), it stops. (Bug#5711 [<http://bugs.mysql.com/5711>])
- Fixed that if a write to a `MyISAM` table fails because of a full disk or an exceeded disk quota, it prints a message to the error log every 10 minutes, and waits until disk becomes free. (Bug#3248 [<http://bugs.mysql.com/3248>])
- Fixed problem with symlinked databases on Windows being shown with `SHOW DATABASES` even if the database name doesn't match the given wildcard (Bug#5539 [<http://bugs.mysql.com/5539>])

## D.2.8. Changes in release 4.0.21 (06 September 2004)

Functionality added or changed:

- Print `version_comment` (from `./configure --comment` during compilation) when starting the server. Example: `Version: '4.0.21-debug' socket: '/tmp/mysql.sock' port: 0 Official MySQL Binary`

- Made the MySQL server not react to signals `SIGHUP` and `SIGQUIT` on Mac OS X 10.3. This is needed because under this OS, the MySQL server receives lots of these signals (reported as [Bug#2030](http://bugs.mysql.com/2030) [<http://bugs.mysql.com/2030>]).
- On Windows, the `mysqld-nt` and `mysqld-max-nt` servers now write error messages to the Windows event log in addition to the MySQL error log.

Bugs fixed:

- Fixed an old bug in concurrent accesses to `MERGE` tables (even one `MERGE` table and `MyISAM` tables), that could've resulted in a crash or hang of the server. ([Bug#2408](http://bugs.mysql.com/2408) [<http://bugs.mysql.com/2408>], [CVE-2004-0837](http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2004-0837) [<http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2004-0837>])
- Fixed a bug that caused incorrect results from `GROUP BY` queries with expression in `HAVING` clause that refers to a columns such as `BLOB`, `TEXT`, or `TINYBLOB`. ([Bug#4358](http://bugs.mysql.com/4358) [<http://bugs.mysql.com/4358>])
- Fixed a bug when memory was not released when `HEAP` table is dropped. It could only happen on Windows when a symlink file (`.sym`) is used and if that symlink file contained double backslashes (`\\`). ([Bug#4973](http://bugs.mysql.com/4973) [<http://bugs.mysql.com/4973>])
- Fixed a bug which prevented `TIMESTAMP(19)` fields from being created. ([Bug#4491](http://bugs.mysql.com/4491) [<http://bugs.mysql.com/4491>])
- Fixed a bug that caused wrong results in queries that were using index to search for `NULL` values in `BLOB` (`TINYBLOB`, `TEXT`, `TINYTEXT`, etc) columns of `MyISAM` tables. ([Bug#4816](http://bugs.mysql.com/4816) [<http://bugs.mysql.com/4816>])
- Fixed a bug in the function `ROUND()` reporting incorrect metadata (number of digits after the decimal point). It can be seen, for example, in `CREATE TABLE t1 SELECT ROUND(1, 34)`. ([Bug#4393](http://bugs.mysql.com/4393) [<http://bugs.mysql.com/4393>])
- Fixed precision loss bug in some mathematical functions such as `SQRT()` and `LOG()`. ([Bug#4356](http://bugs.mysql.com/4356) [<http://bugs.mysql.com/4356>])
- Fixed a long-standing problem with `LOAD DATA` with the `LOCAL` option. The problem occurs when an error happens during the `LOAD DATA` operation. Previously, the connection was broken. Now the error message is returned and connection stays open.
- Optimizer now treats `col IN (val)` the same way it does for `col = val`.
- Fixed a problem with `net_buffer_length` when building the `DBD:mysql` Perl module. ([Bug#4206](http://bugs.mysql.com/4206) [<http://bugs.mysql.com/4206>])
- `lower_case_table_names=2` (keep case for table names) was not honored with `ALTER TABLE` and `CREATE/DROP INDEX`. ([Bug#3109](http://bugs.mysql.com/3109) [<http://bugs.mysql.com/3109>])
- Fixed a crash on declaration of `DECIMAL(0, ...)` column. ([Bug#4046](http://bugs.mysql.com/4046) [<http://bugs.mysql.com/4046>])
- Fixed a bug in `IF()` function incorrectly determining the result type if aggregate functions were involved. ([Bug#3987](http://bugs.mysql.com/3987) [<http://bugs.mysql.com/3987>])
- Fixed bug in privilege checking where, under some conditions, one was able to grant privileges on the database, he has no privileges on. ([Bug#3933](http://bugs.mysql.com/3933) [<http://bugs.mysql.com/3933>])

- Fixed crash in `MATCH ... AGAINST()` on a phrase search operator with a missing closing double quote. (Bug#3870 [<http://bugs.mysql.com/3870>], CVE-2004-0956 [<http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2004-0956>])
- Fixed a bug with truncation of big values (> 4294967295) of 64-bit system variables. (Bug#3754 [<http://bugs.mysql.com/3754>])
- If `server-id` was not set using startup options but with `SET GLOBAL`, the replication slave still complained that it was not set. (Bug#3829 [<http://bugs.mysql.com/3829>])
- Fixed potential memory overrun in `mysql_real_connect()` (which required a compromised DNS server and certain operating systems). (Bug#4017 [<http://bugs.mysql.com/4017>], CVE-2004-0836 [<http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2004-0836>])
- During the installation process of the server RPM on Linux, `mysqld` was run as the `root` system user, and if you had `--log-bin=somewhere_out_of_var_lib_mysql` it created binary log files owned by `root` in this directory, which remained owned by `root` after the installation. This is now fixed by starting `mysqld` as the `mysql` system user instead. (Bug#4038 [<http://bugs.mysql.com/4038>])
- Made `DROP DATABASE` honor the value of `lower_case_table_names`. (Bug#4066 [<http://bugs.mysql.com/4066>])
- The slave SQL thread refused to replicate `INSERT ... SELECT` if it examined more than 4 billion rows. (Bug#3871 [<http://bugs.mysql.com/3871>])
- Fixed incorrect destruction of expression which led to crash of server on complex `AND/OR` expressions if query was ignored (either by a replication server because of `--replicate-*-table` rules, or by any MySQL server because of a syntax error). (Bug#3969 [<http://bugs.mysql.com/3969>], Bug#4494 [<http://bugs.mysql.com/4494>])
- Fixed that `mysqlbinlog --position --read-from-remote-server` had wrong # at lines. (Bug#4506 [<http://bugs.mysql.com/4506>])
- If `CREATE TEMPORARY TABLE t SELECT` failed while loading the data, the temporary table was not dropped. (Bug#4551 [<http://bugs.mysql.com/4551>])
- Fixed that when a multiple-table `DROP TABLE` failed to drop a table on the master server, the error code was not written to the binary log. (Bug#4553 [<http://bugs.mysql.com/4553>])
- When the slave SQL thread was replicating a `LOAD DATA INFILE` statement, it didn't show the statement in the output of `SHOW PROCESSLIST`. (Bug#4326 [<http://bugs.mysql.com/4326>])
- Fixed that `CREATE TABLE ... TYPE=HEAP ... AS SELECT...` caused replication slave to stop. (Bug#4971 [<http://bugs.mysql.com/4971>])
- Fixed that `disable-local-infile` option had no effect if client read it from a configuration file using `mysql_options(...,MYSQL_READ_DEFAULT,...)`. (Bug#5073 [<http://bugs.mysql.com/5073>])
- Fixed that `mysql-test-run` failed on the `rpl_trunc_binlog` test if running test from the installed (the target of 'make install') directory. (Bug#5050 [<http://bugs.mysql.com/5050>])
- Fixed an unlikely deadlock which could happen when using `KILL`. (Bug#4810 [<http://bugs.mysql.com/4810>])
- Fixed a crash when one connection got `KILLED` while it was doing `START SLAVE`. (Bug#4827 [<http://bugs.mysql.com/4827>])

- Made `FLUSH TABLES WITH READ LOCK` block `COMMIT` if server is running with binary logging; this ensures that the binary log position is trustable when doing a full backup of tables and the binary log. ([Bug#4953](http://bugs.mysql.com/4953) [<http://bugs.mysql.com/4953>])
- Fixed that the counter of an `auto_increment` column was not reset by `TRUNCATE TABLE` if the table was a temporary one. ([Bug#5033](http://bugs.mysql.com/5033) [<http://bugs.mysql.com/5033>])
- Made database names to compare case-insensitively in fully qualified column names (`database.table.column`) when `lower_case_table_names=1`. ([Bug#4792](http://bugs.mysql.com/4792) [<http://bugs.mysql.com/4792>])
- Fixed that `SET CHARACTER SET` was not replicated correctly. MySQL 4.1 does not have that bug. ([Bug#4500](http://bugs.mysql.com/4500) [<http://bugs.mysql.com/4500>])

## D.2.9. Changes in release 4.0.20 (17 May 2004)

**Note:** The windows packages had to be repackaged and re-released several times to resolve packaging issues (such as missing files). This did not affect the binaries included (they have not been recompiled), therefore the installation packages are of version 4.0.20d, while the binaries included still identify themselves as version 4.0.20b.

Functionality added or changed:

- From the Windows distribution, predefined accounts without passwords for remote users ("root@%", "@%") were removed (other distributions never had them).
- Phrase search in `MATCH ... AGAINST ( ... IN BOOLEAN MODE)` no longer matches partial words.

Bugs fixed:

- A crashing bug (race condition) was fixed in InnoDB diagnostic logging. It was introduced in 4.0.19. ([Bug#3596](http://bugs.mysql.com/3596) [<http://bugs.mysql.com/3596>])
- Fixed a bug in division `/` reporting incorrect metadata (number of digits after the decimal point). It can be seen, for example, in `CREATE TABLE t1 SELECT "0.01" / "3"`. ([Bug#3612](http://bugs.mysql.com/3612) [<http://bugs.mysql.com/3612>])
- Fixed a problem with non-working `DROP DATABASE` on some configurations (in particular, Linux 2.6.5 with ext3 are known to expose this bug). ([Bug#3594](http://bugs.mysql.com/3594) [<http://bugs.mysql.com/3594>])
- Fixed a symlink vulnerability in the `mysqlhotcopy` script. ([CVE-2004-0457](http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2004-0457) [<http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2004-0457>])
- Fixed that in some replication error messages, a very long query caused the rest of the message to be invisible (truncated), by putting the query last in the message. ([Bug#3357](http://bugs.mysql.com/3357) [<http://bugs.mysql.com/3357>])

## D.2.10. Changes in release 4.0.19 (04 May 2004)

**Note:** The MySQL 4.0.19 binaries were uploaded to the download mirrors on May, 10th. However, a potential crashing bug was found just before the 4.0.19 release was publicly announced and published from the 4.0 download pages at <http://dev.mysql.com/>.

A fix for the bug was pushed into the MySQL source tree shortly after it could be reproduced and is included in MySQL 4.0.20. Users upgrading from MySQL 4.0.18 should upgrade directly to MySQL 4.0.20 or later.

See [Bug#3596](http://bugs.mysql.com/3596) [http://bugs.mysql.com/3596] for details (it was reported against MySQL-4.1, but was confirmed to affect 4.0.19 as well).

Functionality added or changed:

- If length of a timestamp field is defined as 19, the timestamp is displayed as "YYYY-MM-DD HH:MM:SS". This is done to make it easier to use tables created in MySQL 4.1 to be used in MySQL 4.0.
- If you use `RAID_CHUNKS` with a value > 255 it is set to 255. This was made to ensure that all raid directories are always 2 hex bytes. ([Bug#3182](http://bugs.mysql.com/3182) [http://bugs.mysql.com/3182])
- Changed that the optimizer now considers the index specified in `FORCE INDEX` clause as a candidate to resolve `ORDER BY` as well.
- The `--log-warnings` server option now is enabled by default. Disable with `--skip-log-warnings`.
- Until now, in `SELECT ... UNION SELECT ... ORDER BY ...`, it was possible to qualify a column name in the `ORDER BY` clause with a table name. This is no longer possible. Column names in `ORDER BY` should refer to names established in the first `SELECT` of the `UNION`. ([Bug#3064](http://bugs.mysql.com/3064) [http://bugs.mysql.com/3064])
- Added `max_insert_delayed_threads` system variable as a synonym for `max_delayed_threads`.
- Added `query_cache_wlock_invalidate` system variable. It allows emulation of `MyISAM` table write-locking behavior, even for queries in the query cache. ([Bug#2693](http://bugs.mysql.com/2693) [http://bugs.mysql.com/2693])
- The keyword `MASTER_SERVER_ID` is not reserved anymore.
- The following is relevant mainly for Mac OS X users who use a case-insensitive filesystem. This is not relevant for Windows users as `InnoDB` in this case always stores filenames in lower case:

You can now force `lower_case_table_names` to 0 from the command line or a configuration file. This is useful with case-insensitive filesystems when you have previously not used `lower_case_table_names=1` or `lower_case_table_names=2` and you have created `InnoDB` tables. With `lower_case_table_names=0`, `InnoDB` tables were stored in mixed case while setting `lower_case_table_names` to a non-zero value now forces it to lower case (to make the table names case insensitive).

Because it's possible to crash `MyISAM` tables by referring to them with different case on a case-insensitive filesystem, we recommend that you use `lower_case_table_names` or `lower_case_table_names=2` on such filesystems.

The easiest way to convert to use `lower_case_table_names=2` is to dump all your `InnoDB` tables with `mysqldump`, drop them and then restore them.

- Changed that the relay log is flushed to disk by the slave I/O thread every time it reads a relay log event. This reduces the risk of losing some part of the relay log in case of brutal crash.
- When a session having open temporary tables terminates, the statement automatically written to the binary log is now `DROP TEMPORARY TABLE IF EXISTS` instead of `DROP TEMPORARY TA-`

`BLE`, for more robustness.

- Added option `--replicate-same-server-id`.

Bugs fixed:

- Added missing full-text variable `ft_stopword_file` to `myisamchk`.
- Don't allow stray `' , '` at the end of field specifications. (Bug#3481 [http://bugs.mysql.com/3481])
- `INTERVAL` now can handle big values for seconds, minutes and hours. (Bug#3498 [http://bugs.mysql.com/3498])
- Blank hostname did not work as documented for table and column privileges. Now it works the same way as `' % '`. (Bug#3473 [http://bugs.mysql.com/3473])
- Fixed a harmless buffer overflow in `replace` utility. (Bug#3541 [http://bugs.mysql.com/3541])
- Fixed `SOUNDEX()` to ignore non-alphabetic characters also in the beginning of the string. (Bug#3556 [http://bugs.mysql.com/3556])
- Fixed a bug in `MATCH . . . AGAINST()` searches when another thread was doing concurrent inserts into the `MyISAM` table in question. The first `---` full-text search `---` query could return incorrect results in this case (for example, “phantom” rows or not all matching rows, even an empty result set). The easiest way to check whether you are affected is to start `mysqld` with `-skip-concurrent-insert` switch and see whether it helps.
- Fixed bug when doing `DROP DATABASE` on a directory containing non-MySQL files. Now a proper error message is returned.
- Fixed bug in `ANALYZE TABLE` on a `BDB` table inside a transaction that hangs server thread. (Bug#2342 [http://bugs.mysql.com/2342])
- Fixed a symlink vulnerability in the `mysqlbug` script. (Bug#3284 [http://bugs.mysql.com/3284], CVE-2004-0381 [http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2004-0381])
- Fixed core dump bug in `SELECT DISTINCT` where all selected parts were constants and there were hidden columns in the created temporary table. (Bug#3203 [http://bugs.mysql.com/3203])
- Fixed core dump bug in `COUNT(DISTINCT)` when there was a lot of values and one had a big value for `max_heap_table_size`.
- Fixed problem with multiple-table-update and `BDB` tables. (Bug: #3098)
- Fixed memory leak when dropping database with `RAID` tables. (Bug#2882 [http://bugs.mysql.com/2882])
- Fixed core dump crash in replication during relay-log switch when the relay log went over `max_relay_log_size` and the slave thread did a `flush_io_cache()` at the same time.
- Fixed hangup bug when issuing multiple `SLAVE START` from different threads at the same time. (Bug#2921 [http://bugs.mysql.com/2921])
- Fixed bug when using `DROP DATABASE` with `lower_case_table_names=2`.
- Fixed wrong result in `UNION` when using `lower_case_table_names=2`. (Bug#2858 [http://bugs.mysql.com/2858])



- One can now kill threads that is 'stuck' in the join optimizer (can happen when there is MANY tables in the join in which case the optimizer can take really long time). ([Bug#2825](#) [<http://bugs.mysql.com/2825>])
- Rollback `DELETE` and `UPDATE` statements if thread is killed. ([Bug#2422](#) [<http://bugs.mysql.com/2422>])
- Ensure that all rows in an `INSERT DELAYED` statement is written at once if binary logging is enabled. ([Bug#2491](#) [<http://bugs.mysql.com/2491>]).
- Fixed bug in query cache statistic, more accurate formula linked statistic variables mentioned in the manual.
- Fixed a bug in parallel repair (`myisamchk -p, myisam_repair_threads`) - sometimes repair process failed to repair a table. ([Bug#1334](#) [<http://bugs.mysql.com/1334>])
- Fixed bugs with names of tables, databases, and columns that end to space ([Bug#2985](#) [<http://bugs.mysql.com/2985>])
- Fixed a bug in multiple-table `UPDATE` statements involving at least one constant table. Bug was exhibited in allowing non matching row to be updated. ([Bug#2996](#) [<http://bugs.mysql.com/2996>]).
- Fixed all bugs in scripts for creating/upgrading system database ([Bug#2874](#) [<http://bugs.mysql.com/2874>]) Added tests which guarantee against such bugs in the future.
- Fixed bug in `mysql` command-line client in interpreting quotes within comments. ([Bug#539](#) [<http://bugs.mysql.com/539>])
- `--set-character-set` and `--character-sets-dir` options in `myisamchk` now work.
- Fixed a bug in `mysqlbinlog` that caused one pointer to be free'd twice in some cases.
- Fixed a bug in boolean full-text search, that sometimes could lead to false matches in queries with several levels of subexpressions using `+` operator (for example, `MATCH ... AGAINST( '(+(word1 word2)) +word3*' IN BOOLEAN MODE)`).
- Fixed Windows-specific portability bugs in `myisam_ftdump`.
- Fixed a bug in multiple-table `DELETE` that was caused by foreign key constraints. If the order of the tables established by MySQL optimizer did not match parent-child order, no rows were deleted and no error message was provided. ([Bug#2799](#) [<http://bugs.mysql.com/2799>])
- Fixed a few years old bug in the range optimizer that caused a segmentation fault on some very rare queries. ([Bug#2698](#) [<http://bugs.mysql.com/2698>])
- Replication: If a client connects to a slave server and issues an administrative statement for a table (for example, `OPTIMIZE TABLE` or `REPAIR TABLE`), this could sometimes stop the slave SQL thread. This does not lead to any corruption, but you must use `START SLAVE` to get replication going again. ([Bug#1858](#) [<http://bugs.mysql.com/1858>]) The bug was accidentally not fixed in 4.0.17 as it was unfortunately earlier said.
- Fixed that when a `Rotate` event is found by the slave SQL thread in the middle of a transaction, the value of `Relay_Log_Pos` in `SHOW SLAVE STATUS` remains correct. ([Bug#3017](#) [<http://bugs.mysql.com/3017>])
- Corrected the master's binary log position that `InnoDB` reports when it is doing a crash recovery on a slave server. ([Bug#3015](#) [<http://bugs.mysql.com/3015>])
- Changed that when a `DROP TEMPORARY TABLE` statement is automatically written to the binary

log when a session ends, the statement is recorded with an error code of value zero (this ensures that killing a `SELECT` on the master does not result in a superfluous error on the slave). (Bug#3063 [http://bugs.mysql.com/3063])

- Changed that when a thread handling `INSERT DELAYED` (also known as a `delayed_insert` thread) is killed, its statements are recorded with an error code of value zero (killing such a thread does not endanger replication, so we thus avoid a superfluous error on the slave). (Bug#3081 [http://bugs.mysql.com/3081])
- Fixed deadlock when two `START SLAVE` commands were run at the same time. (Bug#2921 [http://bugs.mysql.com/2921])
- Fixed that a statement never triggers a superfluous error on the slave, if it must be excluded given the `--replicate-*` options. The bug was that if the statement had been killed on the master, the slave would stop. (Bug#2983 [http://bugs.mysql.com/2983])
- The `--local-load` option of `mysqlbinlog` now requires an argument.
- Fixed a segmentation fault when running `LOAD DATA FROM MASTER` after `RESET SLAVE`. (Bug#2922 [http://bugs.mysql.com/2922])
- Fixed a rare error condition that caused the slave SQL thread spuriously to print the message `Bin-log has bad magic number` and stop when it was not necessary to do so. (Bug#3401 [http://bugs.mysql.com/3401])
- Fixed bug in privilege checking of `ALTER TABLE RENAME`. (Bug#3270 [http://bugs.mysql.com/3270], CVE-2004-0835 [http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2004-0835])
- Fixed the column `Exec_master_log_pos` (and its disk image in the `relay-log.info` file) to be correct if the master had version 3.23 (it was too big by 6 bytes). This bug does not exist in the 5.0 version. (Bug#3400 [http://bugs.mysql.com/3400])
- Fixed that `mysqlbinlog` does not forget to print a `USE` command under rare circumstances where the binary log contained a `LOAD DATA INFILE` command. (Bug#3415 [http://bugs.mysql.com/3415])
- Fixed a memory corruption when replicating a `LOAD DATA INFILE` when the master had version 3.23. Some smaller problems remain in this setup, See Section 6.7, “Replication Features and Known Problems”. (Bug#3422 [http://bugs.mysql.com/3422])
- Multiple-table `DELETE` statements were always replicated by the slave if there were some `--replicate-*-ignore-table` options and no `--replicate-*-do-table` options. (Bug#3461 [http://bugs.mysql.com/3461])
- Fixed a crash of the MySQL slave server when it was built with `--with-debug` and replicating itself. (Bug#3568 [http://bugs.mysql.com/3568])

## D.2.11. Changes in release 4.0.18 (12 February 2004)

Functionality added or changed:

- Fixed processing of `LOAD DATA` by `mysqlbinlog` in remote mode. (Bug#1378 [http://bugs.mysql.com/1378])
- The `ft_dump` utility program was renamed to `myisam_ftdump`, and is included in binary distri-

butions.

- `ENGINE` is now a synonym for the `TYPE` option for `CREATE TABLE` and `ALTER TABLE`.
- `lower_case_table_names` system variable now can take a value of 2, to store table names in mixed case on case-insensitive filesystems. It's forced to 2 if the database directory is located on a case-insensitive filesystem.
- For replication of `MEMORY (HEAP)` tables: Made the master automatically write a `DELETE FROM` statement to its binary log when a `MEMORY` table is opened for the first time since master's startup. This is for the case where the slave has replicated a non-empty `MEMORY` table, then the master is shut down and restarted: the table is now empty on master; the `DELETE FROM` empties it on slave too. Note that even with this fix, between the master's restart and the first use of the table on master, the slave still has out-of-date data in the table. But if you use the `init-file` option to populate the `MEMORY` table on the master at startup, it ensures that the failing time interval is zero. (Bug#2477 [http://bugs.mysql.com/2477])
- Optimizer is now better tuned for the case where the first used key part (of many) is a constant. (Bug#1679 [http://bugs.mysql.com/1679])
- Removed old non-working `--old-rpl-compat` server option, which was a holdover from the very first 4.0.x versions. (Bug#2428 [http://bugs.mysql.com/2428])
- Added option `--sync-frm`. It's on by default, to instruct MySQL to sync to disk each time `.frm` file is created. Use `--disable-sync-frm` to disable.

Bugs fixed:

- `mysqlhotcopy` now works on NetWare.
- `DROP DATABASE` could not drop databases with RAID tables that had more than nine `RAID_CHUNKS`. (Bug#2627 [http://bugs.mysql.com/2627])
- Fixed bug in range optimizer when using overlapping ranges. (Bug#2448 [http://bugs.mysql.com/2448])
- Limit `wait_timeout` to 2147483 on Windows (OS limit). (Bug#2400 [http://bugs.mysql.com/2400])
- Fixed bug when `--init-file` crashes MySQL if it contains a large `SELECT`. (Bug#2526 [http://bugs.mysql.com/2526])
- `SHOW KEYS` now shows `NULL` in the `Sub_part` column for `FULLTEXT` indexes.
- The signal thread's stack size was increased to enable `mysqld` to run on Debian/IA-64 with a TLS-enabled `glibc`. (Bug#2599 [http://bugs.mysql.com/2599])
- Now only the `SELECT` privilege is needed for tables that are only read in multiple-table `UPDATE` statements. (Bug#2377 [http://bugs.mysql.com/2377])
- Give proper error message if one uses `LOCK TABLES ... ; INSERT ... SELECT` and one used the same table in the `INSERT` and `SELECT` part. (Bug#2296 [http://bugs.mysql.com/2296])
- `SELECT INTO ... DUMPFILE` now deletes the generated file on error.
- Fixed foreign key reference handling to allow references to column names that contain spaces. (Bug#1725 [http://bugs.mysql.com/1725])

- Fixed problem with index reads on character columns with BDB tables. The symptom was that data could be returned in the wrong lettercase. ([Bug#2509](http://bugs.mysql.com/2509) [<http://bugs.mysql.com/2509>])
- Fixed a spurious table corruption problem that could sometimes appear on tables with indexed TEXT columns if these columns happened to contain values having trailing spaces. This bug was introduced in 4.0.17.
- Fixed a problem where some queries could hang if a condition like `indexed_TEXT_column = expr` was present and the column contained values having trailing spaces. This bug was introduced in 4.0.17.
- Fixed a bug that could cause incorrect results from a query that involved range conditions on indexed TEXT columns that happened to contain values having trailing spaces. This bug was introduced in 4.0.17. ([Bug#2295](http://bugs.mysql.com/2295) [<http://bugs.mysql.com/2295>])
- Fixed incorrect path names in some of the manual pages. ([Bug#2270](http://bugs.mysql.com/2270) [<http://bugs.mysql.com/2270>])
- Fixed spurious “table corrupted” errors in parallel repair operations. See [Section 5.2.2, “System Variables”](#).
- Fixed a crashing bug in parallel repair operations. See [Section 5.2.2, “System Variables”](#).
- Fixed bug in updating MyISAM tables for BLOB values longer than 16MB. ([Bug#2159](http://bugs.mysql.com/2159) [<http://bugs.mysql.com/2159>])
- Fixed bug in `mysqld_safe` when running multiple instances of MySQL. ([Bug#2114](http://bugs.mysql.com/2114) [<http://bugs.mysql.com/2114>])
- Fixed a bug in using `HANDLER` statement with tables not from a current database. ([Bug#2304](http://bugs.mysql.com/2304) [<http://bugs.mysql.com/2304>])
- Fixed a crashing bug that occurred due to the fact that multiple-table `UPDATE` statements did not check that there was only one table to be updated. ([Bug#2103](http://bugs.mysql.com/2103) [<http://bugs.mysql.com/2103>])
- Fixed a crashing bug that occurred due to BLOB data type index size being calculated incorrectly in `MIN( )` and `MAX( )` optimizations. ([Bug#2189](http://bugs.mysql.com/2189) [<http://bugs.mysql.com/2189>])
- Fixed a bug with incorrect syntax for `LOCK TABLES` in `mysqldump`. ([Bug#2242](http://bugs.mysql.com/2242) [<http://bugs.mysql.com/2242>])
- Fixed a bug in `mysqld_safe` that caused `mysqld` to generate a warning about duplicate `user=xxx` options if this option was specified in the `[mysqld]` or `[server]` sections of `my.cnf`. ([Bug#2163](http://bugs.mysql.com/2163) [<http://bugs.mysql.com/2163>])
- `INSERT DELAYED ... SELECT ...` could cause table corruption because tables were not locked properly. This is now fixed by ignoring `DELAYED` in this context. ([Bug#1983](http://bugs.mysql.com/1983) [<http://bugs.mysql.com/1983>])
- Replication: Sometimes the master gets a non-fatal error during the execution of a statement that does not immediately succeed. (For example, a write to a MyISAM table may first receive “no space left on device,” but later complete when disk space becomes available. See [Section A.4.3, “How MySQL Handles a Full Disk”](#).) The bug was that the master forgot to reset the error code to 0 after success, so the error code got into its binary log, thus causing the slave to issue false alarms such as “did not get the same error as on master.” ([Bug#2083](http://bugs.mysql.com/2083) [<http://bugs.mysql.com/2083>])
- Removed a misleading “check permissions on master.info” from a replication error message, because the cause of the problem could be something other than permissions. ([Bug#2121](http://bugs.mysql.com/2121) [<http://bugs.mysql.com/2121>])

- Fixed a crash when the replication slave was unable to create the first relay log. ([Bug#2145](#) [<http://bugs.mysql.com/2145>])
- Replication of `LOAD DATA INFILE` for an empty file from a 3.23 master to a 4.0 slave caused the slave to print an error. ([Bug#2452](#) [<http://bugs.mysql.com/2452>])
- When automatically forcing `lower_case_table_names` to 1 if the file system was case insensitive, `mysqld` could crash. This bug existed only in MySQL 4.0.17. ([Bug#2481](#) [<http://bugs.mysql.com/2481>])
- Restored ability to specify default values for `TIMESTAMP` columns that was erroneously disabled in previous release. ([Bug#2539](#) [<http://bugs.mysql.com/2539>]) Fixed `SHOW CREATE TABLE` to reflect these values. ([Bug#1885](#) [<http://bugs.mysql.com/1885>]) Note that because of the auto-update feature for the first `TIMESTAMP` column in a table, it makes no sense to specify a default value for the column. Any such default is silently ignored (unless another `TIMESTAMP` column is added before this one). Also fixed the meaning of the `DEFAULT` keyword when it is used to specify the value to be inserted into a `TIMESTAMP` column other than the first. ([Bug#2464](#) [<http://bugs.mysql.com/2464>])
- Fixed bug for out-of-range arguments on QNX platform that caused `UNIX_TIMESTAMP( )` to produce incorrect results or that caused non-zero values to be inserted into `TIMESTAMP` columns. ([Bug#2523](#) [<http://bugs.mysql.com/2523>]) Also, current time zone now is taken into account when checking if datetime values satisfy both range boundaries for `TIMESTAMP` columns. The range allowed for a `TIMESTAMP` column is time zone-dependent and equivalent to a range of `1970-01-01 00:00:01 UTC to 2037-12-31 23:59:59 UTC`.
- Multiple-table `DELETE` statements were never replicated by the slave if there were any `-replicate-*-table` options. ([Bug#2527](#) [<http://bugs.mysql.com/2527>])
- Changes to session counterparts of variables `query_prealloc_size`, `query_alloc_block_size`, `trans_prealloc_size`, `trans_alloc_block_size` now have an effect. ([Bug#1948](#) [<http://bugs.mysql.com/1948>])
- Fixed bug in `ALTER TABLE RENAME`, when rename to the table with the same name in another database silently dropped destination table if it existed. ([Bug#2628](#) [<http://bugs.mysql.com/2628>])

## D.2.12. Changes in release 4.0.17 (14 December 2003)

Functionality added or changed:

- `mysqldump` no longer dumps data for `MERGE` tables. ([Bug#1846](#) [<http://bugs.mysql.com/1846>])
- `lower_case_table_names` is now forced to 1 if the database directory is located on a case-insensitive filesystem. ([Bug#1812](#) [<http://bugs.mysql.com/1812>])
- Symlink creation is now disabled on systems where `realpath( )` doesn't work. (Before one could use `CREATE TABLE .. DATA DIRECTORY=..` even if `HAVE_BROKEN_REALPATH` was defined. This is now disabled to avoid problems when running `ALTER TABLE`).
- Inserting a negative `AUTO_INCREMENT` value in a `MyISAM` table no longer updates the `AUTO_INCREMENT` counter to a big unsigned value. ([Bug#1366](#) [<http://bugs.mysql.com/1366>])
- Added four new modes to `WEEK( ..., mode)` function. See `WEEK(date: (mode))`. ([Bug#1178](#) [<http://bugs.mysql.com/1178>])
- Allow `UNION DISTINCT` syntax.

- MySQL now syncs to disk each time `.frm` file is created.
- `mysql_server_init()` now returns 1 if it can't initialize the environment. (Previously `mysql_server_init()` called `exit(1)` if it could not create a key with `pthread_key_create()`. (Bug#2062 [<http://bugs.mysql.com/2062>]))
- Allow spaces in Windows service names.
- Changed the default Windows service name for `mysqld` from `MySql` to `MySQL`. This should not affect usage, because service names are not case sensitive.
- When you install `mysqld` as a service on Windows systems, `mysqld` reads startup options in option files from the option group with the same name as the service name. (Except when the service name is `MySQL`).

### Bugs fixed:

- Sending `SIGHUP` to `mysqld` crashed the server if it was running with `--log-bin`. (Bug#2045 [<http://bugs.mysql.com/2045>])
- One can now configure MySQL as a Windows service as a normal user. (Bug#1802 [<http://bugs.mysql.com/1802>]). Thanks to Richard Hansen for fixing this.
- Database names are now compared in lowercase in `ON` clauses when `lower_case_table_names` is set. (Bug#1736 [<http://bugs.mysql.com/1736>])
- `IGNORE ... LINES` option to `LOAD DATA INFILE` didn't work when used with fixed length rows. (Bug#1704 [<http://bugs.mysql.com/1704>])
- Fixed problem with `UNIX_TIMESTAMP()` for timestamps close to 0. (Bug#1998 [<http://bugs.mysql.com/1998>])
- Fixed problem with character values greater than 128 in the `QUOTE()` function. (Bug#1868 [<http://bugs.mysql.com/1868>])
- Fixed searching of `TEXT` with endspace. (Bug#1651 [<http://bugs.mysql.com/1651>])
- Fixed caching bug in multiple-table updates where same table was used twice. (Bug#1711 [<http://bugs.mysql.com/1711>])
- Fixed directory permissions for the MySQL-server RPM documentation directory. (Bug#1672 [<http://bugs.mysql.com/1672>])
- Fixed server crash when updating an `ENUM` column that is set to the empty string (for example, with `REPLACE()`). (Bug#2023 [<http://bugs.mysql.com/2023>])
- `mysql` client program now correctly prints connection identifier returned by `mysql_thread_id()` as unsigned integer rather than as signed integer. (Bug#1951 [<http://bugs.mysql.com/1951>])
- `FOUND_ROWS()` could return incorrect number of rows after a query with an impossible `WHERE` condition. (Bug#1468 [<http://bugs.mysql.com/1468>])
- `SHOW DATABASES` no longer shows `.sym` files (on Windows) that do not point to a valid directory. (Bug#1385 [<http://bugs.mysql.com/1385>])
- Fixed a possible memory leak on Mac OS X when using the shared `libmysql.so` library. (from

`pthread_key_create()`. (Bug#2061 [http://bugs.mysql.com/2061])

- Fixed bug in `UNION` statement with alias `*`. (Bug#1249 [http://bugs.mysql.com/1249])
- Fixed a bug in `DELETE ... ORDER BY ... LIMIT` where the rows were not deleted in the proper order. (Bug#1024 [http://bugs.mysql.com/1024], Bug#1697 [http://bugs.mysql.com/1697]).
- Fixed serious problem with multi-threaded programs on Windows that used the embedded MySQL libraries. (Locks of tables were not handled correctly between different threads).
- Code cleanup: Fixed a few code defects (potential memory leaks, null pointer dereferences, uninitialized variables). Thanks to Reasoning Inc. for informing us about these findings.
- Fixed a buffer overflow error that occurred with prepended '0' characters in some columns of type `DECIMAL`. (Bug#2128 [http://bugs.mysql.com/2128])
- Filesort was never shown in `EXPLAIN` if query contained an `ORDER BY NULL` clause. (Bug#1335 [http://bugs.mysql.com/1335])
- Fixed invalidation of whole query cache on `DROP DATABASE`. (Bug#1898 [http://bugs.mysql.com/1898])
- Fixed bug in range optimizer that caused wrong results for some unlikely `AND/OR` queries. (Bug#1828 [http://bugs.mysql.com/1828])
- Fixed a crash in `ORDER BY` when ordering by expression and identifier. (Bug#1945 [http://bugs.mysql.com/1945])
- Fixed a crash in an open `HANDLER` when an `ALTER TABLE` was executed in a different connection. (Bug#1826 [http://bugs.mysql.com/1826])
- Fixed a bug in `trunc*` operator of full-text search which sometimes caused MySQL not to find all matched rows.
- Fixed bug in prepending '0' characters to `DECIMAL` column values.
- Fixed optimizer bug, introduced in 4.0.16, when `REF` access plan was preferred to more efficient `RANGE` on another column.
- Fixed problem when installing a MySQL server as a Windows service using a command of the form `mysqld --install mysql --defaults-file=path-to-file`. (Bug#1643 [http://bugs.mysql.com/1643])
- Fixed an incorrect result from a query that uses only `const` tables (such as one-row tables) and non-constant expression (such as `RAND()`). (Bug#1271 [http://bugs.mysql.com/1271])
- Fixed bug when the optimizer did not take `SQL_CALC_FOUND_ROWS` into account if `LIMIT` clause was present. (Bug#1274 [http://bugs.mysql.com/1274])
- `mysqlbinlog` now asks for a password at the console when the `-p` or `--password` option is used with no argument. This is consistent with the way that other clients such `mysqladmin` and `mysqldump` behave. **Note:** A consequence of this change is that it is no longer possible to invoke `mysqlbinlog` as `mysqlbinlog -p pass_val` (with a space between the `-p` option and the following password value). (Bug#1595 [http://bugs.mysql.com/1595])
- Fixed bug accidentally introduced in 4.0.16 where the slave SQL thread deleted its replicated temporary tables when `STOP SLAVE` was issued.
- In a "chain" replication setup `A->B->C`, if 2 sessions on A updated temporary tables of the same

name at the same time, the binary log of **B** became incorrect, resulting in **C** becoming confused. ([Bug#1686](http://bugs.mysql.com/1686) [<http://bugs.mysql.com/1686>])

- In a “chain” replication setup **A**->**B**->**C**, if `STOP SLAVE` was issued on **B** while it was replicating a temporary table from **A**, then when `START SLAVE` was issued on **B**, the binary log of **B** became incorrect, resulting in **C** becoming confused. ([Bug#1240](http://bugs.mysql.com/1240) [<http://bugs.mysql.com/1240>])
- When `MASTER_LOG_FILE` and `MASTER_LOG_POS` were not specified, `CHANGE MASTER` used the coordinates of the slave I/O thread to set up replication, which broke replication if the slave SQL thread lagged behind the slave I/O thread. This caused the slave SQL thread to lose some events. The new behavior is to use the coordinates of the slave SQL thread instead. See [Section 13.6.2.1](#), “`CHANGE MASTER TO Syntax`”. ([Bug#1870](http://bugs.mysql.com/1870) [<http://bugs.mysql.com/1870>])
- Now if integer is stored or converted to `TIMESTAMP` or `DATETIME` value checks of year, month, day, hour, minute and second ranges are performed and numbers representing illegal timestamps are converted to 0 value. This behavior is consistent with manual and with behavior of string to `TIMESTAMP/DATETIME` conversion. ([Bug#1448](http://bugs.mysql.com/1448) [<http://bugs.mysql.com/1448>])
- Fixed bug when `BIT_AND( )` and `BIT_OR( )` group functions returned incorrect value if `SELECT` used a temporary table and no rows were found. ([Bug#1790](http://bugs.mysql.com/1790) [<http://bugs.mysql.com/1790>]).
- `BIT_AND( )` is now unsigned in all contexts. This means that it now returns 18446744073709551615 (= 0xffffffffffff) instead of -1 if there were no rows in the result.
- Fixed bug with `BIT_AND( )` still returning signed value for an empty set in some cases. ([Bug#1972](http://bugs.mysql.com/1972) [<http://bugs.mysql.com/1972>])
- Fixed bug with `^` (XOR) and `>>` (bit shift) still returning signed value in some cases. ([Bug#1993](http://bugs.mysql.com/1993) [<http://bugs.mysql.com/1993>])
- Replication: a rare race condition in the slave SQL thread, which could lead to a wrong complain that the relay log is corrupted. ([Bug#2011](http://bugs.mysql.com/2011) [<http://bugs.mysql.com/2011>])
- Replication: in the slave SQL thread, a multiple-table `UPDATE` could produce a wrong complain that some record was not found in one table, if the `UPDATE` was preceded by a `INSERT . . . SELECT`. ([Bug#1701](http://bugs.mysql.com/1701) [<http://bugs.mysql.com/1701>])
- Fixed deficiency in MySQL code which is responsible for scanning directories. This deficiency caused `SHOW TABLE STATUS` to be very slow when a database contained a large number of tables, even if a single particular table were specified. ([Bug#1952](http://bugs.mysql.com/1952) [<http://bugs.mysql.com/1952>])

## D.2.13. Changes in release 4.0.16 (17 October 2003)

Functionality added or changed:

- Option values in option files now may be quoted. This is useful for values that contain whitespace or comment characters.
- Write memory allocation information to error log when doing `mysqladmin debug`. This works only on systems that support the `mallinfo( )` call (like newer Linux systems).
- Added the following new system variables to allow more precise memory allocation: `range_alloc_block_size`, `query_alloc_block_size`, `query_prealloc_size`, `transaction_alloc_block_size`, and `transaction_prealloc_size`.
- `mysqlbinlog` now reads option files. To make this work, you must now specify -



`-read-from-remote-server` when reading binary logs from a MySQL server. (Note that using a remote server is deprecated and may disappear in future `mysqlbinlog` versions).

- Block `SIGPIPE` signals also for non-threaded programs. The blocking is moved from `mysql_init()` to `mysql_server_init()`, which is automatically called on the first call to `mysql_init()`.
- Added `--libs_r` and `--include` options to `mysql_config`.
- New ``>` prompt for `mysql`. This prompt is similar to the `'>` and `">` prompts, but indicates that an identifier quoted with backticks was begun on an earlier line and the closing backtick has not yet been seen.
- Updated `mysql_install_db` to be able to use the local machine's IP address instead of the host-name when building the initial grant tables if `skip-name-resolve` has been specified. This option can be helpful on FreeBSD to avoid thread-safety problems with the FreeBSD resolver libraries. (Thanks to Jeremy Zawodny for the patch.)
- A documentation change: Added a note that when backing up a slave, it is necessary also to back up the `master.info` and `relay-log.info` files, as well as any `SQL_LOAD-*` files located in the directory specified by the `--slave-load-tmpdir` option. All these files are needed when the slave resumes replication after you restore the slave's data.

Bugs fixed:

- Fixed a spurious error `ERROR 14: Can't change size of file (Errcode: 2)` on Windows in `DELETE FROM tbl_name` without a `WHERE` clause or `TRUNCATE TABLE tbl_name`, when `tbl_name` is a `MyISAM` table. (Bug#1397 [http://bugs.mysql.com/1397])
- Fixed a bug that resulted in `thr_alarm queue is full` warnings after increasing the `max_connections` variable with `SET GLOBAL`. (Bug#1435 [http://bugs.mysql.com/1435])
- Made `LOCK TABLES` to work when `Lock_tables_priv` is granted on the database level and `Select_priv` is granted on the table level.
- Fixed crash of `FLUSH QUERY CACHE` on queries that use same table several times (Bug#988 [http://bugs.mysql.com/988]).
- Fixed core dump bug when setting an enum system variable (such as `SQL_WARNINGS`) to `NULL`.
- Extended the default timeout value for Windows clients from 30 seconds to 1 year. (The timeout that was added in MySQL 4.0.15 was way too short). This fixes a bug that caused `ERROR 2013: Lost connection to MySQL server during query` for queries that lasted longer than 30 seconds, if the client didn't specify a limit with `mysql_options()`. Users of 4.0.15 on Windows should upgrade to avoid this problem.
- More “out of memory” checking in range optimizer.
- Fixed and documented a problem when setting and using a user variable within the same `SELECT` statement. (Bug#1194 [http://bugs.mysql.com/1194]).
- Fixed bug in overrun check for `BLOB` values with compressed tables. This was a bug introduced in 4.0.14. It caused MySQL to regard some correct tables containing `BLOB` values as corrupted. (Bug#770 [http://bugs.mysql.com/770], Bug#1304 [http://bugs.mysql.com/1304], and maybe Bug#1295 [http://bugs.mysql.com/1295])
- `SHOW GRANTS` showed `USAGE` instead of the real column-level privileges when no table-level

privileges were given.

- When copying a database from the master, `LOAD DATA FROM MASTER` dropped the corresponding database on the slave, thus erroneously dropping tables that had no counterpart on the master and tables that may have been excluded from replication using `--replicate-*-table` rules. Now `LOAD DATA FROM MASTER` no longer drops the database. Instead, it drops only the tables that have a counterpart on the master and that match the `--replicate-*-table` rules. `--replicate-*-db` rules can still be used to include or exclude a database as a whole from `LOAD DATA FROM MASTER`. A database also is included or excluded as a whole if there are some rules like `--replicate-wild-do-table=db1.%` or `--replicate-wild-ignore-table=db1.%`, as is the case for `CREATE DATABASE` and `DROP DATABASE` in replication. ([Bug#1248](http://bugs.mysql.com/1248) [<http://bugs.mysql.com/1248>])
- Fixed a bug where `mysqlbinlog` crashed with a segmentation fault when used with the `-h` or `-host` option. ([Bug#1258](http://bugs.mysql.com/1258) [<http://bugs.mysql.com/1258>])
- Fixed a bug where `mysqlbinlog` crashed with a segmentation fault when used on a binary log containing only final events for `LOAD DATA`. ([Bug#1340](http://bugs.mysql.com/1340) [<http://bugs.mysql.com/1340>])
- `mysqlbinlog` does not reuse temporary filenames from previous runs. Previously `mysqlbinlog` failed if was used several times on the same binary log file that contained a `LOAD DATA` command.
- Fixed compilation problem when compiling with OpenSSL 0.9.7 with disabled old DES support (If `OPENSSL_DISABLE_OLD_DES_SUPPORT` option was enabled).
- Fixed a bug when two (or more) MySQL servers were running on the same machine, and they were both slaves, and at least one of them was replicating some `LOAD DATA INFILE` command from its master. The bug was that one slave MySQL server sometimes deleted the `SQL_LOAD-*` files (used for replication of `LOAD DATA INFILE` and located in the `slave-load-tmpdir` directory, which defaults to `tmpdir`) belonging to the other slave MySQL server of this machine, if these slaves had the same `slave-load-tmpdir` directory. When that happened, the other slave could not replicate `LOAD DATA INFILE` and complained about not being able to open some `SQL_LOAD-*` file. ([Bug#1357](http://bugs.mysql.com/1357) [<http://bugs.mysql.com/1357>])
- If `LOAD DATA INFILE` failed for a small file, the master forgot to write a marker (a `Delete_file` event) in its binary log, so the slave could not delete 2 files (`SQL_LOAD-*.info` and `SQL_LOAD-*.data` from its `tmpdir`). ([Bug#1391](http://bugs.mysql.com/1391) [<http://bugs.mysql.com/1391>])
- On Windows, the slave forgot to delete a `SQL_LOAD-*.info` file from `tmpdir` after successfully replicating a `LOAD DATA INFILE` command. ([Bug#1392](http://bugs.mysql.com/1392) [<http://bugs.mysql.com/1392>])
- When a connection terminates, MySQL writes `DROP TEMPORARY TABLE` statements to the binary log for all temporary tables which the connection had not explicitly dropped. MySQL forgot to use backticks to quote the database and table names in the statement. ([Bug#1345](http://bugs.mysql.com/1345) [<http://bugs.mysql.com/1345>])
- On some 64-bit machines (some HP-UX and Solaris machines), a slave installed with the 64-bit MySQL binary could not connect to its master (it connected to itself instead). ([Bug#1256](http://bugs.mysql.com/1256) [<http://bugs.mysql.com/1256>], [Bug#1381](http://bugs.mysql.com/1381) [<http://bugs.mysql.com/1381>])
- Code was introduced in MySQL 4.0.15 for the slave to detect that the master had died while writing a transaction to its binary log. This code reported an error in a legal situation: When the slave I/O thread was stopped while copying a transaction to the relay log, the slave SQL thread would later pretend that it found an unfinished transaction. ([Bug#1475](http://bugs.mysql.com/1475) [<http://bugs.mysql.com/1475>])

## D.2.14. Changes in release 4.0.15 (03 September 2003)

---

**IMPORTANT:**

If you are using this release on Windows, you should upgrade at least your clients (any program that uses `libmysql.lib`) to 4.0.16 or above. This is because the 4.0.15 release had a bug in the Windows client library that causes Windows clients using the library to die with a `Lost connection to MySQL server during query` error for queries that take more than 30 seconds. This problem is specific to Windows; clients on other platforms are unaffected.

Functionality added or changed:

- `mysqldump` now correctly quotes all identifiers when communicating with the server. This assures that during the dump process, `mysqldump` never sends queries to the server that result in a syntax error. This problem is **not** related to the `mysqldump` program's output, which was not changed. (Bug#1148 [<http://bugs.mysql.com/1148>])
- Change result set metadata information so that `MIN()` and `MAX()` report that they can return `NULL` (this is true because an empty set returns `NULL`). (Bug#324 [<http://bugs.mysql.com/324>])
- Produce an error message on Windows if a second `mysqld` server is started on the same TCP/IP port as a running `mysqld` server.
- The `mysqld` system variables `wait_timeout`, `net_read_timeout`, and `net_write_timeout` now work on Windows. One can now also set timeouts for read and writes in Windows clients with `mysql_options()`.
- Added option `--sql-mode=NO_DIR_IN_CREATE` to make it possible for slaves to ignore `INDEX DIRECTORY` and `DATA DIRECTORY` options given to `CREATE TABLE`. When this is mode is on, `SHOW CREATE TABLE` does not show the given directories.
- `SHOW CREATE TABLE` now shows the `INDEX DIRECTORY` and `DATA DIRECTORY` options, if they were specified when the table was created.
- The `open_files_limit` system variable now shows the real open files limit.
- `MATCH ... AGAINST()` in natural language mode now treats words that are present in more than 2,000,000 rows as stopwords.
- The Mac OS X installation disk images now include an additional `MySQLStartupItem.pkg` package that enables the automatic startup of MySQL on system startup. See Section 2.5, "Installing MySQL on Mac OS X".
- Most of the documentation included in the binary tarball distributions (`.tar.gz`) has been moved into a subdirectory `docs`. See Section 2.1.5, "Installation Layouts".
- The manual is now included as an additional `info` file in the binary distributions. (Bug#1019 [<http://bugs.mysql.com/1019>])
- The binary distributions now include the embedded server library (`libmysqld.a`) by default. Due to a linking problem with non-gcc compilers, it was not included in all packages of the initial 4.0.15 release. The affected packages were rebuilt and released as 4.0.15a. See Section 1.6.1.2, "The Embedded MySQL Server".
- MySQL can now use range optimization for `BETWEEN` with non-constant limits. (Bug#991 [<http://bugs.mysql.com/991>])
- Replication error messages now include the default database, so that users can check which database the failing query was run for.

- A documentation change: Added a paragraph about how the `binlog-do-db` and `binlog-ignore-db` options are tested against the database on the master (see [Section 5.11.4](#), “The Binary Log”), and a paragraph about how `--replicate-do-db`, `--replicate-do-table` and analogous options are tested against the database and tables on the slave (see [Section 6.8](#), “Replication Startup Options”).
- Now the slave does not replicate `SET PASSWORD` if it is configured to exclude the `mysql` database from replication (using for example `--replicate-wild-ignore-table=mysql.%`). This was the case for `GRANT` and `REVOKE` since version 4.0.13 (although there was [Bug#980](#) [<http://bugs.mysql.com/980>] in 4.0.13 & 4.0.14, which has been fixed in 4.0.15).
- Rewrote the information shown in the `State` column of `SHOW PROCESSLIST` for replication threads and for `MASTER_POS_WAIT()` and added the most common states for these threads to the documentation, see [Section 6.3](#), “Replication Implementation Details”.
- Added a test in replication to detect the case where the master died in the middle of writing a transaction to the binary log; such unfinished transactions now trigger an error message on the slave.
- A `GRANT` command that creates an anonymous user (that is, an account with an empty username) no longer requires `FLUSH PRIVILEGES` for the account to be recognized by the server. ([Bug#473](#) [<http://bugs.mysql.com/473>])
- `CHANGE MASTER` now flushes `relay-log.info`. Previously this was deferred to the next run of `START SLAVE`, so if `mysqld` was shutdown on the slave after `CHANGE MASTER` without having run `START SLAVE`, the relay log’s name and position were lost. At restart they were reloaded from `relay-log.info`, thus reverting to their old (incorrect) values from before `CHANGE MASTER` and leading to error messages (as the old relay log did not exist any more) and the slave threads refusing to start. ([Bug#858](#) [<http://bugs.mysql.com/858>])

Bugs fixed:

- Fixed buffer overflow in password handling which could potentially be exploited by MySQL users with `ALTER` privilege on the `mysql.user` table to execute random code or to gain shell access with the UID of the `mysqld` process (thanks to Jedi/Sector One for spotting and reporting this bug).
- Fixed server crash on `FORCE INDEX` in a query that contained “Range checked for each record” in the `EXPLAIN` output. ([Bug#1172](#) [<http://bugs.mysql.com/1172>])
- Fixed table/column grant handling: The proper sort order (from most specific to less specific, see [Section 5.7.6](#), “Access Control, Stage 2: Request Verification”) was not honored. ([Bug#928](#) [<http://bugs.mysql.com/928>])
- Fixed rare bug in `MYISAM` introduced in 4.0.3 where the index file header was not updated directly after an `UPDATE` of split dynamic rows. The symptom was that the table had a corrupted delete-link if `mysqld` was shut down or the table was checked directly after the update.
- Fixed `Can't unlock file` error when running `myisamchk --sort-index` on Windows. ([Bug#1119](#) [<http://bugs.mysql.com/1119>])
- Fixed possible deadlock when changing `key_buffer_size` while the key cache was actively used. ([Bug#1088](#) [<http://bugs.mysql.com/1088>])
- Fixed overflow bug in `MyISAM` and `ISAM` when a row is updated in a table with a large number of columns and at least one `BLOB/TEXT` column.
- Fixed incorrect result when doing `UNION` and `LIMIT #, #` when braces were not used around the `SELECT` parts.

- Fixed incorrect result when doing `UNION` and `ORDER BY .. LIMIT #` when one didn't use braces around the `SELECT` parts.
- Fixed problem with `SELECT SQL_CALC_FOUND_ROWS ... UNION ALL ... LIMIT #` where `FOUND_ROWS()` returned incorrect number of rows.
- Fixed unlikely stack bug when having a BIG expression of type `1+1-1+1-1...` in certain combinations. (Bug#871 [<http://bugs.mysql.com/871>])
- Fixed the bug that sometimes prevented a table with a `FULLTEXT` index from being marked as "analyzed".
- Fixed MySQL so that the column length (in C API) for the second column in `SHOW CREATE TABLE` is always larger than the data length. The only known application that was affected by the old behavior was Borland dbExpress, which truncated the output from the command. (Bug#1064 [<http://bugs.mysql.com/1064>])
- Fixed crash in comparisons of strings using the `tis620` character set. (Bug#1116 [<http://bugs.mysql.com/1116>])
- Fixed `ISAM` bug in `MAX()` optimization.
- `myisamchk --sort-records=N` no longer marks table as crashed if sorting failed because of an inappropriate key. (Bug#892 [<http://bugs.mysql.com/892>])
- Fixed a minor bug in `MyISAM` compressed table handling that sometimes made it impossible to repair compressed table in "Repair by sort" mode. "Repair with keycache" (`myisamchk -safe-recover`) worked, though. (Bug#1015 [<http://bugs.mysql.com/1015>])
- Fixed bug in propagating the version number to the manual included in the distribution files. (Bug#1020 [<http://bugs.mysql.com/1020>])
- Fixed key sorting problem (a `PRIMARY` key declared for a column that is not explicitly marked `NOT NULL` was sorted after a `UNIQUE` key for a `NOT NULL` column).
- Fixed the result of `INTERVAL` when applied to a `DATE` value. (Bug#792 [<http://bugs.mysql.com/792>])
- Fixed compiling of the embedded server library in the RPM spec file. (Bug#959 [<http://bugs.mysql.com/959>])
- Added some missing files to the RPM spec file and fixed some RPM building errors that occurred on Red Hat Linux 9. (Bug#998 [<http://bugs.mysql.com/998>])
- Fixed incorrect `XOR` evaluation in `WHERE` clause. (Bug#992 [<http://bugs.mysql.com/992>])
- Fixed bug with processing in query cache merged tables constructed from more than 255 tables. (Bug#930 [<http://bugs.mysql.com/930>])
- Fixed incorrect results from outer join query (for example, `LEFT JOIN`) when `ON` condition is always false, and range search in used. (Bug#926 [<http://bugs.mysql.com/926>])
- Fixed a bug causing incorrect results from `MATCH ... AGAINST()` in some joins. (Bug#942 [<http://bugs.mysql.com/942>])
- `MERGE` tables do not ignore `Using index` (from `EXPLAIN` output) anymore.
- Fixed a bug that prevented an empty table from being marked as "analyzed". (Bug#937 [<http://bugs.mysql.com/937>])

- Fixed `myisamchk --sort-records` crash when used on compressed table.
- Fixed slow (as compared to 3.23) `ALTER TABLE` and related commands such as `CREATE INDEX`. (Bug#712 [<http://bugs.mysql.com/712>])
- Fixed segmentation fault resulting from `LOAD DATA FROM MASTER` when the master was running without the `--log-bin` option. (Bug#934 [<http://bugs.mysql.com/934>])
- Fixed a security bug: A server compiled without SSL support still allowed connections by users who had the `REQUIRE SSL` option specified for their accounts.
- Fixed a random bug: Sometimes the slave would replicate `GRANT` or `REVOKE` queries even if it was configured to exclude the `mysql` database from replication (for example, using `--replicate-wild-ignore-table=mysql.%`). (Bug#980 [<http://bugs.mysql.com/980>])
- The `Last_Errno` and `Last_Error` fields in the output of `SHOW SLAVE STATUS` are now cleared by `CHANGE MASTER` and when the slave SQL thread starts. (Bug#986 [<http://bugs.mysql.com/986>])
- A documentation mistake: It said that `RESET SLAVE` does not change connection information (master host, port, user, and password), whereas it does. The statement resets these to the startup options (`master-host` etc) if there were some. (Bug#985 [<http://bugs.mysql.com/985>])
- `SHOW SLAVE STATUS` now shows correct information (master host, port, user, and password) after `RESET SLAVE` (that is, it shows the new values, which are copied from the startup options if there were some). (Bug#985 [<http://bugs.mysql.com/985>])
- Disabled propagation of the original master's log position for events because this caused unexpected values for `Exec_Master_Log_Pos` and problems with `MASTER_POS_WAIT()` in A->B->C replication setup. (Bug#1086 [<http://bugs.mysql.com/1086>])
- Fixed a segmentation fault in `mysqlbinlog` when `--position=x` was used with `x` being between a `Create_file` event and its fellow `Append_block`, `Exec_load` or `Delete_file` events. (Bug#1091 [<http://bugs.mysql.com/1091>])
- `mysqlbinlog` printed superfluous warnings when using `--database`, which caused syntax errors when piped to `mysql`. (Bug#1092 [<http://bugs.mysql.com/1092>])
- Made `mysqlbinlog --database` filter `LOAD DATA INFILE` too (previously, it filtered all queries except `LOAD DATA INFILE`). (Bug#1093 [<http://bugs.mysql.com/1093>])
- `mysqlbinlog` in some cases forgot to put a leading '#' in front of the original `LOAD DATA INFILE` (this command is displayed only for information, not to be run; it is later reworked to `LOAD DATA LOCAL` with a different filename, for execution by `mysql`). (Bug#1096 [<http://bugs.mysql.com/1096>])
- `binlog-do-db` and `binlog-ignore-db` incorrectly filtered `LOAD DATA INFILE` (it was half-written to the binary log). This resulted in a corrupted binary log, which could cause the slave to stop with an error. (Bug#1100 [<http://bugs.mysql.com/1100>])
- When, in a transaction, a transactional table (such as an `InnoDB` table) was updated, and later in the same transaction a non-transactional table (such as a `MyISAM` table) was updated using the updated content of the transactional table (with `INSERT ... SELECT` for example), the queries were written to the binary log in an incorrect order. (Bug#873 [<http://bugs.mysql.com/873>])
- When, in a transaction, `INSERT ... SELECT` updated a non-transactional table, and `ROLLBACK` was issued, no error was returned to the client. Now the client is warned that some changes could not be rolled back, as this was the case for normal `INSERT`. (Bug#1113 [<http://bugs.mysql.com/1113>])

- Fixed a potential bug: When `STOP SLAVE` was run while the slave SQL thread was in the middle of a transaction, and then `CHANGE MASTER` was used to point the slave to some non-transactional statement, the slave SQL thread could get confused (because it would still think, from the past, that it was in a transaction).

## D.2.15. Changes in release 4.0.14 (18 July 2003)

Functionality added or changed:

- Added `default_week_format` system variable. The value is used as the default mode for the `WEEK()` function.
- `mysqld` now reads an additional option file group having a name corresponding to the server's release series: `[mysqld-4.0]` for 4.0.x servers, `[mysqld-4.1]` for 4.1.x servers, and so forth. This allows options to be specified on a series-specific basis.
- The `CONCAT_WS()` function no longer skips empty strings. (Bug#586 [<http://bugs.mysql.com/586>]).
- InnoDB now supports indexing a prefix of a column. This means, in particular, that `BLOB` and `TEXT` columns can be indexed in InnoDB tables, which was not possible before.
- A documentation change: Function `INTERVAL(NULL, ...)` returns `-1`.
- Enabled `INSERT` from `SELECT` when the table into which the records are inserted is also a table listed in the `SELECT`.
- Allow `CREATE TABLE` and `INSERT` from any `UNION`.
- The `SQL_CALC_FOUND_ROWS` option now always returns the total number of rows for any `UNION`.
- Removed `--table` option from `mysqlbinlog` to avoid repeating `mysqldump` functionality.
- Comment lines in option files can now start from the middle of a line, too (like `basedir=c:\mysql # installation directory`).
- Changed optimizer slightly to prefer index lookups over full table scans in some boundary cases.
- Added thread-specific `max_seeks_for_key` variable that can be used to force the optimizer to use keys instead of table scans even if the cardinality of the index is low.
- Added optimization that converts `LEFT JOIN` to normal join in some cases.
- A documentation change: added a paragraph about failover in replication (how to use a surviving slave as the new master, how to resume to the original setup). See Section 6.10, “Replication FAQ”.
- A documentation change: added warning notes about safe use of the `CHANGE MASTER` command. See Section 13.6.2.1, “CHANGE MASTER TO Syntax”.
- MySQL now issues a warning (not an error, as in 4.0.13) when it opens a table that was created with MySQL 4.1.
- Added `--nice` option to `mysqld_safe` to allow setting the niceness of the `mysqld` process. (Thanks to Christian Hammers for providing the initial patch.) (Bug#627 [<http://bugs.mysql.com/627>])

- Added `--read-only` option to cause `mysqld` to allow no updates except from slave threads or from users with the `SUPER` privilege. (Original patch from Markus Benning).
- `SHOW BINLOG EVENTS FROM x` where `x` is less than 4 now silently converts `x` to 4 instead of printing an error. The same change was done for `CHANGE MASTER TO MASTER_LOG_POS=x` and `CHANGE MASTER TO RELAY_LOG_POS=x`.
- `mysqld` now only adds an interrupt handler for the `SIGINT` signal if you start it with the new `-gdb` option. This is done because some MySQL users encountered strange problems when they accidentally sent `SIGINT` to `mysqld` threads.
- `RESET SLAVE` now clears the `Last_Errno` and `Last_Error` fields in the output of `SHOW SLAVE STATUS`.
- Added `max_relay_log_size` variable; the relay log is rotated automatically when its size exceeds `max_relay_log_size`. But if `max_relay_log_size` is 0 (the default), `max_binlog_size` is used (as in older versions). `max_binlog_size` still applies to binary logs in any case.
- `FLUSH LOGS` now rotates relay logs in addition to the other types of logs it rotates.

Bugs fixed:

- Comparison/sorting for `latin1_de` character set was rewritten. The old algorithm could not handle cases like `"sä" > "ßa"`. See Section 5.10.1.1, “Using the German Character Set”. In rare cases it resulted in table corruption.
- Fixed a problem with the password prompt on Windows. (Bug#683 [<http://bugs.mysql.com/683>])
- `ALTER TABLE ... UNION=(...)` for `MERGE` table is now allowed even if some underlying `MyISAM` tables are read-only. (Bug#702 [<http://bugs.mysql.com/702>])
- Fixed a problem with `CREATE TABLE t1 SELECT x'41'`. (Bug#801 [<http://bugs.mysql.com/801>])
- Removed some incorrect lock warnings from the error log.
- Fixed memory overrun when doing `REPAIR TABLE` on a table with a multiple-part auto\_increment key where one part was a packed `CHAR`.
- Fixed a probable race condition in the replication code that could potentially lead to `INSERT` statements not being replicated in the event of a `FLUSH LOGS` command or when the binary log exceeds `max_binlog_size`. (Bug#791 [<http://bugs.mysql.com/791>])
- Fixed a crashing bug in `INTERVAL` and `GROUP BY` or `DISTINCT`. (Bug#807 [<http://bugs.mysql.com/807>])
- Fixed bug in `mysqlhotcopy` so it actually aborts for unsuccessful table copying operations. Fixed another bug so that it succeeds when there are thousands of tables to copy. (Bug#812 [<http://bugs.mysql.com/812>])
- Fixed problem with `mysqlhotcopy` failing to read options from option files. (Bug#808 [<http://bugs.mysql.com/808>])
- Fixed bugs in optimizer that sometimes prevented MySQL from using `FULLTEXT` indexes even though it was possible (for example, in `SELECT * FROM t1 WHERE MATCH a,b AGAINST("index") > 0`).



- Fixed a bug with “table is full” in `UNION` operations.
- Fixed a security problem that enabled users with no privileges to obtain information on the list of existing databases by using `SHOW TABLES` and similar commands.
- Fixed a stack problem on UnixWare/OpenUnix.
- Fixed a configuration problem on UnixWare/OpenUNIX and OpenServer.
- Fixed a stack overflow problem in password verification. (CVE-2003-0780 [<http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2003-0780>])
- Fixed a problem with `max_user_connections`.
- `HANDLER` without an index now works properly when a table has deleted rows. (Bug#787 [<http://bugs.mysql.com/787>])
- Fixed a bug with `LOAD DATA` in `mysqlbinlog`. (Bug#670 [<http://bugs.mysql.com/670>])
- Fixed that `SET CHARACTER SET DEFAULT` works. (Bug#462 [<http://bugs.mysql.com/462>])
- Fixed `MERGE` table behavior in `ORDER BY ... DESC` queries. (Bug#515 [<http://bugs.mysql.com/515>])
- Fixed server crash on `PURGE MASTER LOGS` or `SHOW MASTER LOGS` when the binary log is off. (Bug#733 [<http://bugs.mysql.com/733>])
- Fixed password-checking problem on Windows. (Bug#464 [<http://bugs.mysql.com/464>])
- Fixed the bug in comparison of a `DATETIME` column and an integer constant. (Bug#504 [<http://bugs.mysql.com/504>])
- Fixed remote mode of `mysqlbinlog`. (Bug#672 [<http://bugs.mysql.com/672>])
- Fixed `ERROR 1105: Unknown error` that occurred for some `SELECT` queries, where a column that was declared as `NOT NULL` was compared with an expression that took `NULL` value.
- Changed timeout in `mysql_real_connect()` to use `poll()` instead of `select()` to work around problem with many open files in the client.
- Fixed incorrect results from `MATCH ... AGAINST` used with a `LEFT JOIN` query.
- Fixed a bug that limited the maximum value for `mysqld` variables to 4294967295 when they are specified on the command line.
- Fixed a bug that sometimes caused spurious “Access denied” errors in `HANDLER ... READ` statements, when a table is referenced via an alias.
- Fixed portability problem with `safe_malloc`, which caused MySQL to give “Freeing wrong aligned pointer” errors on SCO 3.2.
- `ALTER TABLE ... ENABLE/DISABLE KEYS` could cause a core dump when done after an `INSERT DELAYED` statement on the same table.
- Fixed problem with conversion of localtime to GMT where some times resulted in different (but correct) timestamps. Now MySQL should use the smallest possible timestamp value in this case. (Bug#316 [<http://bugs.mysql.com/316>])
- Very small query cache sizes could crash `mysqld`. (Bug#549 [<http://bugs.mysql.com/549>])

- Fixed a bug (accidentally introduced by us but present only in version 4.0.13) that made `INSERT ... SELECT` into an `AUTO_INCREMENT` column not replicate well. This bug is in the master, not in the slave. (Bug#490 [<http://bugs.mysql.com/490>])
- Fixed a bug: When an `INSERT ... SELECT` statement inserted rows into a non-transactional table, but failed at some point (for example, due to a “Duplicate key” error), the query was not written to the binary log. Now it is written to the binary log, with its error code, as all other queries are. About the `slave-skip-errors` option for how to handle partially completed queries in the slave, see Section 6.8, “Replication Startup Options”. (Bug#491 [<http://bugs.mysql.com/491>])
- `SET FOREIGN_KEY_CHECKS=0` was not replicated properly. The fix probably will not be backported to 3.23.
- On a slave, `LOAD DATA INFILE` which had no `IGNORE` or `REPLACE` clause on the master, was replicated with `IGNORE`. Although this is not a problem if the master and slave data are identical (a `LOAD` that produces no duplicate conflicts on the master produces none on the slave anyway), which is true in normal operation, it is better for debugging not to silently add the `IGNORE`. That way, you can get an error message on the slave and discover that for some reason, the data on master and slave are different and investigate why. (Bug#571 [<http://bugs.mysql.com/571>])
- On a slave, `LOAD DATA INFILE` printed an incomplete “Duplicate entry '%-.64s' for key %d” message (the key name and value were not mentioned) in case of duplicate conflict (which does not happen in normal operation). (Bug#573 [<http://bugs.mysql.com/573>])
- When using a slave compiled with `--debug`, `CHANGE MASTER TO RELAY_LOG_POS` could cause a debug assertion failure. (Bug#576 [<http://bugs.mysql.com/576>])
- When doing a `LOCK TABLES WRITE` on an InnoDB table, commit could not happen, if the query was not written to the binary log (for example, if `--log-bin` was not used, or `binlog-ignore-db` was used). (Bug#578 [<http://bugs.mysql.com/578>])
- If a 3.23 master had open temporary tables that had been replicated to a 4.0 slave, and the binary log got rotated, these temporary tables were immediately dropped by the slave (which caused problems if the master used them subsequently). This bug had been fixed in 4.0.13, but in a manner which caused an unlikely inconvenience: If the 3.23 master died brutally (power failure), without having enough time to automatically write `DROP TABLE` statements to its binary log, then the 4.0.13 slave would not notice the temporary tables have to be dropped, until the slave `mysqld` server is restarted. This minor inconvenience is fixed in 3.23.57 and 4.0.14 (meaning the master must be upgraded to 3.23.57 and the slave to 4.0.14 to remove the inconvenience). (Bug#254 [<http://bugs.mysql.com/254>])
- If `MASTER_POS_WAIT()` was waiting, and the slave was idle, and the slave SQL thread terminated, `MASTER_POS_WAIT()` would wait forever. Now when the slave SQL thread terminates, `MASTER_POS_WAIT()` immediately returns `NULL` (“slave stopped”). (Bug#651 [<http://bugs.mysql.com/651>])
- After `RESET SLAVE; START SLAVE;`, the `Relay_Log_Space` value displayed by `SHOW SLAVE STATUS` was too big by four bytes. (Bug#763 [<http://bugs.mysql.com/763>])
- If a query was ignored on the slave (because of `--replicate-ignore-table` and other similar rules), the slave still checked if the query got the same error code (0, no error) as on the master. So if the master had an error on the query (for example, “Duplicate entry” in a multiple-row insert), then the slave stopped and warned that the error codes didn't match. (Bug#797 [<http://bugs.mysql.com/797>])

## D.2.16. Changes in release 4.0.13 (16 May 2003)

Functionality added or changed:

- `PRIMARY KEY` now implies `NOT NULL`. (Bug#390 [http://bugs.mysql.com/390])
- The Windows binary packages are now compiled with `--enable-local-infile` to match the Unix build configuration.
- Removed timing of tests from `mysql-test-run.time` does not accept all required parameters on many platforms (for example, QNX) and timing the tests is not really required (it's not a benchmark anyway).
- `SHOW MASTER STATUS` and `SHOW SLAVE STATUS` required the `SUPER` privilege; now they accept `REPLICATION CLIENT` as well. (Bug#343 [http://bugs.mysql.com/343])
- Added multi-threaded `MyISAM` repair optimization and `myisam_repair_threads` variable to enable it. See Section 5.2.2, "System Variables".
- Added `innodb_max_dirty_pages_pct` variable which controls amount of dirty pages allowed in `InnoDB` buffer pool.
- `CURRENT_USER()` and `Access denied` error messages now report the hostname exactly as it was specified in the `GRANT` command.
- Removed benchmark results from the source and binary distributions. They are still available in the BK source tree, though.
- `InnoDB` tables now support `ANALYZE TABLE`.
- MySQL now issues an error when it opens a table that was created with MySQL 4.1.
- Option `--new` now changes binary items (`0xFFDF`) to be treated as binary strings instead of numbers by default. This fixes some problems with character sets where it's convenient to input the string as a binary item. After this change you have to convert the binary string to `INTEGER` with a `CAST` if you want to compare two binary items with each other and know which one is bigger than the other. `SELECT CAST(0xfeff AS UNSIGNED) < CAST(0xff AS UNSIGNED)`. This is the default behavior in MySQL 4.1. (Bug#152 [http://bugs.mysql.com/152])
- Enabled `delayed_insert_timeout` on Linux (most modern `glibc` libraries have a fixed `pthread_cond_timedwait()`). (Bug#211 [http://bugs.mysql.com/211])
- Don't create more insert delayed threads than given by `max_delayed_threads`. (Bug#211 [http://bugs.mysql.com/211])
- Changed `UPDATE ... LIMIT` to apply the limit to rows that were matched, whether or not they actually were changed. Previously the limit was applied as a restriction on the number of rows changed.
- Tuned optimizer to favor clustered index over table scan.
- `BIT_AND()` and `BIT_OR()` now return an unsigned 64-bit value.
- Added warnings to error log of why a secure connection failed (when running with `-log-warnings`).
- Deprecated options `--skip-symlink` and `--use-symbolic-links` and replaced these with `--symbolic-links`.
- The default option for `innodb_flush_log_at_trx_commit` was changed from 0 to 1 to

make InnoDB tables ACID by default. See [Section 14.2.5, “InnoDB Startup Options and System Variables”](#).

- Added a feature to `SHOW KEYS` to display keys that are disabled by `ALTER TABLE DISABLE KEYS` command.
- When using a non-existing table type with `CREATE TABLE`, first try if the default table type exists before falling back to `MyISAM`.
- Added `MEMORY` as an alias for `HEAP`.
- Renamed function `rnd` to `my_rnd` as the name was too generic and is an exported symbol in `libmysqlclient` (thanks to Dennis Haney for the initial patch).
- Portability fix: renamed `include/dbug.h` to `include/my_dbug.h`.
- `mysqldump` no longer silently deletes the binary logs when invoked with the `--master-data` or `--first-slave` option; while this behavior was convenient for some users, others may suffer from it. Now you must explicitly ask for binary logs to be deleted by using the new `--delete-master-logs` option.
- If the slave is configured (using for example `--replicate-wild-ignore-table=mysql.%`) to exclude `mysql.user`, `mysql.host`, `mysql.db`, `mysql.tables_priv` and `mysql.columns_priv` from replication, then `GRANT` and `REVOKE` are not replicated.

Bugs fixed:

- Logged `Access denied` error message had incorrect `Using password` value. ([Bug#398](#) [<http://bugs.mysql.com/398>])
- Fixed bug with `NATURAL LEFT JOIN`, `NATURAL RIGHT JOIN` and `RIGHT JOIN` when using many joined tables. The problem was that the `JOIN` method was not always associated with the tables surrounding the `JOIN` method. If you have a query that uses many `RIGHT JOIN` or `NATURAL ... JOINS` you should verify that they work as you expected after upgrading MySQL to this version. ([Bug#291](#) [<http://bugs.mysql.com/291>])
- Fixed `mysql` parser not to erroneously interpret ‘`'`’ or ‘`"`’ characters within `/* ... */` comment as beginning a quoted string.
- `mysql` command-line client no longer looks for `\*` commands inside backtick-quoted strings.
- Fixed `Unknown error` when using `UPDATE ... LIMIT`. ([Bug#373](#) [<http://bugs.mysql.com/373>])
- Fixed problem with ANSI mode and `GROUP BY` with constants. ([Bug#387](#) [<http://bugs.mysql.com/387>])
- Fixed bug with `UNION` and `OUTER JOIN`. ([Bug#386](#) [<http://bugs.mysql.com/386>])
- Fixed bug if one used a multiple-table `UPDATE` and the query required a temporary table bigger than `tmp_table_size`. ([Bug#286](#) [<http://bugs.mysql.com/286>])
- Run `mysql_install_db` with the `-IN-RPM` option for the Mac OS X installation to not fail on systems with improperly configured hostname configurations.
- `LOAD DATA INFILE` now reads `000000` as a zero date instead of `"2000-00-00"`.

- Fixed bug that caused `DELETE FROM table WHERE const_expression` always to delete the whole table (even if expression result was false). (Bug#355 [<http://bugs.mysql.com/355>])
- Fixed core dump bug when using `FORMAT( 'nan' ,#)`. (Bug#284 [<http://bugs.mysql.com/284>])
- Fixed name resolution bug with `HAVING ... COUNT(DISTINCT ...)`.
- Fixed incorrect result from truncation operator (\*) in `MATCH ... AGAINST()` in some complex joins.
- Fixed a crash in `REPAIR ... USE_FRM` command, when used on read-only, non-existing table or a table with a crashed index file.
- Fixed a crashing bug in `mysql` monitor program. It occurred if program was started with `--no-defaults`, with a prompt that contained the hostname and a connection to a non-existent database was requested.
- Fixed problem when comparing a key for a multi-byte character set. (Bug#152 [<http://bugs.mysql.com/152>])
- Fixed bug in `LEFT`, `RIGHT` and `MID` when used with multi-byte character sets and some `GROUP BY` queries. (Bug#314 [<http://bugs.mysql.com/314>])
- Fix problem with `ORDER BY` being discarded for some `DISTINCT` queries. (Bug#275 [<http://bugs.mysql.com/275>])
- Fixed that `SET SQL_BIG_SELECTS=1` works as documented (This corrects a new bug introduced in 4.0)
- Fixed some serious bugs in `UPDATE ... ORDER BY`. (Bug#241 [<http://bugs.mysql.com/241>])
- Fixed unlikely problem in optimizing `WHERE` clause with constant expression like in `WHERE 1 AND (a=1 AND b=1)`.
- Fixed that `SET SQL_BIG_SELECTS=1` works again.
- Introduced proper backtick quoting for `db.table` in `SHOW GRANTS`.
- `FULLTEXT` index stopped working after `ALTER TABLE` that converts `TEXT` column to `CHAR`. (Bug#283 [<http://bugs.mysql.com/283>])
- Fixed a security problem with `SELECT` and wildcarded select list, when user only had partial column `SELECT` privileges on the table.
- Mark a `MyISAM` table as "analyzed" only when all the keys are indeed analyzed.
- Only ignore world-writable `my.cnf` files that are regular files (and not, for example, named pipes or character devices).
- Fixed few smaller issues with `SET PASSWORD`.
- Fixed error message which contained deprecated text.
- Fixed a bug with two `NATURAL JOINS` in the query.
- `SUM()` didn't return `NULL` when there was no rows in result or when all values was `NULL`.
- On Unix, symbolic link handling was not enabled by default and there was no way to turn this on.
- Added missing dashes to parameter `--open-files-limit` in `mysqld_safe`. (Bug#264

[<http://bugs.mysql.com/264>])

- Fixed incorrect hostname for TCP/IP connections displayed in `SHOW PROCESSLIST`.
- Fixed a bug with `NAN` in `FORMAT( . . . )` function ...
- Fixed a bug with improperly cached database privileges.
- Fixed a bug in `ALTER TABLE ENABLE / DISABLE KEYS` which failed to force a refresh of table data in the cache.
- Fixed bugs in replication of `LOAD DATA INFILE` for custom parameters (`ENCLOSED`, `TERMINATED` and so on) and temporary tables. (Bug#183 [<http://bugs.mysql.com/183>], Bug#222 [<http://bugs.mysql.com/222>])
- Fixed a replication bug when the master is 3.23 and the slave 4.0: the slave lost the replicated temporary tables if `FLUSH LOGS` was issued on the master. (Bug#254 [<http://bugs.mysql.com/254>])
- Fixed a bug when doing `LOAD DATA INFILE IGNORE`: When reading the binary log, `mysql-binlog` and the replication code read `REPLACE` instead of `IGNORE`. This could make the slave's table become different from the master's table. (Bug#218 [<http://bugs.mysql.com/218>])
- Fixed a deadlock when `relay_log_space_limit` was set to a too small value. (Bug#79 [<http://bugs.mysql.com/79>])
- Fixed a bug in `HAVING` clause when an alias is used from the **select list**.
- Fixed overflow bug in `MyISAM` when a row is inserted into a table with a large number of columns and at least one `BLOB/TEXT` column. Bug was caused by incorrect calculation of the needed buffer to pack data.
- Fixed a bug when `SELECT @non_existent_variable` caused an error in the client/server protocol due to `net_printf()` output being sent to the client twice.
- Fixed a bug in setting `SQL_BIG_SELECTS` option.
- Fixed a bug in `SHOW PROCESSLIST` which only displayed a localhost in the "Host" column. This was caused by a glitch that used only current thread information instead of information from the linked list of threads.
- Removed unnecessary Mac OS X helper files from server RPM. (Bug#144 [<http://bugs.mysql.com/144>])
- Allow optimization of multiple-table update for `InnoDB` tables as well.
- Fixed a bug in multiple-table updates that caused some rows to be updated several times.
- Fixed a bug in `mysqldump` when it was called with `--master-data`: the `CHANGE MASTER TO` commands appended to the SQL dump had incorrect coordinates. (Bug#159 [<http://bugs.mysql.com/159>])
- Fixed a bug when an updating query using `USER()` was replicated on the slave; this caused a segmentation fault on the slave. (Bug#178 [<http://bugs.mysql.com/178>]). `USER()` is still badly replicated on the slave (it is replicated to " ").

## D.2.17. Changes in release 4.0.12 (15 March 2003: Production)

Functionality added or changed:

- `mysqld` no longer reads options from world-writable config files. (CVE-2003-0150 [http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2003-0150])
- Integer values between 9223372036854775807 and 9999999999999999 are now regarded as unsigned longlongs, not as floats. This makes these values work similar to values between 1000000000000000000 and 18446744073709551615.
- `SHOW PROCESSLIST` now includes the client TCP port after the hostname to make it easier to know from which client the request originated.
- The `--new` option can be used to make a 4.0 server return `TIMESTAMP` as a string in `'YYYY-MM-DD HH:MM:SS'` format, the way that 4.1 servers do. This is also a `new` system variable that can be set for the same effect. See Section 11.3.1.1, “`TIMESTAMP` Properties Prior to MySQL 4.1”.

Bugs fixed:

- Fixed `mysqld` crash on extremely small values of `sort_buffer` variable.
- `INSERT INTO u SELECT ... FROM t` was written too late to the binary log if `t` was very frequently updated during the execution of this query. This could cause a problem with `mysqlbinlog` or replication. The master must be upgraded, not the slave. (Bug#136 [http://bugs.mysql.com/136])
- Fixed checking of random part of `WHERE` clause. (Bug#142 [http://bugs.mysql.com/142])
- Fixed a bug with multiple-table updates with `InnoDB` tables. This bug occurred as, in many cases, `InnoDB` tables cannot be updated “on the fly,” but offsets to the records have to be stored in a temporary table.
- Added missing file `mysql_secure_installation` to the `server` RPM subpackage. (Bug#141 [http://bugs.mysql.com/141])
- Fixed MySQL (and `myisamchk`) crash on artificially corrupted `.MYI` files.
- Don't allow `BACKUP TABLE` to overwrite existing files.
- Fixed a bug with multiple-table `UPDATE` statements when user had all privileges on the database where tables are located and there were any entries in `tables_priv` table, that is, `grant_option` was true.
- Fixed a bug that allowed a user with table or column grants on some table, `TRUNCATE` any table in the same database.
- Fixed deadlock when doing `LOCK TABLE` followed by `DROP TABLE` in the same thread. In this case one could still kill the thread with `KILL`.
- `LOAD DATA LOCAL INFILE` was not properly written to the binary log (hence not properly replicated). (Bug#82 [http://bugs.mysql.com/82])
- `RAND()` entries were not read correctly by `mysqlbinlog` from the binary log which caused problems when restoring a table that was inserted with `RAND(). INSERT INTO t1 VALUES (RAND())`. In replication this worked okay.
- `SET SQL_LOG_BIN=0` was ignored for `INSERT DELAYED` queries. (Bug#104

[<http://bugs.mysql.com/104>])

- `SHOW SLAVE STATUS` reported too old positions (columns `Relay_Master_Log_File` and `Exec_Master_Log_Pos`) for the last executed statement from the master, if this statement was the `COMMIT` of a transaction. The master must be upgraded for that, not the slave. ([Bug#52](http://bugs.mysql.com/52) [<http://bugs.mysql.com/52>])
- `LOAD DATA INFILE` was not replicated by the slave if `replicate_*_table` was set on the slave. ([Bug#86](http://bugs.mysql.com/86) [<http://bugs.mysql.com/86>])
- After `RESET SLAVE`, the coordinates displayed by `SHOW SLAVE STATUS` looked un-reset (although they were, but only internally). ([Bug#70](http://bugs.mysql.com/70) [<http://bugs.mysql.com/70>])
- Fixed query cache invalidation on `LOAD DATA`.
- Fixed memory leak on `ANALYZE` procedure with error.
- Fixed a bug in handling `CHAR(0)` columns that could cause incorrect results from the query.
- Fixed rare bug with incorrect initialization of `AUTO_INCREMENT` column, as a secondary column in a multi-column key (see [Section 3.6.9, “Using AUTO\\_INCREMENT”](#)), when data was inserted with `INSERT ... SELECT` or `LOAD DATA` into an empty table.
- On Windows, `STOP SLAVE` didn't stop the slave until the slave got one new command from the master (this bug has been fixed for MySQL 4.0.11 by releasing updated 4.0.11a Windows packages, which include this individual fix on top of the 4.0.11 sources). ([Bug#69](http://bugs.mysql.com/69) [<http://bugs.mysql.com/69>])
- Fixed a crash when no database was selected and `LOAD DATA` command was issued with full table name specified, including database prefix.
- Fixed a crash when shutting down replication on some platforms (for example, Mac OS X).
- Fixed a portability bug with `pthread_attr_getstacksize` on HP-UX 10.20 (Patch was also included in 4.0.11a sources).
- Fixed the `bigint` test to not fail on some platforms (for example, HP-UX and Tru64) due to different return values of the `atof()` function.
- Fixed the `rpl_rotate_logs` test to not fail on certain platforms (such as Mac OS X) due to a too-long file name (changed `slave-master-info.opt` to `.slave-mi`).

## D.2.18. Changes in release 4.0.11 (20 February 2003)

Functionality added or changed:

- `NULL` is now sorted **LAST** if you use `ORDER BY ... DESC` (as it was before MySQL 4.0.2). This change was required to comply with the SQL standard. (The original change was made because we thought that standard SQL required `NULL` to be always sorted at the same position, but this was incorrect).
- Added `START TRANSACTION` (standard SQL syntax) as alias for `BEGIN`. This is recommended to use instead of `BEGIN` to start a transaction.
- Added `OLD_PASSWORD()` as a synonym for `PASSWORD()`.
- Allow keyword `ALL` in group functions.



- Added support for some new `INNER JOIN` and `JOIN` syntaxes. For example, `SELECT * FROM t1 INNER JOIN t2` didn't work before.
- Novell NetWare 6.0 porting effort completed, Novell patches merged into the main source tree.

Bugs fixed:

- Fixed problem with multiple-table delete and `InnoDB` tables.
- Fixed a problem with `BLOB NOT NULL` columns used with `IS NULL`.
- Re-added missing pre- and post(un)install scripts to the Linux RPM packages (they were missing after the renaming of the server subpackage).
- Fixed that table locks are not released with multiple-table updates and deletes with `InnoDB` storage engine.
- Fixed bug in updating `BLOB` columns with long strings.
- Fixed integer-wraparound when giving big integer ( $\geq 10$  digits) to function that requires an unsigned argument, like `CREATE TABLE (...) AUTO_INCREMENT=N`.
- `MIN(key_column)` could in some cases return `NULL` on a column with `NULL` and other values.
- `MIN(key_column)` and `MAX(key_column)` could in some cases return incorrect values when used in `OUTER JOIN`.
- `MIN(key_column)` and `MAX(key_column)` could return incorrect values if one of the tables was empty.
- Fixed rare crash in compressed `MyISAM` tables with blobs.
- Fixed bug in using aggregate functions as argument for `INTERVAL`, `CASE`, `FIELD`, `CONCAT_WS`, `ELT` and `MAKE_SET` functions.
- When running with `--lower-case-table-names` (default on Windows) and you had tables or databases with mixed case on disk, then executing `SHOW TABLE STATUS` followed with `DROP DATABASE` or `DROP TABLE` could fail with `Errcode 13`.

## D.2.19. Changes in release 4.0.10 (29 January 2003)

Functionality added or changed:

- Added option `--log-error[=file_name]` to `mysqld_safe` and `mysqld`. This option forces all error messages to be put in a log file if the option `--console` is not given. On Windows `--log-error` is enabled as default, with a default name of `host_name.err` if the name is not specified.
- Changed some things from `Warning:` to `Note:` in the log files.
- The `mysqld` server should now compile on NetWare.
- Added optimization that if one does `GROUP BY ... ORDER BY NULL` then result is not sorted.
- New `--ft-stopword-file` command-line option for `mysqld` to replace/disable the built-in

stopword list that is used in full-text searches. See [Section 5.2.2, “System Variables”](#).

- Changed default stack size from 64KB to 192KB; This fixes a core dump problem on Red Hat 8.0 and other systems with a `glibc` that requires a stack size larger than 128K for `gethostbyaddr()` to resolve a hostname. You can fix this for earlier MySQL versions by starting `mysqld` with `--thread-stack=192K`.
- Added `mysql_waitpid` to the binary distribution and the `MySQL-client` RPM subpackage (required for `mysql-test-run`).
- Renamed the main `MySQL` RPM package to `MySQL-server`. When updating from an older version, `MySQL-server.rpm` simply replaces `MySQL.rpm`.
- If a slave is configured with `replicate_wild_do_table=db.%` or `replicate_wild_ignore_table=db.%`, these rules are applied to `CREATE/DROP DATABASE`, too.
- Added timeout value for `MASTER_POS_WAIT()`.

Bugs fixed:

- Fixed initialization of the random seed for newly created threads to give a better `rand()` distribution from the first call.
- Fixed a bug that caused `mysqld` to hang when a table was opened with the `HANDLER` command and then dropped without being closed.
- Fixed bug in logging to binary log (which affects replication) a query that inserts a `NULL` in an `AUTO_INCREMENT` column and also uses `LAST_INSERT_ID()`.
- Fixed an unlikely bug that could cause a memory overrun when using `ORDER BY constant_expression`.
- Fixed a table corruption in `myisamchk` parallel repair mode.
- Fixed bug in query cache invalidation on simple table renaming.
- Fixed bug in `mysqladmin --relative`.
- On some 64-bit systems, `show status` reported a strange number for `Open_files` and `Open_streams`.
- Fixed incorrect number of columns in `EXPLAIN` on empty table.
- Fixed bug in `LEFT JOIN` that caused zero rows to be returned in the case the `WHERE` condition was evaluated as `FALSE` after reading `const` tables. (Unlikely condition).
- `FLUSH PRIVILEGES` didn't correctly flush table/column privileges when `mysql.tables_priv` is empty.
- Fixed bug in replication when using `LOAD DATA INFILE` one a file that updated an `AUTO_INCREMENT` column with `NULL` or `0`. This bug only affected MySQL 4.0 masters (not slaves or MySQL 3.23 masters). **Note:** If you have a slave that has replicated a file with generated `AUTO_INCREMENT` columns then the slave data is corrupted and you should reinitialize the affected tables from the master.
- Fixed possible memory overrun when sending a `BLOB` value larger than 16M to the client.

- Fixed incorrect error message when setting a `NOT NULL` column to an expression that returned `NULL`.
- Fixed core dump bug in `str LIKE "%other_str%"` where `str` or `other_str` contained characters  $\geq 128$ .
- Fixed bug: When executing on master `LOAD DATA` and `InnoDB` failed with `table full` error the binary log was corrupted.

## D.2.20. Changes in release 4.0.9 (09 January 2003)

Functionality added or changed:

- `OPTIMIZE TABLE` for `MyISAM` tables treats all `NULL` values as different when calculating cardinality. This helps in optimizing joins between tables where one of the tables has a lot of `NULL` values in a indexed column:

```
SELECT * from t1, t2 where t1.a=t2.key_with_a_lot_of_null;
```

- Added join operator `FORCE INDEX (key_list)`. This acts like `USE INDEX (key_list)` but with the addition that a table scan is assumed to be `VERY` expensive. One bad thing with this is that it makes `FORCE` a reserved word.
- Reset internal row buffer in `MyISAM` after each query. This reduces memory in case you have a lot of big blobs in a table.

Bugs fixed:

- A security patch in 4.0.8 causes the `mysqld` server to die if the remote hostname can't be resolved. This is now fixed.
- Fixed crash when replication big `LOAD DATA INFILE` statement that caused log rotation.

## D.2.21. Changes in release 4.0.8 (07 January 2003)

Functionality added or changed:

- Default `max_packet_length` for `libmysqld.c` is now `1024*1024*1024`.
- You can now specify `max_allowed_packet` in a file read by `mysql_options(MYSQL_READ_DEFAULT_FILE)` for clients.
- When sending a too big packet to the server with the not compressed protocol, the client now gets an error message instead of a lost connection.
- We now send big queries/result rows in bigger hunks, which should give a small speed improvement.
- Fixed some bugs with the compressed protocol for rows  $> 16\text{MB}$ .
- `InnoDB` tables now also support `ON UPDATE CASCADE` in `FOREIGN KEY` constraints. See the `InnoDB` section in the manual for the `InnoDB` changelog.

Bugs fixed:

- Fixed bug in `ALTER TABLE` with `BDB` tables.
- Fixed core dump bug in `QUOTE ( )` function.
- Fixed a bug in handling communication packets bigger than 16MB. Unfortunately this required a protocol change; If you upgrade the server to 4.0.8 and above and have clients that use packets  $\geq 255*255*255$  bytes (=16581375) you must also upgrade your clients to at least 4.0.8. If you don't upgrade, the clients hang when sending a big packet.
- Fixed bug when sending blobs longer than 16MB to client.
- Fixed bug in `GROUP BY` when used on `BLOB` column with `NULL` values.
- Fixed a bug in handling `NULL` values in `CASE ... WHEN ...`

## D.2.22. Changes in release 4.0.7 (20 December 2002)

Functionality added or changed:

- `mysqlbug` now also reports the compiler version used for building the binaries (if the compiler supports the option `--version`).

Bugs fixed:

- Fixed compilation problems on OpenUnix and HP-UX 10.20.
- Fixed some optimization problems when compiling MySQL with `-DBIG_TABLES` on a 32-bit system.
- `mysql_drop_db ( )` didn't check permissions properly so anyone could drop another users database. `DROP DATABASE` is checked properly.

## D.2.23. Changes in release 4.0.6 (14 December 2002: Gamma)

Functionality added or changed:

- Added syntax support for `CHARACTER SET xxx` and `CHARSET=xxx` table options (to be able to read table dumps from 4.1).
- Fixed replication bug that caused the slave to lose its position in some cases when the replication log was rotated.
- Fixed that a slave restarts from the start of a transaction if it's killed in the middle of one.
- Moved the manual pages from `man` to `man/man1` in the binary distributions.
- The default type returned by `IFNULL (A, B)` is now set to be the more 'general' of the types of `A` and `B`. (The order is `STRING`, `REAL` or `INTEGER`).
- Moved the `mysql.server` startup script in the RPM packages from /

`etc/rc.d/init.d/mysql` to `/etc/init.d/mysql` (which almost all current Linux distributions support for LSB compliance).

- Added `Qcache_lowmem_prunes` status variable (number of queries that were deleted from the cache because of low memory).
- Fixed `mysqlcheck` so it can deal with table names containing dashes.
- Bulk insert optimization (see [Section 5.2.2, “System Variables”](#)) is no longer used when inserting small (less than 100) number of rows.
- Optimization added for queries like `SELECT ... FROM merge_table WHERE indexed_column=constant_expr`.
- Added functions `LOCALTIME` and `LOCALTIMESTAMP` as synonyms for `NOW()`.
- `CEIL` is now an alias for `CEILING`.
- The `CURRENT_USER()` function can be used to get a `user@host` value as it was matched in the `GRANT` system. See [Section 12.9.3, “Information Functions”](#).
- Fixed `CHECK` constraints to be compatible with standard SQL. This made `CHECK` a reserved word. (Checking of `CHECK` constraints is still not implemented).
- Added `CAST(... as CHAR)`.
- Added PostgreSQL compatible `LIMIT` syntax: `SELECT ... LIMIT row_count OFFSET offset`
- `mysql_change_user()` now resets the connection to the state of a fresh connect (Ie, `ROLLBACK` any active transaction, close all temporary tables, reset all user variables etc..)
- `CHANGE MASTER` and `RESET SLAVE` now require that slave threads both be stopped; these commands return an error if at least one of these two threads is running.

Bugs fixed:

- Fixed number of found rows returned in `multi table updates`
- Make `--lower-case-table-names` default on Mac OS X as the default filesystem (HFS+) is case insensitive. See [Section 9.2.2, “Identifier Case Sensitivity”](#).
- Transactions in `AUTOCOMMIT=0` mode didn't rotate binary log.
- A fix for the bug in a `SELECT` with joined tables with `ORDER BY` and `LIMIT` clause when `file-sort` had to be used. In that case `LIMIT` was applied to `filesort` of one of the tables, although it could not be. This fix also solved problems with `LEFT JOIN`.
- `mysql_server_init()` now makes a copy of all arguments. This fixes a problem when using the embedded server in C# program.
- Fixed buffer overrun in `libmysqlclient` library that allowed a malicious MySQL server to crash the client application. ([CVE-2002-1376](#) [<http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2002-1376>])
- Fixed security-related bug in `mysql_change_user()` handling. All users are strongly recommended to upgrade to version 4.0.6. ([CVE-2002-1374](#))

[<http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2002-1374>], [CVE-2002-1375](#)  
[<http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2002-1375>])

- Fixed bug that prevented `--chroot` command-line option of `mysqld` from working.
- Fixed bug in phrase operator `"..."` in boolean full-text search.
- Fixed bug that caused `OPTIMIZE TABLE` to corrupt the table under some rare circumstances.
- Part rewrite of multiple-table-update to optimize it, make it safer and more bug-free.
- `LOCK TABLES` now works together with multiple-table-update and multiple-table-delete.
- `--replicate-do=xxx` didn't work for `UPDATE` commands. (Bug introduced in 4.0.0)
- Fixed shutdown problem on Mac OS X.
- Major InnoDB bugs in `REPLACE`, `AUTO_INCREMENT`, `INSERT INTO ... SELECT ...` were fixed. See the InnoDB changelog in the InnoDB section of the manual.
- `RESET SLAVE` caused a crash if the slave threads were running.

## D.2.24. Changes in release 4.0.5 (13 November 2002)

Functionality added or changed:

- Port number was added to hostname (if it is known) in `SHOW PROCESSLIST` command
- Changed handling of last argument in `WEEK()` so that you can get week number according to the ISO 8601 specification. (Old code should still work).
- Fixed that `INSERT DELAYED` threads don't hang on `Waiting for INSERT` when one sends a `SIGHUP` to `mysqld`.
- Change that `AND` works according to standard SQL when it comes to `NULL` handling. In practice, this affects only queries where you do something like `WHERE ... NOT (NULL AND 0)`.
- `mysqld` now resolves `basedir` to its full path (with `realpath()`). This enables one to use relative symlinks to the MySQL installation directory. This however causes `show variables` to report different directories on systems where there is a symbolic link in the path.
- Fixed that MySQL does not use index scan on index disabled with `IGNORE INDEX` or `USE INDEX`. to be ignored.
- Added `--use-frm` option to `mysqlcheck`. When used with `REPAIR TABLE`, it gets the table structure from the `.frm` file, so the table can be repaired even if the `.MYI` header is corrupted.
- Fixed bug in `MAX()` optimization when used with `JOIN` and `ON` expressions.
- Added support for reading of MySQL 4.1 table definition files.
- `BETWEEN` behavior changed (see [Section 12.1.3, "Comparison Functions and Operators"](#)). Now `datetime_col BETWEEN timestamp AND timestamp` should work as expected.
- One can create `TEMPORARY MERGE` tables now.
- `DELETE FROM myisam_table` now shrinks not only the `.MYD` file but also the `.MYI` file.

- When one uses the `--open-files-limit=val` option to `mysqld_safe` it's now passed on to `mysqld`.
- Changed output from `EXPLAIN` from `'where used'` to `'Using where'` to make it more in line with other output.
- Removed variable `safe_show_database` as it was no longer used.
- Updated source tree to be built using `automake 1.5` and `libtool 1.4`.
- Fixed an inadvertently changed option (`--ignore-space`) back to the original `--ignore-spaces` in `mysqlclient`. (Both syntaxes work).
- Don't require `UPDATE` privilege when using `REPLACE`.
- Added support for `DROP TEMPORARY TABLE ...`, to be used to make replication safer.
- When transactions are enabled, all commands that update temporary tables inside a `BEGIN/COMMIT` are now stored in the binary log on `COMMIT` and not stored if one does `ROLLBACK`. This fixes some problems with non-transactional temporary tables used inside transactions.
- Allow braces in joins in all positions. Formerly, things like `SELECT * FROM (t2 LEFT JOIN t3 USING (a)), t1` worked, but not `SELECT * FROM t1, (t2 LEFT JOIN t3 USING (a))`. Note that braces are simply removed, they do not change the way the join is executed.
- `InnoDB` now supports also isolation levels `READ UNCOMMITTED` and `READ COMMITTED`. For a detailed `InnoDB` changelog, see [Section D.4, "Changes in InnoDB"](#).

Bugs fixed:

- Fixed bug in `MAX()` optimization when used with `JOIN` and `ON` expressions.
- Fixed that `INSERT DELAY` threads don't hang on `Waiting for INSERT` when one sends a `SIGHUP` to `mysqld`.
- Fixed that MySQL does not use an index scan on an index that has been disabled with `IGNORE INDEX` or `USE INDEX`.
- Corrected test for `root` user in `mysqld_safe`.
- Fixed error message issued when storage engine cannot do `CHECK TABLE` or `REPAIR TABLE`.
- Fixed rare core dump problem in complicated `GROUP BY` queries that didn't return any result.
- Fixed `mysqlshow` to work properly with wildcarded database names and with database names that contain underscores.
- Portability fixes to get MySQL to compile cleanly with Sun Forte 5.0.
- Fixed `MyISAM` crash when using dynamic-row tables with huge numbers of packed columns.
- Fixed query cache behavior with `BDB` transactions.
- Fixed possible floating point exception in `MATCH` relevance calculations.
- Fixed bug in full-text search `IN BOOLEAN MODE` that made `MATCH` to return incorrect relevance value in some complex joins.

- Fixed a bug that limited `MyISAM` key length to a value slightly less than 500. It is exactly 500 now.
- Fixed that `GROUP BY` on columns that may have a `NULL` value doesn't always use disk based temporary tables.
- The filename argument for the `--des-key-file` argument to `mysqld` is interpreted relative to the data directory if given as a relative pathname.
- Removed a condition that temp table with index on column that can be `NULL` has to be `MyISAM`. This was okay for 3.23, but not needed in 4.\*. This resulted in slowdown in many queries since 4.0.2.
- Small code improvement in multiple-table updates.
- Fixed a newly introduced bug that caused `ORDER BY ... LIMIT row_count` to not return all rows.
- Fixed a bug in multiple-table deletes when outer join is used on an empty table, which gets first to be deleted.
- Fixed a bug in multiple-table updates when a single table is updated.
- Fixed bug that caused `REPAIR TABLE` and `myisamchk` to corrupt `FULLTEXT` indexes.
- Fixed bug with caching the `mysql` grant table database. Now queries in this database are not cached in the query cache.
- Small fix in `mysqld_safe` for some shells.
- Give error if a `MyISAM MERGE` table has more than  $2^{32}$  rows and MySQL was not compiled with `-DBIG_TABLES`.
- Fixed some `ORDER BY ... DESC` problems with `InnoDB` tables.

## D.2.25. Changes in release 4.0.4 (29 September 2002)

- Fixed bug where `GRANT/REVOKE` failed if hostname was given in non-matching case.
- Don't give warning in `LOAD DATA INFILE` when setting a `timestamp` to a string value of `'0'`.
- Fixed bug in `myisamchk -R` mode.
- Fixed bug that caused `mysqld` to crash on `REVOKE`.
- Fixed bug in `ORDER BY` when there is a constant in the `SELECT` statement.
- One didn't get an error message if `mysqld` couldn't open the privilege tables.
- `SET PASSWORD FOR ...` closed the connection in case of errors (bug from 4.0.3).
- Increased maximum possible `max_allowed_packet` in `mysqld` to 1GB.
- Fixed bug when doing a multiple-row `INSERT` on a table with an `AUTO_INCREMENT` key which was not in the first part of the key.
- Changed `LOAD DATA INFILE` to not re-create index if the table had rows from before.



- Fixed overrun bug when calling `AES_DECRYPT()` with incorrect arguments.
- `--skip-ssl` can now be used to disable SSL in the MySQL clients, even if one is using other SSL options in an option file or previously on the command line.
- Fixed bug in `MATCH ... AGAINST( ... IN BOOLEAN MODE)` used with `ORDER BY`.
- Added `LOCK TABLES` and `CREATE TEMPORARY TABLES` privilege on the database level. You must run the `mysql_fix_privilege_tables` script on old installations to activate these.
- In `SHOW TABLE ... STATUS`, compressed tables sometimes showed up as `dynamic`.
- `SELECT @@[global|session].var_name` didn't report `global | session` in the result column name.
- Fixed problem in replication that `FLUSH LOGS` in a circular replication setup created an infinite number of binary log files. Now a `rotate-binary-log` command in the binary log does not cause slaves to rotate logs.
- Removed `STOP EVENT` from binary log when doing `FLUSH LOGS`.
- Disable the use of `SHOW NEW MASTER FOR SLAVE` as this needs to be completely reworked in a future release.
- Fixed a bug with constant expression (for example, column of a one-row table, or column from a table, referenced by a `UNIQUE` key) appeared in `ORDER BY` part of `SELECT DISTINCT`.
- `--log-binary=a.b.c` now properly strips off `.b.c`.
- `FLUSH LOGS` removed numerical extension for all future update logs.
- `GRANT ... REQUIRE` didn't store the SSL information in the `mysql.user` table if SSL was not enabled in the server.
- `GRANT ... REQUIRE NONE` can now be used to remove SSL information.
- `AND` is now optional between `REQUIRE` options.
- `REQUIRE` option was not properly saved, which could cause strange output in `SHOW GRANTS`.
- Fixed that `mysqld --help` reports correct values for `--datadir` and `--bind-address`.
- Fixed that one can drop UDFs that didn't exist when `mysqld` was started.
- Fixed core dump problem with `SHOW VARIABLES` on some 64-bit systems (like Solaris SPARC).
- Fixed a bug in `my_getopt()`; `--set-variable` syntax didn't work for those options that didn't have a valid variable in the `my_option` struct. This affected at least the `default-table-type` option.
- Fixed a bug from 4.0.2 that caused `REPAIR TABLE` and `myisamchk --recover` to fail on tables with duplicates in a unique key.
- Fixed a bug from 4.0.3 in calculating the default data type for some functions. This affected queries of type `CREATE TABLE tbl_name SELECT expression(),...`
- Fixed bug in queries of type `SELECT * FROM table-list GROUP BY ...` and `SELECT DISTINCT * FROM ....`
- Fixed bug with the `--slow-log` when logging an administrator command (like `FLUSH`

).

- Fixed a bug that `OPTIMIZE TABLE` of locked and modified table, reported table corruption.
- Fixed a bug in `my_getopt()` in handling of special prefixes (`--skip-`, `--enable-`). `--skip-external-locking` didn't work and the bug may have affected other similar options.
- Fixed bug in checking for output file name of the `tee` option.
- Added some more optimization to use index for `SELECT ... FROM many_tables .. ORDER BY key limit #`
- Fixed problem in `SHOW OPEN TABLES` when a user didn't have access permissions to one of the opened tables.

## D.2.26. Changes in release 4.0.3 (26 August 2002: Beta)

- Fixed problem with types of user variables. (Bug#551 [<http://bugs.mysql.com/551>])
- Fixed problem with `configure ... --localstatedir=...`
- Cleaned up `mysql.server` script.
- Fixed a bug in `mysqladmin shutdown` when pid file was modified while `mysqladmin` was still waiting for the previous one to disappear. This could happen during a very quick restart and caused `mysqladmin` to hang until `shutdown_timeout` seconds had passed.
- Don't increment warnings when setting `AUTO_INCREMENT` columns to `NULL` in `LOAD DATA INFILE`.
- Fixed all boolean type variables/options to work with the old syntax, for example, all of these work: `--lower-case-table-names`, `--lower-case-table-names=1`, `-O lower-case-table-names=1`, `--set-variable=lower-case-table-names=1`
- Fixed shutdown problem (SIGTERM signal handling) on Solaris. (Bug from 4.0.2).
- `SHOW MASTER STATUS` now returns an empty set if binary log is not enabled.
- `SHOW SLAVE STATUS` now returns an empty set if slave is not initialized.
- Don't update `MyISAM` index file on update if not strictly necessary.
- Fixed bug in `SELECT DISTINCT ... FROM many_tables ORDER BY not-used-column`.
- Fixed a bug with `BIGINT` values and quoted strings.
- Added `QUOTE()` function that performs SQL quoting to produce values that can be used as data values in queries.
- Changed variable `DELAY_KEY_WRITE` to an enumeration to allow it to be set for all tables without taking down the server.
- Changed behavior of `IF(condition, column, NULL)` so that it returns the value in the column's data type.
- Made `safe mysqld` a symlink to `mysqld_safe` in binary distribution.

- Fixed security bug when having an empty database name in the `user.db` table.
- Fixed some problems with `CREATE TABLE ... SELECT function()`.
- `mysqld` now has the option `--temp-pool` enabled by default as this gives better performance with some operating systems.
- Fixed problem with too many allocated alarms on slave when connecting to master many times (normally not a very critical error).
- Fixed hang in `CHANGE MASTER TO` if the slave thread died very quickly.
- Big cleanup in replication code (less logging, better error messages, etc..)
- If the `--code-file` option is specified, the server calls `setrlimit()` to set the maximum allowed core file size to unlimited, so core files can be generated.
- Fixed bug in query cache after temporary table creation.
- Added `--count=N (-c)` option to `mysqladmin`, to make the program do only *N* iterations. To be used with `--sleep (-i)`. Useful in scripts.
- Fixed bug in multiple-table `UPDATE`: when updating a table, `do_select()` became confused about reading records from a cache.
- Fixed bug in multiple-table `UPDATE` when several columns were referenced from a single table
- Fixed bug in truncating non-existing table.
- Fixed bug in `REVOKE` that caused user resources to be randomly set.
- Fixed bug in `GRANT` for the new `CREATE TEMPORARY TABLE` privilege.
- Fixed bug in multiple-table `DELETE` when tables are re-ordered in the table initialization method and `ref_lengths` are of different sizes.
- Fixed two bugs in `SELECT DISTINCT` with large tables.
- Fixed bug in query cache initialization with very small query cache size.
- Allow `DEFAULT` with `INSERT` statement.
- The startup parameters `myisam_max_sort_file_size` and `myisam_max_extra_sort_file_size` are now given in bytes, not megabytes.
- External system locking of `MyISAM/ISAM` files is now turned off by default. One can turn this on with `--external-locking`. (For most users this is never needed).
- Fixed core dump bug with `INSERT ... SET db_name.tbl_name.col_name= ''`.
- Fixed client hangup bug when using some SQL commands with incorrect syntax.
- Fixed a timing bug in `DROP DATABASE`
- New `SET [GLOBAL | SESSION]` syntax to change thread-specific and global system variables at runtime.
- Added variable `slave_compressed_protocol`.
- Renamed variable `query_cache_startup_type` to `query_cache_type`, `myis-`

`am_bulk_insert_tree_size` to `bulk_insert_buffer_size`, `record_buffer` to `read_buffer_size` and `record_rnd_buffer` to `read_rnd_buffer_size`.

- Renamed some SQL variables, but old names still work until 5.0. See [Section 2.11.2, “Upgrading from MySQL 3.23 to 4.0”](#).
- Renamed `--skip-locking` to `--skip-external-locking`.
- Removed unused variable `query_buffer_size`.
- Fixed a bug that made the pager option in the `mysql` client non-functional.
- Added full `AUTO_INCREMENT` support to `MERGE` tables.
- Extended `LOG()` function to accept an optional arbitrary base parameter. See [Section 12.4.2, “Mathematical Functions”](#).
- Added `LOG2()` function (useful for finding out how many bits a number would require for storage).
- Added `LN()` natural logarithm function for compatibility with other databases. It is synonymous with `LOG(X)`.

## D.2.27. Changes in release 4.0.2 (01 July 2002)

- Cleaned up `NULL` handling for default values in `DESCRIBE tbl_name`.
- Fixed `TRUNCATE()` to round up negative values to the nearest integer.
- Fixed buffer overflow problem if someone specified a too-long `datadir` parameter to `mysqld`. ([CVE-2002-0969](http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2002-0969) [<http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2002-0969>])
- Changed `--chroot=path` option to execute `chroot()` immediately after all options have been parsed.
- Don't allow database names that contain `'\'`.
- `lower_case_table_names` now also applies to database names.
- Added `XOR` operator (logical and bitwise `XOR`) with `^` as a synonym for bitwise `XOR`.
- Added function `IS_FREE_LOCK("lock_name")`. Based on code contributed by Hartmut Holzgraefe <[hartmut@six.de](mailto:hartmut@six.de)>.
- Removed `mysql_ssl_clear()` from C API, as it was not needed.
- `DECIMAL` and `NUMERIC` types can now read exponential numbers.
- Added `SHA1()` function to calculate 160 bit hash value as described in RFC 3174 (Secure Hash Algorithm). This function can be considered a cryptographically more secure equivalent of `MD5()`. See [Section 12.9.2, “Encryption and Compression Functions”](#).
- Added `AES_ENCRYPT()` and `AES_DECRYPT()` functions to perform encryption according to AES standard (Rijndael). See [Section 12.9.2, “Encryption and Compression Functions”](#).
- Added `--single-transaction` option to `mysqldump`, allowing a consistent dump of `InnoDB` tables. See [Section 8.12, “mysqldump — A Database Backup Program”](#).

- Fixed bug in `innodb_log_group_home_dir` in `SHOW VARIABLES`.
- Fixed a bug in optimizer with merge tables when non-unique values are used in summing up (causing crashes).
- Fixed a bug in optimizer when a range specified makes index grouping impossible (causing crashes).
- Fixed a rare bug when `FULLTEXT` index is present and no tables are used.
- Added privileges `CREATE TEMPORARY TABLES`, `EXECUTE`, `LOCK TABLES`, `REPLICATION CLIENT`, `REPLICATION SLAVE`, `SHOW DATABASES` and `SUPER`. To use these, you must have run the `mysql_fix_privilege_tables` script after upgrading.
- Fixed query cache align data bug.
- Fixed mutex bug in replication when reading from master fails.
- Added missing mutex in `TRUNCATE TABLE`. This fixes some core dump/hangup problems when using `TRUNCATE TABLE`.
- Fixed bug in multiple-table `DELETE` when optimizer uses only indexes.
- Fixed that `ALTER TABLE tbl_name RENAME new_tbl_name` is as fast as `RENAME TABLE`.
- Fixed bug in `GROUP BY` with two or more columns, where at least one column can contain `NULL` values.
- Use Turbo Boyer-Moore algorithm to speed up `LIKE "%keyword%"` searches.
- Fixed bug in `DROP DATABASE` with symlink.
- Fixed crash in `REPAIR ... USE_FRM`.
- Fixed bug in `EXPLAIN` with `LIMIT offset != 0`.
- Fixed bug in phrase operator `"..."` in boolean full-text search.
- Fixed bug that caused duplicated rows when using truncation operator `*` in boolean full-text search.
- Fixed bug in truncation operator of boolean full-text search (incorrect results when there are only `+word*s` in the query).
- Fixed bug in boolean full-text search that caused a crash when an identical `MATCH` expression that did not use an index appeared twice.
- Query cache is now automatically disabled in `mysqldump`.
- Fixed problem on Windows 98 that made sending of results very slow.
- Boolean full-text search weighting scheme changed to something more reasonable.
- Fixed bug in boolean full-text search that caused MySQL to ignore queries of `ft_min_word_len` characters.
- Boolean full-text search now supports “phrase searches.”
- New configure option `--without-query-cache`.
- Memory allocation strategy for “root memory” changed. Block size now grows with number of al-

located blocks.

- `INET_NTOA()` now returns `NULL` if you give it an argument that is too large (greater than the value corresponding to `255.255.255.255`).
- Fix `SQL_CALC_FOUND_ROWS` to work with `UNION`. It works only if the first `SELECT` has this option and if there is global `LIMIT` for the entire statement. For the moment, this requires using parentheses for individual `SELECT` queries within the statement.
- Fixed bug in `SQL_CALC_FOUND_ROWS` and `LIMIT`.
- Don't give an error for `CREATE TABLE ... (... VARCHAR(0))`.
- Fixed `SIGINT` and `SIGQUIT` problems in `mysql.cc` on Linux with some `glibc` versions.
- Fixed bug in `convert.cc`, which is caused by having an incorrect `net_store_length()` linked in the `CONVERT::store()` method.
- `DOUBLE` and `FLOAT` columns now honor the `UNSIGNED` flag on storage.
- `InnoDB` now retains foreign key constraints through `ALTER TABLE` and `CREATE/DROP INDEX`.
- `InnoDB` now allows foreign key constraints to be added through the `ALTER TABLE` syntax.
- `InnoDB` tables can now be set to automatically grow in size (auto-extend).
- Added `--ignore-lines=n` option to `mysqlimport`. This has the same effect as the `IGNORE n LINES` clause for `LOAD DATA`.
- Fixed bug in `UNION` with last offset being transposed to total result set.
- `REPAIR ... USE_FRM` added.
- Fixed that `DEFAULT_SELECT_LIMIT` is always imposed on `UNION` result set.
- Fixed that some `SELECT` options can appear only in the first `SELECT`.
- Fixed bug with `LIMIT` with `UNION`, where last select is in the braces.
- Fixed that full-text works fine with `UNION` operations.
- Fixed bug with indexless boolean full-text search.
- Fixed bug that sometimes appeared when full-text search was used with `const` tables.
- Fixed incorrect error value when doing a `SELECT` with an empty `HEAP` table.
- Use `ORDER BY column DESC` now sorts `NULL` values first. (In other words, `NULL` values sort first in all cases, whether or not `DESC` is specified.) This is changed back in 4.0.10.
- Fixed bug in `WHERE key_name='constant' ORDER BY key_name DESC`.
- Fixed bug in `SELECT DISTINCT ... ORDER BY DESC` optimization.
- Fixed bug in `... HAVING 'GROUP_FUNCTION'(xxx) IS [NOT] NULL`.
- Fixed bug in truncation operator for boolean full-text search.
- Allow value of `--user` option for `mysqld` to be specified as a numeric user ID

`user_id`).

- Fixed a bug where `SQL_CALC_ROWS` returned an incorrect value when used with one table and `ORDER BY` and with `InnoDB` tables.
- Fixed that `SELECT 0 LIMIT 0` doesn't hang thread.
- Fixed some problems with `USE/IGNORE INDEX` when using many keys with the same start column.
- Don't use table scan with `BerkeleyDB` and `InnoDB` tables when we can use an index that covers the whole row.
- Optimized `InnoDB` sort-buffer handling to take less memory.
- Fixed bug in multiple-table `DELETE` and `InnoDB` tables.
- Fixed problem with `TRUNCATE` and `InnoDB` tables that produced the error `Can't execute the given command because you have active locked tables or an active transaction`.
- Added `NO_UNSIGNED_SUBTRACTION` to the set of flags that may be specified with the `-sql-mode` option for `mysqld`. It disables unsigned arithmetic rules when it comes to subtraction. (This makes MySQL 4.0 behave more like 3.23 with `UNSIGNED` columns).
- The result returned for all bit functions (`|`, `<<`, ...) is now of type `unsigned integer`.
- Added detection of `nan` values in `MyISAM` to make it possible to repair tables with `nan` in float or double columns.
- Fixed new bug in `myisamchk` where it didn't correctly update number of “parts” in the `MyISAM` index file.
- Changed to use `autoconf 2.52` (from `autoconf 2.13`).
- Fixed optimization problem where the MySQL Server was in “preparing” state for a long time when selecting from an empty table which had contained a lot of rows.
- Fixed bug in complicated join with `const` tables. This fix also improves performance a bit when referring to another table from a `const` table.
- First pre-version of multiple-table `UPDATE` statement.
- Fixed bug in multiple-table `DELETE`.
- Fixed bug in `SELECT CONCAT(argument_list) ... GROUP BY 1`.
- `INSERT ... SELECT` did a full rollback in case of an error. Fixed so that we only roll back the last statement in the current transaction.
- Fixed bug with empty expression for boolean full-text search.
- Fixed core dump bug in updating full-text key from/to `NULL`.
- ODBC compatibility: Added `BIT_LENGTH()` function.
- Fixed core dump bug in `GROUP BY BINARY column`.
- Added support for `NULL` keys in `HEAP` tables.

- Use index for `ORDER BY` in queries of type: `SELECT * FROM t WHERE key_part1=1 ORDER BY key_part1 DESC, key_part2 DESC`
- Fixed bug in `FLUSH QUERY CACHE`.
- Added `CAST()` and `CONVERT()` functions. The `CAST` and `CONVERT` functions are nearly identical and mainly useful when you want to create a column with a specific type in a `CREATE ... SELECT` statement. For more information, read [Section 12.8, “Cast Functions and Operators”](#).
- `CREATE ... SELECT` on `DATE` and `TIME` functions now create columns of the expected type.
- Changed order in which keys are created in tables.
- Added new columns `Null` and `Index_type` to `SHOW INDEX` output.
- Added `--no-beep` and `--prompt` options to `mysql` command-line client.
- New feature: management of user resources.

```
GRANT ... WITH MAX_QUERIES_PER_HOUR N1
             MAX_UPDATES_PER_HOUR N2
             MAX_CONNECTIONS_PER_HOUR N3;
```

See [Section 5.8.4, “Limiting Account Resources”](#).

- Added `mysql_secure_installation` to the `scripts/` directory.

## D.2.28. Changes in release 4.0.1 (23 December 2001)

- Added `system` command to `mysql`.
- Fixed bug when `HANDLER` was used with some unsupported table type.
- `mysqldump` now puts `ALTER TABLE tbl_name DISABLE KEYS` and `ALTER TABLE tbl_name ENABLE KEYS` in the sql dump.
- Added `mysql_fix_extensions` script.
- Fixed stack overrun problem with `LOAD DATA FROM MASTER` on OSF/1.
- Fixed shutdown problem on HP-UX.
- Added `DES_ENCRYPT()` and `DES_DECRYPT()` functions.
- Added `FLUSH DES_KEY_FILE` statement.
- Added `--des-key-file` option to `mysqld`.
- `HEX(str)` now returns the characters in `str` converted to hexadecimal.
- Fixed problem with `GRANT` when using `lower_case_table_names=1`.
- Changed `SELECT ... IN SHARE MODE` to `SELECT ... LOCK IN SHARE MODE` (as in MySQL 3.23).
- A new query cache to cache results from identical `SELECT` queries.



- Fixed core dump bug on 64-bit machines when it got an incorrect communication packet.
- `MATCH ... AGAINST(... IN BOOLEAN MODE)` can now work without `FULLTEXT` index.
- Fixed slave to replicate from 3.23 master.
- Miscellaneous replication fixes/cleanup.
- Got shutdown to work on Mac OS X.
- Added `myisam/ft_dump` utility for low-level inspection of `FULLTEXT` indexes.
- Fixed bug in `DELETE ... WHERE ... MATCH ...`.
- Added support for `MATCH ... AGAINST(... IN BOOLEAN MODE)`. **Note:** You must rebuild your tables with `ALTER TABLE tbl_name TYPE=MyISAM` to be able to use boolean full-text search.
- `LOCATE()` and `INSTR()` are now case sensitive if either argument is a binary string.
- Changed `RAND()` initialization so that `RAND(N)` and `RAND(N+1)` are more distinct.
- Fixed core dump bug in `UPDATE ... ORDER BY`.
- In 3.23, `INSERT INTO ... SELECT` always had `IGNORE` enabled. Now MySQL stops (and possibly rolls back) by default in case of an error unless you specify `IGNORE`.
- Ignore `DATA DIRECTORY` and `INDEX DIRECTORY` directives on Windows.
- Added boolean full-text search code. It should be considered early alpha.
- Extended `MODIFY` and `CHANGE` in `ALTER TABLE` to accept the `FIRST` and `AFTER` keywords.
- Indexes are now used with `ORDER BY` on a whole InnoDB table.

## D.2.29. Changes in release 4.0.0 (October 2001: Alpha)

- Added `--xml` option to `mysql` for producing XML output.
- Added full-text variables `ft_min_word_len`, `ft_max_word_len`, and `ft_max_word_len_for_sort` system variables.
- Added full-text variables `ft_min_word_len`, `ft_max_word_len`, and `ft_max_word_len_for_sort` variables to `myisamchk`.
- Added documentation for `libmysqld`, the embedded MySQL server library. Also added example programs (a `mysql` client and `mysqltest` test program) which use `libmysqld`.
- Removed all Gemini hooks from MySQL server.
- Removed `my_thread_init()` and `my_thread_end()` from `mysql_com.h`, and added `mysql_thread_init()` and `mysql_thread_end()` to `mysql.h`.
- Support for communication packets > 16MB. In 4.0.1 we extended `MyISAM` to be able to handle these.
- Secure connections (with SSL).

- Unsigned `BIGINT` constants now work. `MIN()` and `MAX()` now handle signed and unsigned `BIGINT` numbers correctly.
- New character set `latin1_de` which provides correct German sorting.
- `STRCMP()` now uses the current character set when doing comparisons, which means that the default comparison behavior now is case insensitive.
- `TRUNCATE TABLE` and `DELETE FROM tbl_name` are now separate functions. One bonus is that `DELETE FROM tbl_name` now returns the number of deleted rows, rather than zero.
- `DROP DATABASE` now executes a `DROP TABLE` on all tables in the database, which fixes a problem with `InnoDB` tables.
- Added support for `UNION`.
- Added support for multiple-table `DELETE` operations.
- A new `HANDLER` interface to `MyISAM` tables.
- Added support for `INSERT` on `MERGE` tables. Patch from Benjamin Pflugmann.
- Changed `WEEK(date, 0)` to match the calendar in the USA.
- `COUNT(DISTINCT)` is about 30% faster.
- Speed up all internal list handling.
- Speed up `IS NULL`, `ISNULL()` and some other internal primitives.
- Full-text index creation now is much faster.
- Tree-like cache to speed up bulk inserts and `myisam_bulk_insert_tree_size` variable.
- Searching on packed (`CHAR/VARCHAR`) keys is now much faster.
- Optimized queries of type: `SELECT DISTINCT * FROM tbl_name ORDER by key_part1 LIMIT row_count`.
- `SHOW CREATE TABLE` now shows all table attributes.
- `ORDER BY ... DESC` can now use keys.
- `LOAD DATA FROM MASTER` “automatically” sets up a slave.
- Renamed `safe_mysqld` to `mysqld_safe` to make this name more in line with other MySQL scripts/commands.
- Added support for symbolic links to `MyISAM` tables. Symlink handling is now enabled by default for Windows.
- Added `SQL_CALC_FOUND_ROWS` and `FOUND_ROWS()`. This makes it possible to know how many rows a query would have returned without a `LIMIT` clause.
- Changed output format of `SHOW OPEN TABLES`.
- Allow `SELECT expression LIMIT ....`
- Added the `IDENTITY` variable as a synonym for the `LAST_INSERT_ID` variable (like Sybase).

- Added `ORDER BY` syntax to `UPDATE` and `DELETE`.
- `SHOW INDEXES` is now a synonym for `SHOW INDEX`.
- Added `ALTER TABLE tbl_name DISABLE KEYS` and `ALTER TABLE tbl_name ENABLE KEYS` commands.
- Allow use of `IN` as a synonym for `FROM` in `SHOW` commands.
- Implemented “repair by sort” for `FULLTEXT` indexes. `REPAIR TABLE`, `ALTER TABLE`, and `OPTIMIZE TABLE` for tables with `FULLTEXT` indexes are now up to 100 times faster.
- Allow standard SQL syntax `X'hexadecimal-number'`.
- Cleaned up global lock handling for `FLUSH TABLES WITH READ LOCK`.
- Fixed problem with `DATETIME = constant` in `WHERE` optimization.
- Added `--master-data` and `--no-autocommit` options to `mysqldump`. (Thanks to Brian Aker for this.)
- Added script `mysql_explain_log.sh` to distribution. (Thanks to mobile.de).

## D.3. Changes in release 3.23.x (Recent; still supported)

Please note that since release 4.0 is now production level, only critical fixes are done in the 3.23 release series. You are recommended to upgrade when possible, to take advantage of all speed and feature improvements in 4.0. See [Section 2.11.2, “Upgrading from MySQL 3.23 to 4.0”](#).

The 3.23 release has several major features that are not present in previous versions. We have added three new table types:

- [MyISAM](#)  
A new [ISAM](#) library which is tuned for SQL and supports large files.
- [InnoDB](#)  
A transaction-safe storage engine that supports row level locking, and many Oracle-like features.
- [BerkeleyDB](#) or [BDB](#)  
Uses the Berkeley DB library from Sleepycat Software to implement transaction-safe tables.

Note that only [MyISAM](#) is available in the standard binary distribution.

The 3.23 release also includes support for database replication between a master and many slaves, full-text indexing, and much more.

All new features are being developed in the 4.1.x and 5.0.x versions. Only serious bugfixes are added to 3.23.

### D.3.1. Changes in release 3.23.59 (Not yet released)

- Fixed an old bug in concurrent accesses to `MERGE` tables (even one `MERGE` table and `MyISAM` tables), that could've resulted in a crash or hang of the server. (Bug#2408 [http://bugs.mysql.com/2408], CVE-2004-0837 [http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2004-0837])
- Fixed incorrect destruction of expression which led to crash of server on complex `AND/OR` expressions if query was ignored (either by a replication server because of `--replicate-*-table` rules, or by any MySQL server because of a syntax error). (Bug#3969 [http://bugs.mysql.com/3969], Bug#4494 [http://bugs.mysql.com/4494])
- Fixed problem with parsing complex queries on 64-bit architectures. (Bug#4204 [http://bugs.mysql.com/4204])
- Fixed a symlink vulnerability in the `mysqlbug` script. (Bug#3284 [http://bugs.mysql.com/3284], CVE-2004-0381 [http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2004-0381])
- Fixed bug in privilege checking of `ALTER TABLE RENAME`. (Bug#3270 [http://bugs.mysql.com/3270], CVE-2004-0835 [http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2004-0835])
- Fixed bugs in `ACOS()`, `ASIN()` (Bug#2338 [http://bugs.mysql.com/2338]) and in `FLOOR()` (Bug#3051 [http://bugs.mysql.com/3051]). The cause of the problem is an overly strong optimization done by `gcc` in this case.
- Fixed bug in `INSERT ... SELECT` statements where, if a `NOT NULL` column is assigned a value of `NULL`, the following columns in the row might be assigned a value of zero. (Bug#2012 [http://bugs.mysql.com/2012])
- If a query was ignored on the slave (because of `--replicate-ignore-table` and other similar rules), the slave still checked if the query got the same error code (0, no error) as on the master. So if the master had an error on the query (for example, "Duplicate entry" in a multiple-row insert), then the slave stopped and warned that the error codes didn't match. This is a backport of the fix for MySQL 4.0. (Bug#797 [http://bugs.mysql.com/797])
- `mysqlbinlog` now asks for a password at console when the `-p/--password` option is used with no argument. This is how the other clients (`mysqladmin`, `mysqldump`..) behave. Note that one now has to use `mysqlbinlog -p<my_password>`; `mysqlbinlog -p <my_password>` does not work anymore (in other words, put no space after `-p`). (Bug#1595 [http://bugs.mysql.com/1595])
- On some 64-bit machines (some HP-UX and Solaris machines), a slave installed with the 64-bit MySQL binary could not connect to its master (it connected to itself instead). (Bug#1256 [http://bugs.mysql.com/1256], Bug#1381 [http://bugs.mysql.com/1381])
- Fixed a Windows-specific bug present since MySQL 3.23.57 and 3.23.58 that caused Windows slaves to crash when they started replication if a `master.info` file existed. (Bug#1720 [http://bugs.mysql.com/1720])
- Fixed bug in `ALTER TABLE RENAME`, when rename to the table with the same name in another database silently dropped destination table if it existed. (Bug#2628 [http://bugs.mysql.com/2628])
- Fixed potential memory overrun in `mysql_real_connect()` (which required a compromised DNS server and certain operating systems). (Bug#4017 [http://bugs.mysql.com/4017], CVE-2004-0836 [http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2004-0836])

## D.3.2. Changes in release 3.23.58 (11 September 2003)

- Fixed buffer overflow in password handling which could potentially be exploited by MySQL users with `ALTER` privilege on the `mysql.user` table to execute random code or to gain shell access with the UID of the `mysqld` process (thanks to Jedi/Sector One for spotting and reporting this bug).
- `mysqldump` now correctly quotes all identifiers when communicating with the server. This assures that during the dump process, `mysqldump` never sends queries to the server that result in a syntax error. This problem is **not** related to the `mysqldump` program's output, which was not changed. (Bug#1148 [<http://bugs.mysql.com/1148>])
- Fixed table/column grant handling: The proper sort order (from most specific to less specific, see [Section 5.7.6, “Access Control, Stage 2: Request Verification”](#)) was not honored. (Bug#928 [<http://bugs.mysql.com/928>])
- Fixed overflow bug in `MyISAM` and `ISAM` when a row is updated in a table with a large number of columns and at least one `BLOB/TEXT` column.
- Fixed MySQL so that field length (in C API) for the second column in `SHOW CREATE TABLE` is always larger than the data length. The only known application that was affected by the old behavior was Borland dbExpress, which truncated the output from the command. (Bug#1064 [<http://bugs.mysql.com/1064>])
- Fixed `ISAM` bug in `MAX( )` optimization.
- Fixed `Unknown error` when doing `ORDER BY` on reference table which was used with `NULL` value on `NOT NULL` column. (Bug#479 [<http://bugs.mysql.com/479>])

### D.3.3. Changes in release 3.23.57 (06 June 2003)

- Fixed problem in alarm handling that could cause problems when getting a packet that is too large.
- Fixed problem when installing MySQL as a service on Windows when two arguments were specified to `mysqld` (option file group name and service name).
- Fixed `kill pid-of-mysqld` to work on Mac OS X.
- `SHOW TABLE STATUS` displayed incorrect `Row_format` value for tables that have been compressed with `myisampack`. (Bug#427 [<http://bugs.mysql.com/427>])
- `SHOW VARIABLES LIKE 'innodb_data_file_path'` displayed only the name of the first data file. (Bug#468 [<http://bugs.mysql.com/468>])
- Fixed security problem where `mysqld` didn't allow one to `UPDATE` rows in a table even if one had a global `UPDATE` privilege and a database `SELECT` privilege.
- Fixed a security problem with `SELECT` and wildcard select list, when user only had partial column `SELECT` privileges on the table.
- Fixed unlikely problem in optimizing `WHERE` clause with a constant expression such as in `WHERE 1 AND (a=1 AND b=1)`.
- Fixed problem on IA-64 with timestamps that caused `mysqlbinlog` to fail.
- The default option for `innodb_flush_log_at_trx_commit` was changed from 0 to 1 to make `InnoDB` tables ACID by default. See [Section 14.2.5, “InnoDB Startup Options and System Variables”](#).
- Fixed problem with too many allocated alarms on slave when connecting to master many times

(normally not a very critical error).

- Fixed a bug in replication of temporary tables. ([Bug#183](http://bugs.mysql.com/183) [<http://bugs.mysql.com/183>])
- Fixed 64-bit bug that affected at least AMD hammer systems.
- Fixed a bug when doing `LOAD DATA INFILE IGNORE`: When reading the binary log, `mysql-binlog` and the replication code read `REPLACE` instead of `IGNORE`. This could make the slave's table become different from the master's table. ([Bug#218](http://bugs.mysql.com/218) [<http://bugs.mysql.com/218>])
- Fixed overflow bug in `MyISAM` when a row is inserted into a table with a large number of columns and at least one `BLOB/TEXT` column. Bug was caused by incorrect calculation of the needed buffer to pack data.
- The binary log was not locked during `TRUNCATE tbl_name` or `DELETE FROM tbl_name` statements, which could cause an `INSERT` to `tbl_name` to be written to the log before the `TRUNCATE` or `DELETE` statements.
- Fixed rare bug in `UPDATE` of `InnoDB` tables where one row could be updated multiple times.
- Produce an error for empty table and column names.
- Changed `PROCEDURE ANALYSE ( )` to report `DATE` instead of `NEWDATE`.
- Changed `PROCEDURE ANALYSE (#)` to restrict the number of values in an `ENUM` column to `#` also for string values.
- `mysqldump` no longer silently deletes the binary logs when invoked with the `--master-data` or `--first-slave` option; while this behavior was convenient for some users, others may suffer from it. Now you must explicitly ask for binary logs to be deleted by using the new `--delete-master-logs` option.
- Fixed a bug in `mysqldump` when it was invoked with the `--master-data` option: The `CHANGE MASTER TO` statements that were appended to the SQL dump had incorrect coordinates. ([Bug#159](http://bugs.mysql.com/159) [<http://bugs.mysql.com/159>])

### D.3.4. Changes in release 3.23.56 (13 March 2003)

- Fixed `mysqld` crash on extremely small values of `sort_buffer` variable.
- Fixed a bug in privilege system for `GRANT UPDATE` on the column level.
- Fixed a rare bug when using a date in `HAVING` with `GROUP BY`.
- Fixed checking of random part of `WHERE` clause. ([Bug#142](http://bugs.mysql.com/142) [<http://bugs.mysql.com/142>])
- Fixed MySQL (and `myisamchk`) crash on artificially corrupted `.MYI` files.
- Security enhancement: `mysqld` no longer reads options from world-writable config files. ([CVE-2003-0150](http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2003-0150) [<http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2003-0150>])
- Security enhancement: `mysqld` and `safe_mysqld` now use only the first `--user` option specified on the command line. Normally this comes from `/etc/my.cnf`. ([CVE-2003-0150](http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2003-0150) [<http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2003-0150>])
- Security enhancement: Don't allow `BACKUP TABLE` to overwrite existing files.

- Fixed unlikely deadlock bug when one thread did a `LOCK TABLE` and another thread did a `DROP TABLE`. In this case one could do a `KILL` on one of the threads to resolve the deadlock.
- `LOAD DATA INFILE` was not replicated by slave if `replicate_*_table` was set on the slave.
- Fixed a bug in handling `CHAR(0)` columns that could cause incorrect results from the query.
- Fixed a bug in `SHOW VARIABLES` on 64-bit platforms. The bug was caused by incorrect declaration of variable `server_id`.
- The `Comment` column in `SHOW TABLE STATUS` now reports that it can contain `NULL` values (which is the case for a crashed `.frm` file).
- Fixed the `rpl_rotate_logs` test to not fail on certain platforms (such as Mac OS X) due to a too-long file name (changed `slave-master-info.opt` to `.slave-mi`).
- Fixed a problem with `BLOB NOT NULL` columns used with `IS NULL`.
- Fixed bug in `MAX()` optimization in `MERGE` tables.
- Better `RAND()` initialization for new connections.
- Fixed bug with connect timeout. This bug was manifested on OS's with `poll()` system call, which resulted in timeout the value specified as it was executed in both `select()` and `poll()`.
- Fixed bug in `SELECT * FROM table WHERE datetime1 IS NULL OR datetime2 IS NULL`.
- Fixed bug in using aggregate functions as argument for `INTERVAL`, `CASE`, `FIELD`, `CONCAT_WS`, `ELT` and `MAKE_SET` functions.
- When running with `--lower-case-table-names=1` (default on Windows) and you had tables or databases with mixed case on disk, then executing `SHOW TABLE STATUS` followed with `DROP DATABASE` or `DROP TABLE` could fail with `Errcode 13`.
- Fixed bug in logging to binary log (which affects replication) a query that inserts a `NULL` in an `auto_increment` field and also uses `LAST_INSERT_ID()`.
- Fixed bug in `mysqladmin --relative`.
- On some 64-bit systems, `show status` reported a strange number for `Open_files` and `Open_streams`.

### D.3.5. Changes in release 3.23.55 (23 January 2003)

- Fixed double `free`'d pointer bug in `mysql_change_user()` handling, that enabled a specially hacked version of MySQL client to crash `mysqld`. **Note** that you must log in to the server by using a valid user account to be able to exploit this bug. (CVE-2003-0073 [http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2003-0073])
- Fixed bug with the `--slow-log` when logging an administrator command (like `FLUSH TABLES`).
- Fixed bug in `GROUP BY` when used on `BLOB` column with `NULL` values.
- Fixed a bug in handling `NULL` values in `CASE ... WHEN ...`.

- Bugfix for `--chroot` (see [Section D.3.6, “Changes in release 3.23.54 \(05 December 2002\)”](#)) is reverted. Unfortunately, there is no way to make it to work, without introducing backward-incompatible changes in `my.cnf`. Those who need `--chroot` functionality, should upgrade to MySQL 4.0. (The fix in the 4.0 branch did not break backward-compatibility).
- Make `--lower-case-table-names` default on Mac OS X as the default filesystem (HFS+) is case insensitive.
- Fixed a bug in `scripts/mysql_safe.sh` in `NOHUP_NICENESS` testing.
- Transactions in `AUTOCOMMIT=0` mode didn't rotate binary log.
- Fixed a bug in `scripts/make_binary_distribution` that resulted in a remaining `@HOSTNAME@` variable instead of replacing it with the correct path to the `hostname` binary.
- Fixed a very unlikely bug that could cause `SHOW PROCESSLIST` to core dump in `pthread_mutex_unlock()` if a new thread was connecting.
- Forbid `SLAVE STOP` if the thread executing the query has locked tables. This removes a possible deadlock situation.

### D.3.6. Changes in release 3.23.54 (05 December 2002)

- Fixed a bug, that allowed to crash `mysqld` with a specially crafted packet. ([CVE-2002-1373](#) [<http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2002-1373>])
- Fixed a rare crash (double `free`'d pointer) when altering a temporary table.
- Fixed buffer overrun in `libmysqlclient` library that allowed malicious MySQL server to crash the client application. ([CVE-2002-1376](#) [<http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2002-1376>])
- Fixed security-related bug in `mysql_change_user()` handling. All users are strongly recommended to upgrade to the version 3.23.54. ([CVE-2002-1374](#) [<http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2002-1374>], [CVE-2002-1375](#) [<http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2002-1375>])
- Fixed bug that prevented `--chroot` command-line option of `mysqld` from working.
- Fixed bug that made `OPTIMIZE TABLE` to corrupt the table under some rare circumstances.
- Fixed `mysqlcheck` so it can deal with table names containing dashes.
- Fixed shutdown problem on Mac OS X.
- Fixed bug with comparing an indexed `NULL` field with `<=> NULL`.
- Fixed bug that caused `IGNORE INDEX` and `USE INDEX` sometimes to be ignored.
- Fixed rare core dump problem in complicated `GROUP BY` queries that didn't return any result.
- Fixed a bug where `MATCH ... AGAINST ( ) >=0` was treated as if it was `>`.
- Fixed core dump in `SHOW PROCESSLIST` when running with an active slave (unlikely timing bug).
- Make it possible to use multiple MySQL servers on Windows (code backported from 4.0.2).



- One can create `TEMPORARY MERGE` tables now.
- Fixed that `--core-file` works on Linux (at least on kernel 2.4.18).
- Fixed a problem with `BDB` and `ALTER TABLE`.
- Fixed reference to freed memory when doing complicated `GROUP BY ... ORDER BY` queries. Symptom was that `mysqld` died in function `send_fields`.
- Allocate heap rows in smaller blocks to get better memory usage.
- Fixed memory allocation bug when storing `BLOB` values in internal temporary tables used for some (unlikely) `GROUP BY` queries.
- Fixed a bug in key optimizing handling where the expression `WHERE col_name = key_col_name` was calculated as true for `NULL` values.
- Fixed core dump bug when doing `LEFT JOIN ... WHERE key_column=NULL`.
- Fixed `MyISAM` crash when using dynamic-row tables with huge numbers of packed fields.
- Updated source tree to be built using `automake` 1.5 and `libtool` 1.4.

### D.3.7. Changes in release 3.23.53 (09 October 2002)

- Fixed crash when `SHOW INNODB STATUS` was used and `skip-innodb` was defined.
- Fixed possible memory corruption bug in binary log file handling when slave rotated the logs (only affected 3.23, not 4.0).
- Fixed problem in `LOCK TABLES` on Windows when one connects to a database that contains uppercase letters.
- Fixed that `--skip-show-database` doesn't reset the `--port` option.
- Small fix in `safe_mysqld` for some shells.
- Fixed that `FLUSH STATUS` doesn't reset `delayed_insert_threads`.
- Fixed core dump bug when using the `BINARY` cast on a `NULL` value.
- Fixed race condition when someone did a `GRANT` at the same time a new user logged in or did a `USE database`.
- Fixed bug in `ALTER TABLE` and `RENAME TABLE` when running with `-O lower_case_table_names=1` (typically on Windows) when giving the table name in uppercase.
- Fixed that `-O lower_case_table_names=1` also converts database names to lowercase.
- Fixed unlikely core dump with `SELECT ... ORDER BY ... LIMIT`.
- Changed `AND/OR` to report that they can return `NULL`. This fixes a bug in `GROUP BY` on `AND/OR` expressions that return `NULL`.
- Fixed a bug that `OPTIMIZE TABLE` of locked and modified `MyISAM` table, reported table corruption.

- Fixed a BDB-related `ALTER TABLE` bug with dropping a column and shutting down immediately thereafter.
- Fixed problem with `configure ... --localstatedir=...`
- Fixed problem with `UNSIGNED BIGINT` on AIX (again).
- Fixed bug in `pthread_mutex_trylock()` on HP-UX 11.0.
- Multi-threaded stress tests for `InnoDB`.

### D.3.8. Changes in release 3.23.52 (14 August 2002)

- Wrap `BEGIN/COMMIT` around transaction in the binary log. This makes replication honor transactions.
- Fixed security bug when having an empty database name in the `user.db` table.
- Changed initialization of `RAND()` to make it less predictable.
- Fixed problem with `GROUP BY` on result with expression that created a `BLOB` field.
- Fixed problem with `GROUP BY` on columns that have `NULL` values. To solve this we now create an `MyISAM` temporary table when doing a `GROUP BY` on a possible `NULL` item. From MySQL 4.0.5 we can use in memory `HEAP` tables for this case.
- Fixed problem with privilege tables when downgrading from 4.0.2 to 3.23.
- Fixed thread bug in `SLAVE START`, `SLAVE STOP` and automatic repair of `MyISAM` tables that could cause table cache to be corrupted.
- Fixed possible thread related key-cache-corruption problem with `OPTIMIZE TABLE` and `REPAIR TABLE`.
- Added name of 'administrator command' logs.
- Fixed bug with creating an auto-increment value on second part of a `UNIQUE()` key where first part could contain `NULL` values.
- Don't write slave-timeout reconnects to the error log.
- Fixed bug with slave net read timeouting
- Fixed a core-dump bug with `MERGE` tables and `MAX()` function.
- Fixed bug in `ALTER TABLE` with `BDB` tables.
- Fixed bug when logging `LOAD DATA INFILE` to binary log with no active database.
- Fixed a bug in range optimizer (causing crashes).
- Fixed possible problem in replication when doing `DROP DATABASE` on a database with `InnoDB` tables.
- Fixed `mysql_info()` to return 0 for `Duplicates` value when using `INSERT DELAYED IGNORE`.

- Added `-DHAVE_BROKEN_REALPATH` to the Mac OS X (darwin) compile options in `configure.in` to fix a failure under high load.

### D.3.9. Changes in release 3.23.51 (31 May 2002)

- Fix bug with closing tags missing slash for `mysqldump` XML output.
- Remove endspace from `ENUM` values. (This fixed a problem with `SHOW CREATE TABLE`.)
- Fixed bug in `CONCAT_WS()` that cut the result.
- Changed name of server variables `Com_show_master_stat` to `Com_show_master_status` and `Com_show_slave_stat` to `Com_show_slave_status`.
- Changed handling of `gethostbyname()` to make the client library thread-safe even if `gethostbyname_r` doesn't exist.
- Fixed core-dump problem when giving a wrong password string to `GRANT`.
- Fixed bug in `DROP DATABASE` with symlinked directory.
- Fixed optimization problem with `DATETIME` and value outside `DATETIME` range.
- Removed Sleepycat's `BDB` doc files from the source tree, as they're not needed (MySQL covers `BDB` in its own documentation).
- Fixed MIT-pthreads to compile with `glibc 2.2` (needed for `make dist`).
- Fixed the `FLOAT(X+1, X)` is not converted to `FLOAT(X+2, X)`. (This also affected `DECIMAL`, `DOUBLE` and `REAL` types)
- Fixed the result from `IF()` is case in-sensitive if the second and third arguments are case sensitive.
- Fixed core dump problem on OSF/1 in `gethostbyname_r`.
- Fixed that underflowed decimal fields are not zero filled.
- If we get an overflow when inserting `'+11111'` for `DECIMAL(5,0) UNSIGNED` columns, we just drop the sign.
- Fixed optimization bug with `ISNULL(expression_which_cannot_be_null)` and `ISNULL(constant_expression)`.
- Fixed host lookup bug in the `glibc` library that we used with the 3.23.50 Linux-x86 binaries.

### D.3.10. Changes in release 3.23.50 (21 April 2002)

- Fixed buffer overflow problem if someone specified a too-long `datadir` parameter to `mysqld`. (CVE-2002-0969 [<http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2002-0969>])
- Add missing `<row>` tags for `mysqldump` XML output.
- Fixed problem with `crash-me` and `gcc 3.0.4`.
- Fixed that `@@unknown_variable` doesn't hang server.

- Added `@@VERSION` as a synonym for `VERSION()`.
- `SHOW VARIABLES LIKE 'xxx'` is now case-insensitive.
- Fixed timeout for `GET_LOCK()` on HP-UX with DCE threads.
- Fixed memory allocation bug in the glibc library used to build Linux binaries, which caused `mysqld` to die in `free()`.
- Fixed `SIGINT` and `SIGQUIT` problems in `mysql`.
- Fixed bug in character table converts when used with big (larger than 64KB) strings.
- `InnoDB` now retains foreign key constraints through `ALTER TABLE` and `CREATE/DROP INDEX`.
- `InnoDB` now allows foreign key constraints to be added through the `ALTER TABLE` syntax.
- `InnoDB` tables can now be set to automatically grow in size (auto-extend).
- Our Linux RPMS and binaries are now compiled with `gcc 3.0.4`, which should make them a bit faster.
- Fixed some buffer overflow problems when reading startup parameters.
- Because of problems on shutdown we have now disabled named pipes on Windows by default. One can enable named pipes by starting `mysqld` with `--enable-named-pipe`.
- Fixed bug when using `WHERE key_column = 'J' or key_column='j'`.
- Fixed core-dump bug when using `--log-bin` with `LOAD DATA INFILE` without an active database.
- Fixed bug in `RENAME TABLE` when used with `lower_case_table_names=1` (default on Windows).
- Fixed unlikely core-dump bug when using `DROP TABLE` on a table that was in use by a thread that also used queries on only temporary tables.
- Fixed problem with `SHOW CREATE TABLE` and `PRIMARY KEY` when using 32 indexes.
- Fixed that one can use `SET PASSWORD` for the anonymous user.
- Fixed core dump bug when reading client groups from option files using `mysql_options()`.
- Memory leak (16 bytes per every **corrupted** table) closed.
- Fixed binary builds to use `--enable-local-infile`.
- Update source to work with new version of `bison`.
- Updated shell scripts to now agree with new POSIX standard.
- Fixed bug where `DATE_FORMAT()` returned empty string when used with `GROUP BY`.

### D.3.11. Changes in release 3.23.49 (14 February 2002)

- For a `MERGE` table, `DELETE FROM merge_table` used without a `WHERE` clause no longer clears the mapping for the table by emptying the `.MRG` file. Instead, it deletes records from the mapped tables.
- Don't give warning for a statement that is only a comment; this is needed for `mysqldump --disable-keys` to work.
- Fixed unlikely caching bug when doing a join without keys. In this case, the last used field for a table always returned `NULL`.
- Added options to make `LOAD DATA LOCAL INFILE` more secure.
- MySQL binary release 3.23.48 for Linux contained a new `glibc` library, which has serious problems under high load and Red Hat 7.2. The 3.23.49 binary release doesn't have this problem.
- Fixed shutdown problem on NT.

### D.3.12. Changes in release 3.23.48 (07 February 2002)

- Added `--xml` option to `mysqldump` for producing XML output.
- Changed to use `autoconf 2.52` (from `autoconf 2.13`)
- Fixed bug in complicated join with `const` tables.
- Added internal safety checks for `InnoDB`.
- Some `InnoDB` variables were always shown in `SHOW VARIABLES` as `OFF` on high-byte-first systems (like SPARC).
- Fixed problem with one thread using an `InnoDB` table and another thread doing an `ALTER TABLE` on the same table. Before that, `mysqld` could crash with an assertion failure in `row0row.c`, line 474.
- Tuned the `InnoDB` SQL optimizer to favor index searches more often over table scans.
- Fixed a performance problem with `InnoDB` tables when several large `SELECT` queries are run concurrently on a multiprocessor Linux computer. Large CPU-bound `SELECT` queries now also generally run faster on all platforms.
- If MySQL binary logging is used, `InnoDB` now prints after crash recovery the latest MySQL binary log name and the offset `InnoDB` was able to recover to. This is useful, for example, when resynchronizing a master and a slave database in replication.
- Added better error messages to help in installation problems of `InnoDB` tables.
- It is now possible to recover MySQL temporary tables that have become orphaned inside the `InnoDB` tablespace.
- `InnoDB` now prevents a `FOREIGN KEY` declaration where the signedness is not the same in the referencing and referenced integer columns.
- Calling `SHOW CREATE TABLE` or `SHOW TABLE STATUS` could cause memory corruption and make `mysqld` crash. Especially at risk was `mysqldump`, because it frequently calls `SHOW CREATE TABLE`.
- If inserts to several tables containing an `AUTO_INCREMENT` column were wrapped inside one

`LOCK TABLES`, InnoDB asserted in `lock0lock.c`.

- In 3.23.47 we allowed several `NULL` values in a `UNIQUE` secondary index for an InnoDB table. But `CHECK TABLE` was not relaxed: it reports the table as corrupt. `CHECK TABLE` no longer complains in this situation.
- `SHOW GRANTS` now shows `REFERENCES` instead of `REFERENCE`.

### D.3.13. Changes in release 3.23.47 (27 December 2001)

- Fixed bug when using the following construct: `SELECT ... WHERE key=@var_name OR key=@var_name2`
- Restrict InnoDB keys to 500 bytes.
- InnoDB now supports `NULL` in keys.
- Fixed shutdown problem on HP-UX. (Introduced in 3.23.46)
- Fixed core dump bug in replication when using `SELECT RELEASE_LOCK()`.
- Added new statement: `DO expr[,expr]...`
- Added `slave-skip-errors` option.
- Added statistics variables for all MySQL commands. (`SHOW STATUS` is now much longer.)
- Fixed default values for InnoDB tables.
- Fixed that `GROUP BY expr DESC` works.
- Fixed bug when using `t1 LEFT JOIN t2 ON t2.key=constant`.
- `mysql_config` now also works with binary (relocated) distributions.

### D.3.14. Changes in release 3.23.46 (29 November 2001)

- Fixed problem with aliased temporary table replication.
- InnoDB and BDB tables now use index when doing an `ORDER BY` on the whole table.
- Fixed bug where one got an empty set instead of a DEADLOCK error when using BDB tables.
- One can now kill `ANALYZE TABLE`, `REPAIR TABLE`, and `OPTIMIZE TABLE` when the thread is waiting to get a lock on the table.
- Fixed race condition in `ANALYZE TABLE`.
- Fixed bug when joining with caching (unlikely to happen).
- Fixed race condition when using the binary log and `INSERT DELAYED` which could cause the binary log to have rows that were not yet written to MyISAM tables.
- Changed caching of binary log to make replication slightly faster.

- Fixed bug in replication on Mac OS X.

### D.3.15. Changes in release 3.23.45 (22 November 2001)

- `(UPDATE | DELETE) ... WHERE MATCH` bugfix.
- shutdown should now work on Darwin (Mac OS X).
- Fixed core dump when repairing corrupted packed `MyISAM` files.
- `--core-file` now works on Solaris.
- Fix a bug which could cause `InnoDB` to complain if it cannot find free blocks from the buffer cache during recovery.
- Fixed bug in `InnoDB` insert buffer B-tree handling that could cause crashes.
- Fixed bug in `InnoDB` lock timeout handling.
- Fixed core dump bug in `ALTER TABLE` on a `TEMPORARY InnoDB` table.
- Fixed bug in `OPTIMIZE TABLE` that reset index cardinality if it was up to date.
- Fixed problem with `t1 LEFT JOIN t2 ... WHERE t2.date_column IS NULL` when `date_column` was declared as `NOT NULL`.
- Fixed bug with `BDB` tables and keys on `BLOB` columns.
- Fixed bug in `MERGE` tables on OS with 32-bit file pointers.
- Fixed bug in `TIME_TO_SEC()` when using negative values.

### D.3.16. Changes in release 3.23.44 (31 October 2001)

- Fixed `Rows_examined` count in slow query log.
- Fixed bug when using a reference to an `AVG()` column in `HAVING`.
- Fixed that date functions that require correct dates, like `DAYOFYEAR(column)`, return `NULL` for `0000-00-00` dates.
- Fixed bug in const-propagation when comparing columns of different types. (`SELECT * FROM date_col="2001-01-01" and date_col=time_col`)
- Fixed bug that caused error message `Can't write, because of unique constraint` with some `GROUP BY` queries.
- Fixed problem with `sjis` character strings used within quoted table names.
- Fixed core dump when using `CREATE ... FULLTEXT` keys with other storage engines than `MyISAM`.
- Don't use `signal()` on Windows because this appears to not be 100% reliable.
- Fixed bug when doing `WHERE col_name=NULL` on an indexed column that had `NULL` values.

- Fixed bug when doing `LEFT JOIN ... ON (col_name = constant) WHERE col_name = constant`.
- When using replications, aborted queries that contained `%` could cause a core dump.
- `TCP_NODELAY` was not used on some systems. (Speed problem.)
- Applied portability fixes for OS/2. (Patch by Yuri Dario.)

The following changes are for `InnoDB` tables:

- Add missing `InnoDB` variables to `SHOW VARIABLES`.
- Foreign key checking is now done for `InnoDB` tables.
- `DROP DATABASE` now works also for `InnoDB` tables.
- `InnoDB` now supports data files and raw disk partitions bigger than 4GB on those operating systems that have big files.
- `InnoDB` calculates better table cardinality estimates for the MySQL optimizer.
- Accent characters in the default character set `latin1` are ordered according to the MySQL ordering.

Note: If you are using `latin1` and have inserted characters whose code is greater than 127 into an indexed `CHAR` column, you should run `CHECK TABLE` on your table when you upgrade to 3.23.44, and drop and reimport the table if `CHECK TABLE` reports an error!

- A new `my.cnf` parameter, `innodb_thread_concurrency`, helps in performance tuning in heavily concurrent environments.
- A new `my.cnf` parameter, `innodb_fast_shutdown`, speeds up server shutdown.
- A new `my.cnf` parameter, `innodb_force_recovery`, helps to save your data in case the disk image of the database becomes corrupt.
- `innodb_monitor` has been improved and a new `innodb_table_monitor` added.
- Increased maximum key length from 500 to 7000 bytes.
- Fixed a bug in replication of `AUTO_INCREMENT` columns with multiple-line inserts.
- Fixed a bug when the case of letters changes in an update of an indexed secondary column.
- Fixed a hang when there are more than 24 data files.
- Fixed a crash when `MAX(col)` is selected from an empty table, and `col` is not the first column in a multi-column index.
- Fixed a bug in purge which could cause crashes.

### D.3.17. Changes in release 3.23.43 (04 October 2001)

- Fixed a bug in `INSERT DELAYED` and `FLUSH TABLES` introduced in 3.23.42.



- Fixed unlikely bug, which returned non-matching rows, in `SELECT` with many tables and multi-column indexes and 'range' type.
- Fixed an unlikely core dump bug when doing `EXPLAIN SELECT` when using many tables and `ORDER BY`.
- Fixed bug in `LOAD DATA FROM MASTER` when using table with `CHECKSUM=1`.
- Added unique error message when a DEADLOCK occurs during a transaction with `BDB` tables.
- Fixed problem with `BDB` tables and `UNIQUE` columns defined as `NULL`.
- Fixed problem with `myisampack` when using pre-space filled `CHAR` columns.
- Applied patch from Yuri Dario for OS/2.
- Fixed bug in `--safe-user-create`.

### D.3.18. Changes in release 3.23.42 (08 September 2001)

- Fixed problem when using `LOCK TABLES` and `BDB` tables.
- Fixed problem with `REPAIR TABLE` on `MyISAM` tables with row lengths in the range from 65517 to 65520 bytes.
- Fixed rare hang when doing `mysqladmin shutdown` when there was a lot of activity in other threads.
- Fixed problem with `INSERT DELAYED` where delayed thread could be hanging on `upgrading locks` for no apparent reason.
- Fixed problem with `myisampack` and `BLOB`.
- Fixed problem when one edited `.MRG` tables by hand. (Patch from Benjamin Pflugmann).
- Enforce that all tables in a `MERGE` table come from the same database.
- Fixed bug with `LOAD DATA INFILE` and transactional tables.
- Fix bug when using `INSERT DELAYED` with wrong column definition.
- Fixed core dump during `REPAIR TABLE` of some particularly broken tables.
- Fixed bug in `InnoDB` and `AUTO_INCREMENT` columns.
- Fixed bug in `InnoDB` and `RENAME TABLE` columns.
- Fixed critical bug in `InnoDB` and `BLOB` columns. If you have used `BLOB` columns larger than 8000 bytes in an `InnoDB` table, it is necessary to dump the table with `mysqldump`, drop it and restore it from the dump.
- Applied large patch for OS/2 from Yuri Dario.
- Fixed problem with `InnoDB` when one could get the error `Can't execute the given command...` even when no transaction was active.
- Applied some minor fixes that concern Gemini.

- Use real arithmetic operations even in integer context if not all arguments are integers. (Fixes uncommon bug in some integer contexts).
- Don't force everything to lowercase on Windows. (To fix problem with Windows and `ALTER TABLE`.) Now `--lower_case_table_names` also works on Unix.
- Fixed that automatic rollback is done when thread end doesn't lock other threads.

### D.3.19. Changes in release 3.23.41 (11 August 2001)

- Added `--sql-mode=value[ ,value[ ,value ]]` option to `mysqld`. See [Section 5.2.1, “Command Options”](#).
- Fixed possible problem with `shutdown` on Solaris where the `.pid` file wasn't deleted.
- `InnoDB` now supports < 4GB rows. The former limit was 8000 bytes.
- The `doublewrite` file flush method is used in `InnoDB`. It reduces the need for Unix `fsync()` calls to a fraction and improves performance on most Unix flavors.
- You can now use the `InnoDB` Monitor to print a lot of `InnoDB` state information, including locks, to the standard output. This is useful in performance tuning.
- Several bugs which could cause hangs in `InnoDB` have been fixed.
- Split `record_buffer` to `record_buffer` and `record_rnd_buffer`. To make things compatible to previous MySQL versions, if `record_rnd_buffer` is not set, then it takes the value of `record_buffer`.
- Fixed optimizing bug in `ORDER BY` where some `ORDER BY` parts were wrongly removed.
- Fixed overflow bug with `ALTER TABLE` and `MERGE` tables.
- Added prototypes for `my_thread_init()` and `my_thread_end()` to `mysql_com.h`
- Added `--safe-user-create` option to `mysqld`.
- Fixed bug in `SELECT DISTINCT ... HAVING` that caused error message `Can't find record in #...`

### D.3.20. Changes in release 3.23.40 (18 July 2001)

- Fixed problem with `--low-priority-updates` and `INSERT` statements.
- Fixed bug in slave thread when under some rare circumstances it could get 22 bytes ahead on the offset in the master.
- Added `slave_net_timeout` for replication.
- Fixed problem with `UPDATE` and `BDB` tables.
- Fixed hard bug in `BDB` tables when using key parts.
- Fixed problem when using `GRANT FILE ON database.* ...`; previously we added the `DROP` privilege for the database.

- Fixed `DELETE FROM tbl_name ... LIMIT 0` and `UPDATE FROM tbl_name ... LIMIT 0`, which acted as though the `LIMIT` clause was not present (they deleted or updated all selected rows).
- `CHECK TABLE` now checks whether an `AUTO_INCREMENT` column contains the value 0.
- Sending a `SIGHUP` to `mysqld` now only flushes the logs, but does not reset the replication.
- Fixed parser to allow floats of type `1.0e1` (no sign after `e`).
- Option `--force` to `myisamchk` now also updates states.
- Added option `--warnings` to `mysqld`. Now `mysqld` prints the error `Aborted connection` only if this option is used.
- Fixed problem with `SHOW CREATE TABLE` when you didn't have a `PRIMARY KEY`.
- Properly fixed the rename of `innodb_unix_file_flush_method` variable to `innodb_flush_method`.
- Fixed bug when converting `BIGINT UNSIGNED` to `DOUBLE`. This caused a problem when doing comparisons with `BIGINT` values outside of the signed range.
- Fixed bug in `BDB` tables when querying empty tables.
- Fixed a bug when using `COUNT(DISTINCT)` with `LEFT JOIN` and there weren't any matching rows.
- Removed all documentation referring to the `GEMINI` table type. `GEMINI` is not released under an Open Source license.

### D.3.21. Changes in release 3.23.39 (12 June 2001)

- The `AUTO_INCREMENT` sequence wasn't reset when dropping and adding an `AUTO_INCREMENT` column.
- `CREATE ... SELECT` now creates non-unique indexes delayed.
- Fixed problem where `LOCK TABLES tbl_name READ` followed by `FLUSH TABLES` put an exclusive lock on the table.
- `REAL @variable` values were represented with only 2 digits when converted to strings.
- Fixed problem that client “hung” when `LOAD TABLE FROM MASTER` failed.
- `myisamchk --fast --force` no longer repairs tables that only had the open count wrong.
- Added functions to handle symbolic links to make life easier in 4.0.
- We are now using the `-lcma` thread library on HP-UX 10.20 so that MySQL is more stable on HP-UX.
- Fixed problem with `IF()` and number of decimals in the result.
- Fixed date-part extraction functions to work with dates where day and/or month is 0.
- Extended argument length in option files from 256 to 512 chars.

- Fixed problem with shutdown when `INSERT DELAYED` was waiting for a `LOCK TABLE`.
- Fixed core dump bug in `InnoDB` when tablespace was full.
- Fixed problem with `MERGE` tables and big tables (larger than 4GB) when using `ORDER BY`.

### D.3.22. Changes in release 3.23.38 (09 May 2001)

- Fixed a bug when `SELECT` from `MERGE` table sometimes results in incorrectly ordered rows.
- Fixed a bug in `REPLACE()` when using the `ujis` character set.
- Applied Sleepycat `BDB` patches 3.2.9.1 and 3.2.9.2.
- Added `--skip-stack-trace` option to `mysqld`.
- `CREATE TEMPORARY` now works with `InnoDB` tables.
- `InnoDB` now promotes sub keys to whole keys.
- Added option `CONCURRENT` to `LOAD DATA`.
- Better error message when slave `max_allowed_packet` is too low to read a very long log event from the master.
- Fixed bug when too many rows were removed when using `SELECT DISTINCT ... HAVING`.
- `SHOW CREATE TABLE` now returns `TEMPORARY` for temporary tables.
- Added `Rows_examined` to slow query log.
- Fixed problems with function returning empty string when used together with a group function and a `WHERE` that didn't match any rows.
- New program `mysqlcheck`.
- Added database name to output for administrative commands like `CHECK TABLE`, `REPAIR TABLE`, `OPTIMIZE TABLE`.
- Lots of portability fixes for `InnoDB`.
- Changed optimizer so that queries like `SELECT * FROM tbl_name, tbl_name2 ... ORDER BY key_part1 LIMIT row_count` use an index on `key_part1` instead of `file-sort`.
- Fixed bug when doing `LOCK TABLE to_table WRITE, ...; INSERT INTO to_table... SELECT ...` when `to_table` was empty.
- Fixed bug with `LOCK TABLE` and `BDB` tables.

### D.3.23. Changes in release 3.23.37 (17 April 2001)

- Fixed a bug when using `MATCH()` in `HAVING` clause.
- Fixed a bug when using `HEAP` tables with `LIKE`.

- Added `--mysql-version` option to `safe_mysqld`
- Changed `INNOBASE` to `InnoDB` (because the `INNOBASE` name was in use). All `configure` options and `mysqld` start options now use `innodb` instead of `innobase`. This means that before upgrading to this version, you have to change any configuration files where you have used `innobase` options!
- Fixed bug when using indexes on `CHAR(255) NULL` columns.
- Slave threads now start even if `master-host` is not set, as long as `server-id` is set and valid `master.info` is present.
- Partial updates (terminated with kill) are now logged with a special error code to the binary log. Slave refuses to execute them if the error code indicates the update was terminated abnormally, and has to be recovered with `SET SQL_SLAVE_SKIP_COUNTER=1; SLAVE START` after a manual sanity check/correction of data integrity.
- Fixed bug that erroneously logged a drop of internal temporary table on thread termination to the binary log --- this bug affected replication.
- Fixed a bug in `REGEXP` on 64-bit machines.
- `UPDATE` and `DELETE` with `WHERE unique_key_part IS NULL` didn't update/delete all rows.
- Disabled `INSERT DELAYED` for tables that support transactions.
- Fixed bug when using date functions on `TEXT/BLOB` column with wrong date format.
- UDFs now also work on Windows. (Patch by Ralph Mason.)
- Fixed bug in `ALTER TABLE` and `LOAD DATA INFILE` that disabled key-sorting. These commands should now be faster in most cases.
- Fixed performance bug where reopened tables (tables that had been waiting for `FLUSH` or `REPAIR TABLE`) would not use indexes for the next query.
- Fixed problem with `ALTER TABLE` to `InnoDB` tables on FreeBSD.
- Added `mysqld` variables `myisam_max_sort_file_size` and `myisam_max_extra_sort_file_size`.
- Initialize signals early to avoid problem with signals in `InnoDB`.
- Applied patch for the `tis620` character set to make comparisons case-independent and to fix a bug in `LIKE` for this character set. **Note:** All tables that use the `tis620` character set must be fixed with `myisamchk -r` or `REPAIR TABLE`!
- Added `--skip-safemalloc` option to `mysqld`.

### D.3.24. Changes in release 3.23.36 (27 March 2001)

- Fixed a bug that allowed use of database names containing a `'.` character. This fixes a serious security issue when `mysqld` is run as root. (CVE-2001-0407 [http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2001-0407])
- Fixed bug when thread creation failed (could happen when doing a **lot** of connections in a short

time).

- Fixed some problems with `FLUSH TABLES` and `TEMPORARY` tables. (Problem with freeing the key cache and error `Can't reopen table...`)
- Fixed a problem in `InnoDB` with other character sets than `latin1` and another problem when using many columns.
- Fixed bug that caused a core dump when using a very complex query involving `DISTINCT` and summary functions.
- Added `SET TRANSACTION ISOLATION LEVEL ...`
- Added `SELECT ... FOR UPDATE`.
- Fixed bug where the number of affected rows was not returned when MySQL was compiled without transaction support.
- Fixed a bug in `UPDATE` where keys weren't always used to find the rows to be updated.
- Fixed a bug in `CONCAT_WS( )` where it returned incorrect results.
- Changed `CREATE ... SELECT` and `INSERT ... SELECT` to not allow concurrent inserts as this could make the binary log hard to repeat. (Concurrent inserts are enabled if you are not using the binary or update log.)
- Changed some macros to be able to use fast mutex with `glibc 2.2`.

### D.3.25. Changes in release 3.23.35 (15 March 2001)

- Fixed newly introduced bug in `ORDER BY`.
- Fixed wrong define `CLIENT_TRANSACTIONS`.
- Fixed bug in `SHOW VARIABLES` when using `INNOBASE` tables.
- Setting and using user variables in `SELECT DISTINCT` didn't work.
- Tuned `SHOW ANALYZE` for small tables.
- Fixed handling of arguments in the benchmark script `run-all-tests`.

### D.3.26. Changes in release 3.23.34a (11 March 2001)

- Added extra files to the distribution to allow `INNOBASE` support to be compiled.

### D.3.27. Changes in release 3.23.34 (10 March 2001)

- Added the `INNOBASE` storage engine and the `BDB` storage engine to the MySQL source distribution.
- Updated the documentation about `GEMINI` tables.

- Fixed a bug in `INSERT DELAYED` that caused threads to hang when inserting `NULL` into an `AUTO_INCREMENT` column.
- Fixed a bug in `CHECK TABLE / REPAIR TABLE` that could cause a thread to hang.
- Fixed problem that `REPLACE` would not replace a row that conflicts with an `AUTO_INCREMENT` generated key.
- `mysqld` now only sets `CLIENT_TRANSACTIONS` in `mysql->server_capabilities` if the server supports a transaction-safe storage engine.
- Fixed `LOAD DATA INFILE` to allow numeric values to be read into `ENUM` and `SET` columns.
- Improved error diagnostic for slave thread exit.
- Fixed bug in `ALTER TABLE ... ORDER BY`.
- Added `max_user_connections` variable to `mysqld`.
- Limit query length for replication by `max_allowed_packet`, not the arbitrary limit of 4MB.
- Allow space around `=` in argument to `--set-variable`.
- Fixed problem in automatic repair that could leave some threads in state `Waiting for table`.
- `SHOW CREATE TABLE` now displays the `UNION=()` for `MERGE` tables.
- `ALTER TABLE` now remembers the old `UNION=()` definition.
- Fixed bug when replicating timestamps.
- Fixed bug in bidirectional replication.
- Fixed bug in the `BDB` storage engine that occurred when using an index on multiple-part key where a key part may be `NULL`.
- Fixed `MAX()` optimization on sub-key for `BDB` tables.
- Fixed problem where garbage results were returned when using `BDB` tables and `BLOB` or `TEXT` fields when joining many tables.
- Fixed a problem with `BDB` tables and `TEXT` columns.
- Fixed bug when using a `BLOB` key where a const row wasn't found.
- Fixed that `mysqlbinlog` writes the timestamp value for each query. This ensures that one gets same values for date functions like `NOW()` when using `mysqlbinlog` to pipe the queries to another server.
- Allow `--skip-gemini`, `--skip-bdb`, and `--skip-innodb` options to be specified when invoking `mysqld`, even if these storage engines are not compiled in to `mysqld`.
- You can now use `ASC` and `DESC` with `GROUP BY` columns to specify a sort order.
- Fixed a deadlock in the `SET` code, when one ran `SET @foo=bar`, where `bar` is a column reference, an error was not properly generated.

### D.3.28. Changes in release 3.23.33 (09 February 2001)

- Fixed DNS lookups not to use the same mutex as the hostname cache. This enables known hosts to be quickly resolved even if a DNS lookup takes a long time.
- Added `--character-sets-dir` option to `myisampack`.
- Removed warnings when running `REPAIR TABLE ... EXTENDED`.
- Fixed a bug that caused a core dump when using `GROUP BY` on an alias, where the alias was the same as an existing column name.
- Added `SEQUENCE()` as an example user-defined function.
- Changed `mysql_install_db` to use `BINARY` for `CHAR` columns in the privilege tables.
- Changed `TRUNCATE tbl_name` to `TRUNCATE TABLE tbl_name` to use the same syntax as Oracle. Until 4.0 we also allow `TRUNCATE tbl_name` to not crash old code.
- Fixed “no found rows” bug in `MyISAM` tables when a `BLOB` was first part of a multiple-part key.
- Fixed bug where `CASE` didn't work with `GROUP BY`.
- Added `--sort-recover` option to `myisamchk`.
- `myisamchk -S` and `OPTIMIZE TABLE` now work on Windows.
- Fixed bug when using `DISTINCT` on results from functions that referred to a group function, like:

```
SELECT a, DISTINCT SEC_TO_TIME(SUM(a))  
FROM tbl_name GROUP BY a, b;
```

- Fixed buffer overrun in `libmysqlclient` library. Fixed bug in handling `STOP` event after `ROTATE` event in replication.
- Fixed another buffer overrun in `DROP DATABASE`.
- Added `Table_locks_immediate` and `Table_locks_waited` status variables.
- Fixed bug in replication that broke slave server start with existing `master.info`. This fixes a bug introduced in 3.23.32.
- Added `SET SQL_SLAVE_SKIP_COUNTER=n` command to recover from replication glitches without a full database copy.
- Added `max_binlog_size` variable; the binary log is rotated automatically when the size crosses the limit.
- Added `Last_Error`, `Last_Errno`, and `Slave_skip_counter` variables to `SHOW SLAVE STATUS`.
- Fixed bug in `MASTER_POS_WAIT()` function.
- Execute core dump handler on `SIGILL`, and `SIGBUS` in addition to `SIGSEGV`.
- On x86 Linux, print the current query and thread (connection) id, if available, in the core dump handler.
- Fixed several timing bugs in the test suite.
- Extended `mysqltest` to take care of the timing issues in the test suite.



- `ALTER TABLE` can now be used to change the definition for a `MERGE` table.
- Fixed creation of `MERGE` tables on Windows.
- Portability fixes for OpenBSD and OS/2.
- Added `--temp-pool` option to `mysqld`. Using this option causes most temporary files created to use a small set of names, rather than a unique name for each new file. This is to work around a problem in the Linux kernel dealing with creating a bunch of new files with different names. With the old behavior, Linux seems to "leak" memory, as it's being allocated to the directory entry cache instead of the disk cache.

### D.3.29. Changes in release 3.23.32 (22 January 2001)

- Changed code to get around compiler bug in Compaq C++ on OSF/1, that broke `BACKUP TABLE`, `RESTORE TABLE`, `CHECK TABLE`, `REPAIR TABLE`, and `ANALYZE TABLE`.
- Added option `FULL` to `SHOW COLUMNS`. Now we show the privilege list for the columns only if this option is given.
- Fixed bug in `SHOW LOGS` when there weren't any `BDB` logs.
- Fixed a timing problem in replication that could delay sending an update to the client until a new update was done.
- Don't convert field names when using `mysql_list_fields()`. This is to keep this code compatible with `SHOW FIELDS`.
- `MERGE` tables didn't work on Windows.
- Fixed problem with `SET PASSWORD=...` on Windows.
- Added missing `my_config.h` to RPM distribution.
- `TRIM("foo" from "foo")` didn't return an empty string.
- Added `--with-version-suffix` option to `configure`.
- Fixed core dump when client aborted connection without `mysql_close()`.
- Fixed a bug in `RESTORE TABLE` when trying to restore from a non-existent directory.
- Fixed a bug which caused a core dump on the slave when replicating `SET PASSWORD`.
- Added `MASTER_POS_WAIT()` function.

### D.3.30. Changes in release 3.23.31 (17 January 2001: Production)

- The test suite now tests all reachable `BDB` interface code. During testing we found and fixed many errors in the interface code.
- Using `HAVING` on an empty table could produce one result row when it shouldn't.

- Fixed the MySQL RPM so it no longer depends on Perl5.
- Fixed some problems with `HEAP` tables on Windows.
- `SHOW TABLE STATUS` didn't show correct average row length for tables larger than 4GB.
- `CHECK TABLE ... EXTENDED` didn't check row links for fixed size tables.
- Added option `MEDIUM` to `CHECK TABLE`.
- Fixed problem when using `DECIMAL()` keys on negative numbers.
- `hour()` (and some other `TIME` functions) on a `CHAR` column always returned `NULL`.
- Fixed security bug in `SHOW GRANT` (please upgrade if you are using an earlier MySQL 3.23 version). ([CVE-2001-1275](http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2001-1275) [<http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2001-1275>])
- Fixed buffer overflow bug when writing a certain error message. ([CVE-2001-1274](http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2001-1274) [<http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2001-1274>])
- Added usage of `setrlimit()` on Linux to get `-O --open_files_limit=val` to work on Linux.
- Added `bdb_version` variable to `mysqld`.
- Fixed bug when using expression of type:

```
SELECT ... FROM t1 LEFT JOIN t2 ON (t1.a=t2.a) WHERE t1.a=t2.a
```

In this case the test in the `WHERE` clause was wrongly optimized away.

- Fixed bug in `MyISAM` when deleting keys with possible `NULL` values, but the first key-column was not a prefix-compressed text column.
- Fixed `mysql.server` to read the `[mysql.server]` option file group rather than the `[mysql_server]` group.
- Fixed `safe_mysqld` and `mysql.server` to also read the `server` option section.
- Added `Threads_created` status variable to `mysqld`.

### D.3.31. Changes in release 3.23.30 (04 January 2001)

- Added `SHOW OPEN TABLES` command.
- Fixed that `myisamdump` works against old `mysqld` servers.
- Fixed `myisamchk -kN` so that it works again.
- Fixed a problem with replication when the binary log file went over 2G on 32-bit systems.
- `LOCK TABLES` now automatically starts a new transaction.
- Changed `BDB` tables to not use internal subtransactions and reuse open files to get more speed.
- Added `--mysqld=path` option to `safe_mysqld`.

- Allow hex constants in the `--fields-*--by` and `--lines-terminated-by` options to `mysqldump` and `mysqlimport`. By Paul DuBois.
- Added `--safe-show-database` option to `mysqld`.
- Added `have_bdb`, `have_gemini`, `have_innobase`, `have_raid` and `have_openssl` to `SHOW VARIABLES` to make it easy to test for supported extensions.
- Added `--open-files-limit` option to `mysqld`.
- Changed `--open-files` option to `--open-files-limit` in `safe_mysqld`.
- Fixed a bug where some rows were not found with `HEAP` tables that had many keys.
- Fixed that `--bdb-no-sync` works.
- Changed `--bdb-recover` to `--bdb-no-recover` as recover should be on by default.
- Changed the default number of `BDB` locks to 10000.
- Fixed a bug from 3.23.29 when allocating the shared structure needed for `BDB` tables.
- Changed `mysqld_multi.sh` to use configure variables. Patch by Christopher McCrory.
- Added fixing of include files for Solaris 2.8.
- Fixed bug with `--skip-networking` on Debian Linux.
- Fixed problem that some temporary files were reported as having the name `UNOPENED` in error messages.
- Fixed bug when running two simultaneous `SHOW LOGS` queries.

### D.3.32. Changes in release 3.23.29 (16 December 2000)

- Configure updates for Tru64, large file support, and better TCP wrapper support. By Albert Chin-A-Young.
- Fixed bug in `<=>` operator.
- Fixed bug in `REPLACE` with `BDB` tables.
- `LPAD( )` and `RPAD( )` shortens the result string if it's longer than the length argument.
- Added `SHOW LOGS` command.
- Remove unused `BDB` logs on shutdown.
- When creating a table, put `PRIMARY` keys first, followed by `UNIQUE` keys.
- Fixed a bug in `UPDATE` involving multiple-part keys where you specified all key parts both in the update and the `WHERE` part. In this case MySQL could try to update a record that didn't match the whole `WHERE` part.
- Changed drop table to first drop the tables and then the `.frm` file.
- Fixed a bug in the hostname cache which caused `mysqld` to report the hostname as `' '` in some er-

ror messages.

- Fixed a bug with `HEAP` type tables; the variable `max_heap_table_size` wasn't used. Now either `MAX_ROWS` or `max_heap_table_size` can be used to limit the size of a `HEAP` type table.
- Changed the default `server-id` value to 1 for masters and 2 for slaves to make it easier to use the binary log.
- Renamed `bdb_lock_max` variable to `bdb_max_lock`.
- Added support for `AUTO_INCREMENT` on sub-fields for `BDB` tables.
- Added `ANALYZE TABLE` of `BDB` tables.
- In `BDB` tables, we now store the number of rows; this helps to optimize queries when we need an approximation of the number of rows.
- If we get an error in a multiple-row statement, we now only roll back the last statement, not the entire transaction.
- If you do a `ROLLBACK` when you have updated a non-transactional table you get an error as a warning.
- Added `--bdb-shared-data` option to `mysqld`.
- Added `Slave_open_temp_tables` status variable to `mysqld`
- Added `binlog_cache_size` and `max_binlog_cache_size` variables to `mysqld`.
- `DROP TABLE`, `RENAME TABLE`, `CREATE INDEX` and `DROP INDEX` are now transaction endpoints.
- If you do a `DROP DATABASE` on a symbolically linked database, both the link and the original database are deleted.
- Fixed `DROP DATABASE` to work on OS/2.
- Fixed bug when doing a `SELECT DISTINCT ... table1 LEFT JOIN table2 ...` when `table2` was empty.
- Added `--abort-slave-event-count` and `--disconnect-slave-event-count` options to `mysqld` for debugging and testing of replication.
- Fixed replication of temporary tables. Handles everything except slave server restart.
- `SHOW KEYS` now shows whether key is `FULLTEXT`.
- New script `mysqld_multi`. See [Section 5.4.3, “mysqld\\_multi — Manage Multiple MySQL Servers”](#).
- Added new script, `mysql-multi.server.sh`. Thanks to Tim Bunce <Tim.Bunce@ig.co.uk> for modifying `mysql.server` to easily handle hosts running many `mysqld` processes.
- `safe_mysqld`, `mysql.server`, and `mysql_install_db` have been modified to use `mysql_print_defaults` instead of various hacks to read the `my.cnf` files. In addition, the handling of various paths has been made more consistent with how `mysqld` handles them by default.

- Automatically remove Berkeley DB transaction logs that no longer are in use.
- Fixed bug with several `FULLTEXT` indexes in one table.
- Added a warning if number of rows changes on `REPAIR TABLE/OPTIMIZE TABLE`.
- Applied patches for OS/2 by [Yuri Dario](#).
- `FLUSH TABLES tbl_name` didn't always flush the index tree to disk properly.
- `--bootstrap` is now run in a separate thread. This fixes a problem that caused `mysql_install_db` to core dump on some Linux machines.
- Changed `mi_create()` to use less stack space.
- Fixed bug with optimizer trying to over-optimize `MATCH()` when used with `UNIQUE` key.
- Changed `crash-me` and the MySQL benchmarks to also work with FrontBase.
- Allow `RESTRICT` and `CASCADE` after `DROP TABLE` to make porting easier.
- Reset status variable which could cause problem if one used `--slow-log`.
- Added `connect_timeout` variable to `mysql` and `mysqladmin`.
- Added `connect-timeout` as an alias for `timeout` for option files read by `mysql_options()`.

### D.3.33. Changes in release 3.23.28 (22 November 2000: Gamma)

- Added new options `--pager[=...]`, `--no-pager`, `--tee=...` and `--no-tee` to the `mysql` client. The new corresponding interactive commands are `pager`, `nopager`, `tee` and `notee`. See [Section 8.6, “mysql — The MySQL Command-Line Tool”](#), `mysql --help` and the interactive help for more information.
- Fixed crash when automatic repair of `MyISAM` table failed.
- Fixed a major performance bug in the table locking code when a lot of `SELECT`, `UPDATE` and `INSERT` statements constantly were running. The symptom was that the `UPDATE` and `INSERT` queries were locked for a long time while new `SELECT` statements were executed before the updates.
- When reading `options_files` with `mysql_options()` the `return-found-rows` option was ignored.
- You can now specify `interactive-timeout` in the option file that is read by `mysql_options()`. This makes it possible to force programs that run for a long time (like `mysqlhotcopy`) to use the `interactive_timeout` time instead of the `wait_timeout` time.
- Added to the slow query log the time and the username for each logged query. If you are using `--log-long-format` then also queries that do not use an index are logged, even if the query takes less than `long_query_time` seconds.
- Fixed a problem in `LEFT JOIN` which caused all columns in a reference table to be `NULL`.

- Fixed a problem when using `NATURAL JOIN` without keys.
- Fixed a bug when using a multiple-part keys where the first part was of type `TEXT` or `BLOB`.
- `DROP` of temporary tables wasn't stored in the update/binary log.
- Fixed a bug where `SELECT DISTINCT * ... LIMIT row_count` only returned one row.
- Fixed a bug in the assembler code in `strstr()` for SPARC and cleaned up the `global.h` header file to avoid a problem with bad aliasing with the compiler submitted with Red Hat 7.0. (Reported by Trond Eivind Glomsrød)
- The `--skip-networking` option now works properly on NT.
- Fixed a long outstanding bug in the `ISAM` tables when a row with a length of more than 65KB was shortened by a single byte.
- Fixed a bug in `MyISAM` when running multiple updating processes on the same table.
- Allow one to use `FLUSH TABLE tbl_name`.
- Added `--replicate-ignore-table`, `--replicate-do-table`, `--replicate-wild-ignore-table`, and `--replicate-wild-do-table` options to `mysqld`.
- Changed all log files to use our own `IO_CACHE` mechanism instead of `FILE` to avoid OS problems when there are many files open.
- Added `--open-files` and `--timezone` options to `safe_mysqld`.
- Fixed a fatal bug in `CREATE TEMPORARY TABLE ... SELECT ...`
- Fixed a problem with `CREATE TABLE ... SELECT NULL`.
- Added variables `large_file_support`, `net_read_timeout`, `net_write_timeout` and `query_buffer_size` to `SHOW VARIABLES`.
- Added status variables `Created_tmp_files` and `Sort_merge_passes` to `SHOW STATUS`.
- Fixed a bug where we didn't allow an index name after the `FOREIGN KEY` definition.
- Added `TRUNCATE tbl_name` as a synonym for `DELETE FROM tbl_name`.
- Fixed a bug in a `BDB` key compare function when comparing part keys.
- Added `bdb_lock_max` variable to `mysqld`.
- Added more tests to the benchmark suite.
- Fixed an overflow bug in the client code when using overly long database names.
- `mysql_connect()` now aborts on Linux if the server doesn't answer in `timeout` seconds.
- `SLAVE START` did not work if you started with `--skip-slave-start` and had not explicitly run `CHANGE MASTER TO`.
- Fixed the output of `SHOW MASTER STATUS` to be consistent with `SHOW SLAVE STATUS`. (It now has no directory in the log name.)
- Added `PURGE MASTER LOGS TO`.

- Added `SHOW MASTER LOGS` statement to display a list of binary log files.
- Added `--safemalloc-mem-limit` option to `mysqld` to simulate memory shortage when compiled with the `--with-debug=full` option.
- Fixed several core dumps in out-of-memory conditions.
- `SHOW SLAVE STATUS` was using an uninitialized mutex if the slave had not been started yet.
- Fixed bug in `ELT()` and `MAKE_SET()` when the query used a temporary table.
- `CHANGE MASTER TO` without specifying `MASTER_LOG_POS` would set it to 0 instead of 4 and hit the magic number in the master binary log.
- `ALTER TABLE ... ORDER BY ...` syntax added. This creates the new table with the rows in a specific order.

### D.3.34. Changes in release 3.23.27 (24 October 2000)

- Fixed a bug where the automatic repair of `MyISAM` tables sometimes failed when the data file was corrupt.
- Fixed a bug in `SHOW CREATE` when using `AUTO_INCREMENT` columns.
- Changed `BDB` tables to use new compare function in Berkeley DB 3.2.3.
- You can now use Unix socket files with MIT-pthreads.
- Added the `latin5` (turkish) character set.
- Small portability fixes.

### D.3.35. Changes in release 3.23.26 (18 October 2000)

- Renamed `FLUSH MASTER` and `FLUSH SLAVE` to `RESET MASTER` and `RESET SLAVE`.
- Fixed `<>` to work properly with `NULL`.
- Fixed a problem with `SUBSTRING_INDEX()` and `REPLACE()`. (Patch by Alexander Igonitchev)
- Fix `CREATE TEMPORARY TABLE IF NOT EXISTS` not to produce an error if the table exists.
- If you don't create a `PRIMARY KEY` in a `BDB` table, a hidden `PRIMARY KEY` is created.
- Added read-only-key optimization to `BDB` tables.
- `LEFT JOIN` in some cases preferred a full table scan when there was no `WHERE` clause.
- When using `--log-slow-queries`, don't count the time waiting for a lock.
- Fixed bug in lock code on Windows which could cause the key cache to report that the key file was crashed even if it was okay.
- Automatic repair of `MyISAM` tables if you start `mysqld` with `--myisam-recover`.

- Removed the `TYPE=` keyword from `CHECK TABLE` and `REPAIR TABLE`. Allow `CHECK TABLE` options to be combined. (You can still use `TYPE=`, but this usage is deprecated.)
- Fixed mutex bug in the binary replication log --- long update queries could be read only in part by the slave if it did it at the wrong time, which was not fatal, but resulted in a performance-degrading reconnect and a scary message in the error log.
- Changed the format of the binary log --- added magic number, server version, binary log version. Added the server ID and query error code for each query event.
- Replication thread from the slave now kills all the stale threads from the same server.
- Long replication usernames were not being handled properly.
- Added `--replicate-rewrite-db` option to `mysqld`.
- Added `--skip-slave-start` option to `mysqld`.
- Updates that generated an error code (such as `INSERT INTO foo(some_key) values (1), (1)`) erroneously terminated the slave thread.
- Added optimization of queries where `DISTINCT` is used only on columns from some of the tables.
- Allow floating-point numbers where there is no sign after the exponent (like `1e1`).
- `SHOW GRANTS` didn't always show all column grants.
- Added `--default-extra-file=file_name` option to all MySQL clients.
- Columns referenced in `INSERT` statements now are initialized properly.
- `UPDATE` didn't always work when used with a range on a timestamp that was part of the key that was used to find rows.
- Fixed a bug in `FULLTEXT` index when inserting a `NULL` column.
- Changed to use `mkstemp()` instead of `tempnam()`. Based on a patch from John Jones.

### D.3.36. Changes in release 3.23.25 (29 September 2000)

- Fixed that `dbname` works as second argument to `mysqlhotcopy`.
- The values for the `UMASK` and `UMASK_DIR` environment variables now can be specified in octal by beginning the value with a zero.
- Added `RIGHT JOIN`. This makes `RIGHT` a reserved word.
- Added `@@IDENTITY` as a synonym for `LAST_INSERT_ID()`. (This is for MSSQL compatibility.)
- Fixed a bug in `myisamchk` and `REPAIR TABLE` when using `FULLTEXT` index.
- `LOAD DATA INFILE` now works with FIFOs. (Patch by Toni L. Harbaugh-Blackford.)
- `FLUSH LOGS` broke replication if you specified a log name with an explicit extension as the value of the `log-bin` option.



- Fixed a bug in `MyISAM` with packed multiple-part keys.
- Fixed crash when using `CHECK TABLE` on Windows.
- Fixed a bug where `FULLTEXT` index always used the `koi8_ukr` character set.
- Fixed privilege checking for `CHECK TABLE`.
- The `MyISAM` repair/reindex code didn't use the `--tmpdir` option for its temporary files.
- Added `BACKUP TABLE` and `RESTORE TABLE`.
- Fixed core dump on `CHANGE MASTER TO` when the slave did not have the master to start with.
- Fixed incorrect `Time` in the processlist for `Connect` of the slave thread.
- The slave now logs when it connects to the master.
- Fixed a core dump bug when doing `FLUSH MASTER` if you didn't specify a filename argument to `-log-bin`.
- Added missing `ha_berkeley.x` files to the MySQL Windows distribution.
- Fixed some mutex bugs in the log code that could cause thread blocks if new log files couldn't be created.
- Added lock time and number of selected processed rows to slow query log.
- Added `--memlock` option to `mysqld` to lock `mysqld` in memory on systems with the `mlock-all()` call (as in Solaris).
- `HEAP` tables didn't use keys properly. (Bug from 3.23.23.)
- Added better support for `MERGE` tables (keys, mapping, creation, documentation...). See [Section 14.3, “The MERGE Storage Engine”](#).
- Fixed bug in `mysqldump` from 3.23 which caused some `CHAR` columns not to be quoted.
- Merged `analyze`, `check`, `optimize` and repair code.
- `OPTIMIZE TABLE` is now mapped to `REPAIR TABLE` with statistics and sorting of the index tree. This means that for the moment it only works on `MyISAM` tables.
- Added a pre-allocated block to `root_malloc` to get fewer mallocs.
- Added a lot of new statistics variables.
- Fixed `ORDER BY` bug with `BDB` tables.
- Removed warning that `mysqld` couldn't remove the `.pid` file under Windows.
- Changed `--log-isam` to log `MyISAM` tables instead of `isam` tables.
- Fixed `CHECK TABLE` to work on Windows.
- Added file mutexes to make `pwrite()` safe on Windows.

### D.3.37. Changes in release 3.23.24 (08 September 2000)

- Added `Created_tmp_disk_tables` variable to `mysqld`.
- To make it possible to reliably dump and restore tables with `TIMESTAMP(X)` columns, MySQL now reports columns with `X` other than 14 or 8 to be strings.
- Changed sort order for `latin1` as it was before MySQL 3.23.23. Any table that was created or modified with 3.23.22 must be repaired if it has `CHAR` columns that may contain characters with ASCII values greater than 128!
- Fixed small memory leak introduced from 3.23.22 when creating a temporary table.
- Fixed problem with `BDB` tables and reading on a unique (not primary) key.
- Restored the `win1251` character set (it's now only marked deprecated).

### D.3.38. Changes in release 3.23.23 (01 September 2000)

- Changed sort order for 'German'; all tables created with 'German' sortorder must be repaired with `REPAIR TABLE` or `myisamchk` before use!
- Added `--core-file` option to `mysqld` to get a core file on Linux if `mysqld` dies on the `SIGSEGV` signal.
- MySQL client `mysql` now starts with option `--no-named-commands (-g)` by default. This option can be disabled with `--enable-named-commands (-G)`. This may cause incompatibility problems in some cases, for example, in SQL scripts that use named commands without a semi-colon! Long format commands still work from the first line.
- Fixed a problem when using many pending `DROP TABLE` statements at the same time.
- Optimizer didn't use keys properly when using `LEFT JOIN` on an empty table.
- Added shorter help text when invoking `mysqld` with incorrect options.
- Fixed non-fatal `free()` bug in `mysqlimport`.
- Fixed bug in MyISAM index handling of `DECIMAL/NUMERIC` keys.
- Fixed a bug in concurrent insert in MyISAM tables. In some contexts, usage of `MIN(key_part)` or `MAX(key_part)` returned an empty set.
- Updated `mysqlhotcopy` to use the new `FLUSH TABLES table_list` syntax. Only tables which are being backed up are flushed now.
- Changed behavior of `--enable-thread-safe-client` so that both non-threaded (`-lmysqlclient`) and threaded (`-lmysqlclient_r`) libraries are built. Users who linked against a threaded `-lmysqlclient` need to link against `-lmysqlclient_r` now.
- Added atomic `RENAME TABLE` command.
- Don't count `NULL` values in `COUNT(DISTINCT ...)`.
- Changed `ALTER TABLE`, `LOAD DATA INFILE` on empty tables and `INSERT ... SELECT ...` on empty tables to create non-unique indexes in a separate batch with sorting. This makes these statements much faster when you have many indexes.
- `ALTER TABLE` now logs the first used `insert_id` correctly.

- Fixed crash when adding a default value to a `BLOB` column.
- Fixed a bug with `DATE_ADD/DATE_SUB` where it returned a datetime instead of a date.
- Fixed a problem with the thread cache which made some threads show up as `***DEAD***` in `SHOW PROCESSLIST`.
- Fixed a lock in our `thr_rwlock` code, which could make selects that run at the same time as concurrent inserts crash. This affects only systems that don't have the `pthread_rwlock_rdlock` code.
- When deleting rows with a non-unique key in a `HEAP` table, all rows weren't always deleted.
- Fixed bug in range optimizer for `HEAP` tables for searches on a part index.
- Fixed `SELECT` on part keys to work with `BDB` tables.
- Fixed `INSERT INTO bdb_table ... SELECT` to work with `BDB` tables.
- `CHECK TABLE` now updates key statistics for the table.
- `ANALYZE TABLE` now only updates tables that have been changed since the last `ANALYZE TABLE`. Note that this is a new feature and tables are not marked to be analyzed until they are updated in any way with 3.23.23 or newer. For older tables, you have to do `CHECK TABLE` to update the key distribution.
- Fixed some minor privilege problems with `CHECK TABLE`, `ANALYZE TABLE`, `REPAIR TABLE` and `SHOW CREATE` commands.
- Added `CHANGE MASTER TO` statement.
- Added `FAST`, `QUICK` `EXTENDED` check types to `CHECK TABLES`.
- Changed `myisamchk` so that `--fast` and `--check-only-changed` are also honored with `--sort-index` and `--analyze`.
- Fixed fatal bug in `LOAD TABLE FROM MASTER` that did not lock the table during index re-build.
- `LOAD DATA INFILE` broke replication if the database was excluded from replication.
- More variables in `SHOW SLAVE STATUS` and `SHOW MASTER STATUS`.
- `SLAVE STOP` now does not return until the slave thread actually exits.
- Full-text search via the `MATCH()` function and `FULLTEXT` index type (for `MyISAM` files). This makes `FULLTEXT` a reserved word.

### D.3.39. Changes in release 3.23.22 (31 July 2000)

- Fixed that `lex_hash.h` is created properly for each MySQL distribution.
- Fixed that `MASTER` and `COLLECTION` are not reserved words.
- The log generated by `--slow-query-log` didn't contain the whole queries.
- Fixed that open transactions in `BDB` tables are rolled back if the connection is closed unexpectedly.
- Added workaround for a bug in `gcc 2.96` (intel) and `gcc 2.9` (IA-64) in `gen_lex_hash.c`.

- Fixed memory leak in the client library when using `host=` in the `my.cnf` file.
- Optimized functions that manipulate the hours/minutes/seconds.
- Fixed bug when comparing the result of `DATE_ADD()`/`DATE_SUB()` against a number.
- Changed the meaning of `-F`, `--fast` for `myisamchk`. Added `-C`, `--check-only-changed` option to `myisamchk`.
- Added `ANALYZE tbl_name` to update key statistics for tables.
- Changed binary items `0x...` to be regarded as integers by default.
- Fix for SCO and `SHOW PROCESSLIST`.
- Added `auto-rehash` on reconnect for the `mysql` client.
- Fixed a newly introduced bug in `MyISAM`, where the index file couldn't get bigger than 64MB.
- Added `SHOW MASTER STATUS` and `SHOW SLAVE STATUS`.

### D.3.40. Changes in release 3.23.21 (04 July 2000)

- Added `mysql_character_set_name()` function to the MySQL C API.
- Made the update log ASCII 0 safe.
- Added the `mysql_config` script.
- Fixed problem when using `<` or `>` with a char column that was only partly indexed.
- One would get a core dump if the log file was not readable by the MySQL user.
- Changed `mysqladmin` to use `CREATE DATABASE` and `DROP DATABASE` statements instead of the old deprecated API calls.
- Fixed `chown` warning in `safe_mysqld`.
- Fixed a bug in `ORDER BY` that was introduced in 3.23.19.
- Only optimize the `DELETE FROM tbl_name` to do a drop+create of the table if we are in `AUTO-COMMIT` mode (needed for `BDB` tables).
- Added extra checks to avoid index corruption when the `ISAM/MyISAM` index files get full during an `INSERT/UPDATE`.
- `myisamchk` didn't correctly update row checksum when used with `-ro` (this only gave a warning in subsequent runs).
- Fixed bug in `REPAIR TABLE` so that it works with tables without indexes.
- Fixed buffer overrun in `DROP DATABASE`.
- `LOAD TABLE FROM MASTER` is sufficiently bug-free to announce it as a feature.
- `MATCH` and `AGAINST` are now reserved words.

### D.3.41. Changes in release 3.23.20 (28 June 2000: Beta)

- Fixed bug in 3.23.19; `DELETE FROM tbl_name` removed the `.frm` file.
- Added `SHOW CREATE TABLE`.

### D.3.42. Changes in release 3.23.19

- Changed copyright for all files to GPL for the server code and utilities and to LGPL for the client libraries. See <http://www.fsf.org/licenses/>.
- Fixed bug where all rows matching weren't updated on a `MyISAM` table when doing update based on key on a table with many keys and some key changed values.
- The Linux MySQL RPMs and binaries are now statically linked with a `linuxthread` version that has faster mutex handling when used with MySQL.
- `ORDER BY` can now use `REF` keys to find subsets of the rows that need to be sorted.
- Changed name of `print_defaults` program to `my_print_defaults` to avoid name confusion.
- Fixed `NULLIF()` to work as required by standard SQL.
- Added `net_read_timeout` and `net_write_timeout` as startup parameters to `mysqld`.
- Fixed bug that destroyed index when doing `myisamchk --sort-records` on a table with prefix compressed index.
- Added `pack_isam` and `myisampack` to the standard MySQL distribution.
- Added the syntax `BEGIN WORK` (the same as `BEGIN`).
- Fixed core dump bug when using `ORDER BY` on a `CONV()` expression.
- Added `LOAD TABLE FROM MASTER`.
- Added `FLUSH MASTER` and `FLUSH SLAVE`.
- Fixed big/little endian problem in the replication.

### D.3.43. Changes in release 3.23.18 (11 June 2000)

- Fixed a problem from 3.23.17 when choosing character set on the client side.
- Added `FLUSH TABLES WITH READ LOCK` to make a global lock suitable for making a copy of MySQL data files.
- `CREATE TABLE ... SELECT ... PROCEDURE` now works.
- Internal temporary tables now use compressed index when using `GROUP BY` on `VARCHAR/CHAR` columns.
- Fixed a problem when locking the same table with both a `READ` and a `WRITE` lock.

- Fixed problem with `myisamchk` and `RAID` tables.

### D.3.44. Changes in release 3.23.17 (07 June 2000)

- Fixed a bug in `FIND_IN_SET()` when the first argument was `NULL`.
- Added table locks to Berkeley DB.
- Fixed a bug with `LEFT JOIN` and `ORDER BY` where the first table had only one matching row.
- Added 4 sample `my.cnf` example files in the `support-files` directory.
- Fixed `duplicated key` problem when doing big `GROUP BY` operations. (This bug was probably introduced in 3.23.15.)
- Changed syntax for `INNER JOIN` to match standard SQL.
- Added `NATURAL JOIN` syntax.
- A lot of fixes in the `BDB` interface.
- Added handling of `--no-defaults` and `--defaults-file` to `safe_mysqld.sh` and `mysql_install_db.sh`.
- Fixed bug in reading compressed tables with many threads.
- Fixed that `USE INDEX` works with `PRIMARY` keys.
- Added `BEGIN` statement to start a transaction in `AUTOCOMMIT` mode.
- Added support for symbolic links for Windows.
- Changed protocol to let client know if the server is in `AUTOCOMMIT` mode and if there is a pending transaction. If there is a pending transaction, the client library gives an error before reconnecting to the server to let the client know that the server did a rollback. The protocol is still backward-compatible with old clients.
- `KILL` now works on a thread that is locked on a 'write' to a dead client.
- Fixed memory leak in the replication slave thread.
- Added new `log-slave-updates` option to `mysqld`, to allow daisy-chaining the slaves.
- Fixed compile error on FreeBSD and other systems where `pthread_t` is not the same as `int`.
- Fixed master shutdown aborting the slave thread.
- Fixed a race condition in `INSERT DELAYED` code when doing `ALTER TABLE`.
- Added deadlock detection sanity checks to `INSERT DELAYED`.

### D.3.45. Changes in release 3.23.16 (16 May 2000)

- Added `SLAVE START` and `SLAVE STOP` statements.

- Added `TYPE=QUICK` option to `CHECK TABLE` and to `REPAIR TABLE`.
- Fixed bug in `REPAIR TABLE` when the table was in use by other threads.
- Added a thread cache to make it possible to debug MySQL with `gdb` when one does a lot of reconnects. This also improves systems where you can't use persistent connections.
- Lots of fixes in the Berkeley DB interface.
- `UPDATE IGNORE` does not abort if an update results in a `DUPLICATE_KEY` error.
- Put `CREATE TEMPORARY TABLE` commands in the update log.
- Fixed bug in handling of masked IP numbers in the privilege tables.
- Fixed bug with `delay_key_write` tables and `CHECK TABLE`.
- Added `--replicate-do-db` and `--replicate-ignore-db` options to `mysqld`, to restrict which databases get replicated.
- Added `SQL_LOG_BIN` option.

### D.3.46. Changes in release 3.23.15 (08 May 2000)

- To start `mysqld` as `root`, you must now use the `--user=root` option.
- Added interface to Berkeley DB. (This is not yet functional; play with it at your own risk!)
- Replication between master and slaves.
- Fixed bug that other threads could steal a lock when a thread had a lock on a table and did a `FLUSH TABLES` command.
- Added the `slow_launch_time` variable and the `Slow_launch_threads` status variable to `mysqld`. These can be examined with `mysqladmin variables` and `mysqladmin extended-status`.
- Added functions `INET_NTOA()` and `INET_ATON()`.
- The default type of `IF()` now depends on the second and third arguments and not only on the second argument.
- Fixed case when `myisamchk` could go into a loop when trying to repair a crashed table.
- Don't write `INSERT DELAYED` to update log if `SQL_LOG_UPDATE=0`.
- Fixed problem with `REPLACE` on `HEAP` tables.
- Added possible character sets and time zone to `SHOW VARIABLES` output.
- Fixed bug in locking code that could result in locking problems with concurrent inserts under high load.
- Fixed a problem with `DELETE` of many rows on a table with compressed keys where MySQL scanned the index to find the rows.
- Fixed problem with `CHECK TABLE` on table with deleted keyblocks.

- Fixed a bug in reconnect (at the client side) where it didn't free memory properly in some contexts.
- Fixed problems in update log when using `LAST_INSERT_ID()` to update a table with an `AUTO_INCREMENT` key.
- Added `NULLIF()` function.
- Fixed bug when using `LOAD DATA INFILE` on a table with `BLOB/TEXT` columns.
- Optimized `MyISAM` to be faster when inserting keys in sorted order.
- `EXPLAIN SELECT . . .` now also prints out whether MySQL needs to create a temporary table or use file sorting when resolving the `SELECT`.
- Added optimization to skip `ORDER BY` parts where the part is a constant expression in the `WHERE` part. Indexes can now be used even if the `ORDER BY` doesn't match the index exactly, as long as all the unused index parts and all the extra `ORDER BY` columns are constants in the `WHERE` clause. See [Section 7.4.5, "How MySQL Uses Indexes"](#).
- `UPDATE` and `DELETE` on a whole unique key in the `WHERE` part are now faster than before.
- Changed `RAID_CHUNKSIZE` to be in 1024-byte increments.
- Fixed core dump in `LOAD_FILE(NULL)`.

### D.3.47. Changes in release 3.23.14 (09 April 2000)

- Added `mysqlbinlog` program for displaying binary log files in text format.
- Added `mysql_real_escape_string()` function to the MySQL C API.
- Fixed a bug in `CONCAT()` where one of the arguments was a function that returned a modified argument.
- Fixed a critical bug in `myisamchk`, where it updated the header in the index file when one only checked the table. This confused the `mysqld` daemon if it updated the same table at the same time. Now the status in the index file is only updated if one uses `--update-state`. With older `myisamchk` versions you should use `--read-only` when only checking tables, if there is the slightest chance that the `mysqld` server is working on the table at the same time!
- Fixed that `DROP TABLE` is logged in the update log.
- Fixed problem when searching on `DECIMAL()` key field where the column data contained leading zeros.
- Fix bug in `myisamchk` when the `AUTO_INCREMENT` column isn't the first key.
- Allow `DATETIME` in ISO8601 format: 2000-03-12T12:00:00
- Dynamic character sets. A `mysqld` binary can now handle many different character sets (you can choose which when starting `mysqld`).
- Added `REPAIR TABLE` statement.
- Added `mysql_thread_safe()` function to the MySQL C API.
- Added the `UMASK_DIR` environment variable.



- Added `CONNECTION_ID()` function to return the client connection thread ID.
- When using `=` on `BLOB` or `VARCHAR BINARY` keys, where only a part of the column was indexed, the whole column of the result row wasn't compared.
- Fix for `sjis` character set and `ORDER BY`.
- When running in ANSI mode, don't allow columns to be used that aren't in the `GROUP BY` part.

### D.3.48. Changes in release 3.23.13 (14 March 2000)

- Fixed problem when doing locks on the same table more than 2 times in the same `LOCK TABLE` command; this fixed the problem one got when running the test-ATIS test with `--fast` or `-check-only-changed`.
- Added `SQL_BUFFER_RESULT` option to `SELECT`.
- Removed endspace from double/float numbers in results from temporary tables.
- Added `CHECK TABLE` command.
- Added changes for `MyISAM` in 3.23.12 that didn't get into the source distribution because of CVS problems.
- Fixed bug so that `mysqladmin shutdown` waits for the local server to close down.
- Fixed a possible endless loop when calculating timestamp.
- Added `print_defaults` program to the `.rpm` files. Removed `mysqlbug` from the client `.rpm` file.

### D.3.49. Changes in release 3.23.12 (07 March 2000)

- Fixed bug in `MyISAM` involving `REPLACE ... SELECT ...` which could give a corrupted table.
- Fixed bug in `myisamchk` where it incorrectly reset the `AUTO_INCREMENT` value.
- LOTS of patches for Linux Alpha. MySQL now appears to be relatively stable on Alpha.
- Changed `DISTINCT` on `HEAP` temporary tables to use hashed keys to quickly find duplicated rows. This mostly concerns queries of type `SELECT DISTINCT ... GROUP BY ...`. This fixes a problem where not all duplicates were removed in queries of the above type. In addition, the new code is MUCH faster.
- Added patches to make MySQL compile on Mac OS X.
- Added `IF NOT EXISTS` clause to `CREATE DATABASE`.
- Added `--all-databases` and `--databases` options to `mysqldump` to allow dumping of many databases at the same time.
- Fixed bug in compressed `DECIMAL()` index in `MyISAM` tables.
- Fixed bug when storing 0 into a timestamp.

- When doing `mysqladmin shutdown` on a local connection, `mysqladmin` now waits until the PID file is gone before terminating.
- Fixed core dump with some `COUNT(DISTINCT ...)` queries.
- Fixed that `myisamchk` works properly with RAID tables.
- Fixed problem with `LEFT JOIN` and `key_col IS NULL`.
- Fixed bug in `net_clear()` which could give the error `Aborted connection` in the MySQL clients.
- Added options `USE INDEX (key_list)` and `IGNORE INDEX (key_list)` as parameters in `SELECT`.
- `DELETE` and `RENAME` should now work on RAID tables.

### D.3.50. Changes in release 3.23.11 (16 February 2000)

- Added `HIGH_PRIORITY` option to `INSERT`. This overrides the effect of the `-low-priority-updates` server option and does not perform concurrent inserts.
- Allow the `ALTER TABLE tbl_name ADD (field_list)` syntax.
- Fixed problem with optimizer that could sometimes use incorrect keys.
- Fixed that `GRANT/REVOKE ALL PRIVILEGES` doesn't affect `GRANT OPTION`.
- Removed extra `'` from the output of `SHOW GRANTS`.
- Fixed problem when storing numbers in timestamps.
- Fix problem with time zones that have half hour offsets.
- Allow the syntax `UNIQUE INDEX` in `CREATE` statements.
- `mysqlhotcopy` - fast online hot-backup utility for local MySQL databases. By Tim Bunce.
- New more secure `mysqlaccess`. Thanks to Steve Harvey for this.
- Added `--i-am-a-dummy` and `--safe-updates` options to `mysql`.
- Added `select_limit` and `max_join_size` variables to `mysql`.
- Added `SQL_MAX_JOIN_SIZE` and `SQL_SAFE_UPDATES` options.
- Added `READ LOCAL` lock that doesn't lock the table for concurrent inserts. (This is used by `mysqldump`.)
- Changed that `LOCK TABLES ... READ` no longer allows concurrent inserts.
- Added `--skip-delay-key-write` option to `mysqld`.
- Fixed security problem in the protocol regarding password checking.
- `_rowid` can now be used as an alias for an integer type unique indexed column.

- Added back blocking of `SIGPIPE` when compiling with `--thread-safe-clients` to make things safe for old clients.

### D.3.51. Changes in release 3.23.10 (30 January 2000)

- Fixed bug in 3.23.9 where memory wasn't properly freed when using `LOCK TABLES`.

### D.3.52. Changes in release 3.23.9 (29 January 2000)

- Fixed problem that affected queries that did arithmetic on group functions.
- Fixed problem with timestamps and `INSERT DELAYED`.
- Fixed that `date_col BETWEEN const_date AND const_date` works.
- Fixed problem when only changing a 0 to `NULL` in a table with `BLOB/TEXT` columns.
- Fixed bug in range optimizer when using many key parts and or on the middle key parts: `WHERE K1=1 and K3=2 and (K2=2 and K4=4 or K2=3 and K4=5)`
- Added `source` command to `mysql` to allow reading of batch files inside the `mysql` client. Original patch by Matthew Vanecek.
- Fixed critical problem with the `WITH GRANT OPTION` option.
- Don't give an unnecessary `GRANT` error when using tables from many databases in the same query.
- Added VIO wrapper (needed for SSL support; by Andrei Errapart and Tõnu Samuel).
- Fixed optimizer problem on `SELECT` when using many overlapping indexes. MySQL should now be able to choose keys even better when there are many keys to choose from.
- Changed optimizer to prefer a range key instead of a ref key when the range key can use more columns than the ref key (which only can use columns with `=`). For example, the following type of queries should now be faster: `SELECT * from key_part_1=const and key_part_2 > const2`
- Fixed bug that a change of all `VARCHAR` columns to `CHAR` columns didn't change row type from dynamic to fixed.
- Disabled floating-point exceptions for FreeBSD to fix core dump when doing `SELECT FLOOR(POW(2,63))`.
- Renamed `mysqld` startup option from `--delay-key-write` to `--delay-key-write-for-all-tables`.
- Added `read-next-on-key` to `HEAP` tables. This should fix all problems with `HEAP` tables when using non-`UNIQUE` keys.
- Added option to print default arguments to all clients.
- Added `--log-slow-queries` option to `mysqld` to log all queries that take a long time to a separate log file with a time indicating how long the query took.

- Fixed core dump when doing `WHERE key_col=RAND(...)`.
- Fixed optimization bug in `SELECT ... LEFT JOIN ... key_col IS NULL`, when `key_col` could contain `NULL` values.
- Fixed problem with 8-bit characters as separators in `LOAD DATA INFILE`.

### D.3.53. Changes in release 3.23.8 (02 January 2000)

- Fixed problem when handling indexfiles larger than 8GB.
- Added latest patches to MIT-pthreads for NetBSD.
- Fixed problem with time zones that are < GMT - 11.
- Fixed a bug when deleting packed keys in `NISAM`.
- Fixed problem with `ISAM` when doing some `ORDER BY ... DESC` queries.
- Fixed bug when doing a join on a text key which didn't cover the whole key.
- Option `--delay-key-write` didn't enable delayed key writing.
- Fixed update of `TEXT` column which involved only case changes.
- Fixed that `INSERT DELAYED` doesn't update timestamps that are given.
- Added function `YEARWEEK()` and options `x`, `X`, `v` and `V` to `DATE_FORMAT()`.
- Fixed problem with `MAX(indexed_column)` and `HEAP` tables.
- Fixed problem with `BLOB NULL` keys and `LIKE "prefix%"`.
- Fixed problem with `MyISAM` and fixed-length rows < 5 bytes.
- Fixed problem that could cause MySQL to touch freed memory when doing very complicated `GROUP BY` queries.
- Fixed core dump if you got a crashed table where an `ENUM` field value was too big.

### D.3.54. Changes in release 3.23.7 (10 December 1999)

- Fixed workaround under Linux to avoid problems with `pthread_mutex_timedwait()`, which is used with `INSERT DELAYED`. See [Section 2.13.1, "Linux Notes"](#).
- Fixed that one get a 'disk full' error message if one gets disk full when doing sorting (instead of waiting until we got more disk space).
- Fixed a bug in `MyISAM` with keys > 250 characters.
- In `MyISAM` one can now do an `INSERT` at the same time as other threads are reading from the table.
- Added `max_write_lock_count` variable to `mysqld` to force a `READ` lock after a certain number of `WRITE` locks.

- Inverted flag `delay_key_write` on `show variables`.
- Renamed `concurrency` variable to `thread_concurrency`.
- The following functions are now multi-byte-safe: `LOCATE(substr, str)`, `POSITION(substr IN str)`, `LOCATE(substr, str, pos)`, `INSTR(str, substr)`, `LEFT(str, len)`, `RIGHT(str, len)`, `SUBSTRING(str, pos, len)`, `SUBSTRING(str FROM pos FOR len)`, `MID(str, pos, len)`, `SUBSTRING(str, pos)`, `SUBSTRING(str FROM pos)`, `SUBSTRING_INDEX(str, delim, count)`, `RTRIM(str)`, `TRIM([[BOTH | TRAILING] [remstr] FROM] str)`, `REPLACE(str, from_str, to_str)`, `REVERSE(str)`, `INSERT(str, pos, len, newstr)`, `LCASE(str)`, `LOWER(str)`, `UCASE(str)` and `UPPER(str)`; patch by Wei He.
- Fix core dump when releasing a lock from a non-existent table.
- Remove locks on tables before starting to remove duplicates.
- Added option `FULL` to `SHOW PROCESSLIST`.
- Added option `--verbose` to `mysqladmin`.
- Fixed problem when automatically converting `HEAP` to `MyISAM`.
- Fixed bug in `HEAP` tables when doing insert + delete + insert + scan the table.
- Fixed bugs on Alpha with `REPLACE()` and `LOAD DATA INFILE`.
- Added `interactive_timeout` variable to `mysqld`.
- Changed the argument to `mysql_data_seek()` from `ulong` to `ulonglong`.

### D.3.55. Changes in release 3.23.6 (15 December 1999)

- Added `-O lower_case_table_names={0|1}` option to `mysqld` to allow users to force table names to lowercase.
- Added `SELECT ... INTO DUMPFILE`.
- Added `--ansi` option to `mysqld` to make some functions standard SQL compatible.
- Temporary table names now start with `#sql`.
- Added quoting of identifiers with ``` (" in `--ansi` mode).
- Changed to use `snprintf()` when printing floats to avoid some buffer overflows on FreeBSD.
- Made `FLOOR()` overflow safe on FreeBSD.
- Added `--quote-names` option to `mysqldump`.
- Fixed bug that one could make a part of a `PRIMARY KEY NOT NULL`.
- Fixed `encrypt()` to be thread-safe and not reuse buffer.
- Added `mysql_odbc_escape_string()` function to support big5 characters in MyODBC.
- Rewrote the storage engine to use classes. This introduces a lot of new code, but make table hand-

ling faster and better.

- Added patch by Sasha for user-defined variables.
- Changed that `FLOAT` and `DOUBLE` (without any length modifiers) no longer are fixed decimal point numbers.
- Changed the meaning of `FLOAT(X)`: Now this is the same as `FLOAT` if  $X \leq 24$  and a `DOUBLE` if  $24 < X \leq 53$ .
- `DECIMAL(X)` is now an alias for `DECIMAL(X, 0)` and `DECIMAL` is now an alias for `DECIMAL(10, 0)`. The same goes for `NUMERIC`.
- Added option `ROW_FORMAT={DEFAULT | DYNAMIC | FIXED | COMPRESSED}` to `CREATE TABLE`.
- `DELETE FROM tbl_name` didn't work on temporary tables.
- Changed function `CHAR_LENGTH()` to be multi-byte character safe.
- Added function `ORD(string)`.

### D.3.56. Changes in release 3.23.5 (20 October 1999)

- Fixed some Y2K problems in the new date handling in 3.23.
- Fixed problem with `SELECT DISTINCT ... ORDER BY RAND()`.
- Added patches by Sergei A. Golubchik for text searching on the `MyISAM` level.
- Fixed cache overflow problem when using full joins without keys.
- Fixed some configure issues.
- Some small changes to make parsing faster.
- Adding a column after the last field with `ALTER TABLE` didn't work.
- Fixed problem when using an `AUTO_INCREMENT` column in two keys
- With `MyISAM`, you now can have an `AUTO_INCREMENT` column as a key sub part: `CREATE TABLE foo (a INT NOT NULL AUTO_INCREMENT, b CHAR(5), PRIMARY KEY (b,a))`
- Fixed bug in `MyISAM` with packed char keys that could be `NULL`.
- `AS` on field name with `CREATE TABLE tbl_name SELECT ...` didn't work.
- Allow use of `NATIONAL` and `NCHAR` when defining character columns. This is the same as not using `BINARY`.
- Don't allow `NULL` columns in a `PRIMARY KEY` (only in `UNIQUE` keys).
- Clear `LAST_INSERT_ID()` if one uses this in ODBC: `WHERE auto_increment_column IS NULL`. This seems to fix some problems with Access.
- `SET SQL_AUTO_IS_NULL=0 | 1` now turns on/off the handling of searching for the last inserted

row with `WHERE auto_increment_column IS NULL`.

- Added new variable `concurrency` to `mysqld` for Solaris.
- Added `--relative` option to `mysqladmin` to make `extended-status` more useful to monitor changes.
- Fixed bug when using `COUNT(DISTINCT ...)` on an empty table.
- Added support for the Chinese character set GBK.
- Fixed problem with `LOAD DATA INFILE` and `BLOB` columns.
- Added bit operator `~` (negation).
- Fixed problem with user-defined functions.

### D.3.57. Changes in release 3.23.4 (28 September 1999)

- Inserting a `DATETIME` into a `TIME` column no longer try to store 'days' in it.
- Fixed problem with storage of float/double on little endian machines. (This affected `SUM( )`.)
- Added connect timeout on TCP/IP connections.
- Fixed problem with `LIKE "%"` on an index that may have `NULL` values.
- `REVOKE ALL PRIVILEGES` didn't revoke all privileges.
- Allow creation of temporary tables with same name as the original table.
- When granting an account a `GRANT` option for a database, the account couldn't grant privileges to other users.
- New statement: `SHOW GRANTS FOR user` (by Sinisa).
- New `date_add` syntax: `date/datetime + INTERVAL # interval_type`. By Joshua Chamas.
- Fixed privilege check for `LOAD DATA REPLACE`.
- Automatic fixing of broken include files on Solaris 2.7
- Some configure issues to fix problems with big filesystem detection.
- `REGEXP` is now case-insensitive if you use non-binary strings.

### D.3.58. Changes in release 3.23.3 (13 September 1999)

- Added patches for MIT-pthreads on NetBSD.
- Fixed range bug in `MyISAM`.
- `ASC` is now the default again for `ORDER BY`.

- Added `LIMIT` to `UPDATE`.
- Added `mysql_change_user()` function to the MySQL C API.
- Added character set to `SHOW VARIABLES`.
- Added support of `--[whitespace]` comments.
- Allow `INSERT INTO tbl_name VALUES ()`, that is, you may now specify an empty value list to insert a row in which each column is set to its default value.
- Changed `SUBSTRING(text FROM pos)` to conform to standard SQL. (Before this construct returned the rightmost `pos` characters.)
- `SUM()` with `GROUP BY` returned 0 on some systems.
- Changed output for `SHOW TABLE STATUS`.
- Added `DELAY_KEY_WRITE` option to `CREATE TABLE`.
- Allow `AUTO_INCREMENT` on any key part.
- Fixed problem with `YEAR(NOW())` and `YEAR(CURDATE())`.
- Added `CASE` construct.
- New `COALESCE()` function.

### D.3.59. Changes in release 3.23.2 (09 August 1999)

- Fixed range optimizer bug: `SELECT * FROM tbl_name WHERE key_part1 >= const AND (key_part2 = const OR key_part2 = const)`. The bug was that some rows could be duplicated in the result.
- Running `myisamchk` without `-a` updated the index distribution incorrectly.
- `SET SQL_LOW_PRIORITY_UPDATES=1` was causing a parse error.
- You can now update index columns that are used in the `WHERE` clause. `UPDATE tbl_name SET KEY=KEY+1 WHERE KEY > 100`
- Date handling should now be a bit faster.
- Added handling of fuzzy dates (dates where day or month is 0), such as `'1999-01-00'`.
- Fixed optimization of `SELECT ... WHERE key_part1=const1 AND key_part_2=const2 AND key_part1=const4 AND key_part2=const4`; indextype should be `range` instead of `ref`.
- Fixed `egcs` 1.1.2 optimizer bug (when using `BLOB` values) on Linux Alpha.
- Fixed problem with `LOCK TABLES` combined with `DELETE FROM table`.
- `MyISAM` tables now allow keys on `NULL` and `BLOB/TEXT` columns.
- The following join is now much faster: `SELECT ... FROM t1 LEFT JOIN t2 ON ... WHERE t2.not_null_column IS NULL`.



- `ORDER BY` and `GROUP BY` can be done on functions.
- Changed handling of 'const\_item' to allow handling of `ORDER BY RAND()`.
- Indexes are now used for `WHERE key_column = function`.
- Indexes are now used for `WHERE key_column = col_name` even if the columns are not identically packed.
- Indexes are now used for `WHERE col_name IS NULL`.
- Changed heap tables to be stored in `low_byte_first` order (to make it easy to convert to `MyISAM` tables)
- Automatic change of `HEAP` temporary tables to `MyISAM` tables in case of “table is full” errors.
- Added `--init-file=file_name` option to `mysqld`.
- Added `COUNT(DISTINCT value, [value, ...])`.
- `CREATE TEMPORARY TABLE` now creates a temporary table, in its own namespace, that is automatically deleted if connection is dropped.
- New reserved words (required for `CASE`): `CASE`, `THEN`, `WHEN`, `ELSE` and `END`.
- New functions `EXPORT_SET()` and `MD5()`.
- Support for the GB2312 Chinese character set.

### D.3.60. Changes in release 3.23.1 (08 July 1999)

- Fixed some compilation problems.

### D.3.61. Changes in release 3.23.0 (05 July 1999: Alpha)

- A new storage engine library (`MyISAM`) with a lot of new features. See [Section 14.1, “The MyISAM Storage Engine”](#).
- You can create in-memory `HEAP` tables which are extremely fast for lookups.
- Support for big files (63-bit) on OSs that support big files.
- New function `LOAD_FILE(filename)` to get the contents of a file as a string value.
- New `<=>` operator that acts as `=` but returns `TRUE` if both arguments are `NULL`. This is useful for comparing changes between tables.
- Added the ODBC 3.0 `EXTRACT(interval FROM datetime)` function.
- Columns defined as `FLOAT(X)` are not rounded on storage and may be in scientific notation (1.0E+10) when retrieved.
- `REPLACE` is now faster than before.
- Changed `LIKE` character comparison to behave as `=`; This means that `'e' LIKE 'é'` is now true.

(If the line doesn't display correctly, the latter 'e' is a French 'e' with an acute accent above.)

- `SHOW TABLE STATUS` returns a lot of information about the tables.
- Added `LIKE` to the `SHOW STATUS` command.
- Added `Privileges` column to `SHOW COLUMNS`.
- Added `Packed` and `Comment` columns to `SHOW INDEX`.
- Added comments to tables (with `CREATE TABLE ... COMMENT 'xxx'`).
- Added `UNIQUE`, as in `CREATE TABLE tbl_name (col INT NOT NULL UNIQUE)`
- New create syntax: `CREATE TABLE tbl_name SELECT ...`
- New create syntax: `CREATE TABLE IF NOT EXISTS ...`
- Allow creation of `CHAR(0)` columns.
- `DATE_FORMAT()` now requires `'%'` before any format character.
- `DELAYED` is now a reserved word (sorry about that :()).
- An example procedure is added: `analyse`, file: `sql_analyse.c`. This describes the data in your query. Try the following:

```
SELECT ... FROM ...
WHERE ... PROCEDURE ANALYSE([max_elements],[max_memory])
```

This procedure is extremely useful when you want to check the data in your table!

- `BINARY` cast to force a string to be compared in case-sensitive fashion.
- Added `--skip-show-database` option to `mysqld`.
- Check whether a row has changed in an `UPDATE` now also works with `BLOB/TEXT` columns.
- Added the `INNER` join syntax. **Note:** This made `INNER` a reserved word!
- Added support for netmasks to the hostname in the MySQL grant tables. You can specify a netmask using the `IP/NETMASK` syntax.
- If you compare a `NOT NULL DATE/DATETIME` column with `IS NULL`, this is changed to a compare against `0` to satisfy some ODBC applications. (By <shreeve@uci.edu>.)
- `NULL IN (...)` now returns `NULL` instead of `0`. This ensures that `null_column NOT IN (...)` doesn't match `NULL` values.
- Fix storage of floating-point values in `TIME` columns.
- Changed parsing of `TIME` strings to be more strict. Now the fractional second part is detected (and currently skipped). The following formats are supported:
  - `[[DAYS] [H]H:]MM:]SS[.fraction]`
  - `[[[[[H]H]H]H]MM]SS[.fraction]`
- Detect (and ignore) fractional second part from `DATETIME`.

- Added the `LOW_PRIORITY` attribute to `LOAD DATA INFILE`.
- The default index name now uses the same case as the column name on which the index name is based.
- Changed default number of connections to 100.
- Use bigger buffers when using `LOAD DATA INFILE`.
- `DECIMAL(x,y)` now works according to standard SQL.
- Added aggregate user-defined functions. Thanks to Andreas F. Bobak ([<bobak@relog.ch>](mailto:bobak@relog.ch)) for this!
- `LAST_INSERT_ID()` is now updated for `INSERT INTO ... SELECT`.
- Some small changes to the join table optimizer to make some joins faster.
- `SELECT DISTINCT` is much faster; it uses the new `UNIQUE` functionality in `MyISAM`. One difference compared to MySQL 3.22 is that the output of `DISTINCT` is no longer sorted.
- All C client API macros are now functions to make shared libraries more reliable. Because of this, you can no longer call `mysql_num_fields()` on a `MYSQL` object, you must use `mysql_field_count()` instead.
- Added use of `LIBWRAP`; patch by Henning P. Schmiedehausen.
- Don't allow `AUTO_INCREMENT` for other than numerical columns.
- Using `AUTO_INCREMENT` now automatically makes the column `NOT NULL`.
- Show `NULL` as the default value for `AUTO_INCREMENT` columns.
- Added `SQL_BIG_RESULT`; `SQL_SMALL_RESULT` is now default.
- Added a shared library RPM. This enhancement was contributed by David Fox ([<dsfox@cogsci.ucsd.edu>](mailto:dsfox@cogsci.ucsd.edu)).
- Added `--enable-large-files` and `--disable-large-files` options to `configure`. See `configure.in` for some systems where this is automatically turned off because of broken implementations.
- Upgraded `readline` to 4.0.
- New `CREATE TABLE` options: `PACK_KEYS` and `CHECKSUM`.
- Added `--default-table-type` option to `mysqld`.

## D.4. Changes in InnoDB

Starting from 4.0.22 and 4.1.5, all InnoDB changes are included in the MySQL Change History, and this manual section is no longer separately maintained.

### D.4.1. Changes in MySQL/InnoDB-4.0.21, September 10, 2004

Functionality added or changed:

- Renamed the `innodb.status.<pid>` files (created in the data directory) to `innodb_status.<pid>`. This avoids problems on filesystems that do not allow multiple periods in filenames.
- Added `innodb_status_file` system variable to `mysqld` to control whether output from `SHOW INNODB STATUS` is written to a `innodb_status.<pid>` file in the data directory. By default, the file is not created. To create it, start `mysqld` with the `--innodb_status_file=1` option.
- Changes for NetWare to exit InnoDB gracefully on NetWare even in a case of an assertion failure, instead of intentionally crashing the `'mysqld'` server process.

Bugs fixed:

- Fixed a bug in `ON DELETE CASCADE` and `ON UPDATE CASCADE` foreign key constraints: long chains of cascaded operations would cause a stack overflow and crash the server. Cascaded operations are now limited to 15 levels. ([Bug#4446](http://bugs.mysql.com/4446) [<http://bugs.mysql.com/4446>])
- Fixed a possible bug in `LOCK TABLES` introduced in MySQL/InnoDB-4.0.19: The count of tables explicitly locked by a transaction was incremented only after the locks were granted, but decremented when the lock structures were destroyed.
- Fixed a bug in `UNLOCK TABLES` in `AUTOCOMMIT=0` mode, introduced in MySQL/InnoDB-4.0.19: The memory allocated for some locks acquired by the transaction could be deallocated before those locks were released. The bug can lead to crashes and memory corruption of the buffer pool when the transaction acquires a large number of locks (table locks or row-level locks).
- Increment the InnoDB watchdog timeout during `CHECK TABLE`. A long-running `CHECK TABLE` would cause InnoDB to complain about a 'long semaphore wait', and crash the server, if a query had to wait more than 600 seconds behind that `CHECK TABLE` operation. ([Bug#2694](http://bugs.mysql.com/2694) [<http://bugs.mysql.com/2694>])
- If you configure `innodb_additional_mem_pool_size` so small that InnoDB memory allocation spills over from it, then every 4 billionth spill may cause memory corruption. A symptom is a printout like the one following in the `.err` log. The workaround is to make `innodb_additional_mem_pool_size` big enough to hold all memory allocation. Use `SHOW INNODB STATUS` to determine that there is plenty of free space available in the additional mem pool, and the total allocated memory stays rather constant.

```
InnoDB: Error: Mem area size is 0. Possibly a memory overrun of the
InnoDB: previous allocated area!
InnoDB: Apparent memory corruption: mem dump len 500; hex
```

- The special meaning of the table names `innodb_monitor`, `innodb_lock_monitor`, `innodb_tablespace_monitor`, `innodb_table_monitor`, and `innodb_validate` in `CREATE TABLE` and `DROP TABLE` statements was accidentally removed in MySQL/InnoDB-4.0.19. The diagnostic functions attached to these special table names (see [Section 14.2.12.1, “SHOW ENGINE INNODB STATUS and the InnoDB Monitors”](#)) are accessible again in MySQL/InnoDB-4.0.21.
- When the private SQL parser of InnoDB was modified in MySQL/InnoDB-4.0.19 in order to allow the use of the apostrophe (‘ ’) in table and column names, the fix relied on a previously unused function `mem_realloc()`, whose implementation was incorrect. As a result, InnoDB can incorrectly parse column and table names as the empty string. The InnoDB `realloc()` implementation has been corrected in MySQL/InnoDB-4.0.21.
- Fixed a glitch introduced in 4.0.18 and 4.1.2: in `SHOW TABLE STATUS` InnoDB systematically overestimated the row count by 1 if the table fit on a single 16 kB data page.

- InnoDB created temporary files with the C library function `tmpfile()`. On Windows, the files would be created in the root directory of the current file system. To correct this behavior, the invocations of `tmpfile()` were replaced with code that uses the function `create_temp_file()` in the MySQL portability layer. (Bug#3998 [http://bugs.mysql.com/3998])
- If `ALTER TABLE ... DROP FOREIGN KEY ...` fails because of a wrong constraint name, return a table handler error number 150 instead of 152.
- If there was little file I/O in InnoDB, but the insert buffer was used, it could happen that 'Pending normal aio reads' was bigger than 0, but the I/O handler thread did not get waken up in 600 seconds. This resulted in a hang, and crashing of InnoDB.
- If we `RENAME`d a table, InnoDB forgot to load the `FOREIGN KEY` constraints that reference the new table name, and forgot to check that they are compatible with the table.

## D.4.2. Changes in MySQL/InnoDB-4.1.4, August 31, 2004

Functionality added or changed:

- **Important:** Made internal representation of `TIMESTAMP` values in InnoDB in 4.1 to be the same as in 4.0. This difference resulted in incorrect datetime values in `TIMESTAMP` columns in InnoDB tables after an upgrade from 4.0 to 4.1. (Bug#4492 [http://bugs.mysql.com/4492]) **Warning: extra steps during upgrade required!** This means that if you are upgrading from 4.1.x, where  $x \leq 3$ , to 4.1.4 you should use `mysqldump` for saving and then restoring your InnoDB tables with `TIMESTAMP` columns. No conversion is needed if you upgrade from 3.23 or 4.0 to 4.1.4 or later.
- Added a new startup option `innodb_locks_unsafe_for_binlog`. This option forces InnoDB not to use next-key locking in searches and index scans.
- Added `innodb_status_file` system variable to `mysqld` to control whether output from `SHOW INNODB STATUS` is written to a `innodb_status.<pid>` file in the data directory. By default, the file is not created. To create it, start `mysqld` with the `--innodb_status_file=1` option.
- Changes for NetWare to exit InnoDB gracefully on NetWare even in a case of an assertion failure, instead of intentionally crashing the `mysqld` server process.
- “Gap” type row locks without the `LOCK_INSERT_INTENTION` flag do not need to wait for anything. This is because different users can have conflicting lock types on gaps. This change reduces unnecessary deadlocks.

Bugs fixed:

- Fixed a bug in `ON DELETE CASCADE` and `ON UPDATE CASCADE` foreign key constraints: long chains of cascaded operations would cause a stack overflow and crash the server. Cascaded operations are now limited to 15 levels. (Bug#4446 [http://bugs.mysql.com/4446])
- Increment the InnoDB watchdog timeout during `CHECK TABLE`. (Bug#2694 [http://bugs.mysql.com/2694])
- If you configure `innodb_additional_mem_pool_size` so small that InnoDB memory allocation spills over from it, then every 4 billionth spill may cause memory corruption. A symptom is a printout like the one following in the `.err` log.

```
InnoDB: Error: Mem area size is 0. Possibly a memory overrun of the
InnoDB: previous allocated area!
```

```
InnoDB: Apparent memory corruption: mem dump len 500; hex
```

- Fixed a glitch introduced in 4.0.18 and 4.1.2: in `SHOW TABLE STATUS` InnoDB systematically overestimated the row count by 1 if the table fit on a single 16 kB data page.
- InnoDB created temporary files with the C library function `tmpfile()`. On Windows, the files would be created in the root directory of the current file system. To correct this behavior, the invocations of `tmpfile()` were replaced with code that uses the function `create_temp_file()` in the MySQL portability layer. (Bug#3998 [http://bugs.mysql.com/3998])
- If we `RENAME`d a table, InnoDB forgot to load the foreign key constraints that reference the new table name, and forgot to check that they are compatible with the table.
- If there was little file I/O in InnoDB, but the insert buffer was used, it could happen that 'Pending normal aio reads' was bigger than 0, but the I/O handler thread did not get waken up in 600 seconds. This resulted in a hang, and an intentional crashing of `mysqld`.

### D.4.3. Changes in MySQL/InnoDB-4.1.3, June 28, 2004

Functionality added or changed:

- **Important:** Starting from MySQL 4.1.3, InnoDB uses the same character set comparison functions as MySQL for non-`latin1_swedish_ci` character strings that are not `BINARY`. This changes the sorting order of space and characters < ASCII(32) in those character sets. For `latin1_swedish_ci` character strings and `BINARY` strings, InnoDB uses its own pad-spaces-at-end comparison method, which stays unchanged. If you have an InnoDB table created with MySQL 4.1.2 or earlier, with an index on a non-`latin1` character set (in the case of 4.1.0 and 4.1.1 with any character set) `CHAR/VARCHAR`/or `TEXT` column that is not `BINARY` but may contain characters < ASCII(32), then you should do `ALTER TABLE` or `OPTIMIZE TABLE` on it to **regenerate the index, after upgrading to MySQL 4.1.3 or later**.
- `OPTIMIZE TABLE` for InnoDB tables is now mapped to `ALTER TABLE` rather than to `ANALYZE TABLE`.
- Added an interface for storing the binlog offset in the InnoDB log and flushing the log.

Bugs fixed:

- The **critical bug in 4.1.2** (crash recovery skipping all `.ibd` files if you specify `innodb_file_per_table` on Unix) has been fixed. The bug was a combination of two bugs. Crash recovery ignored the files, because the attempt to lock them in the wrong mode failed. From now on, locks are only obtained for regular files opened in read/write mode, and crash recovery stops if an `.ibd` file for a table exists in a database directory but is inaccessible.
- Do not remember the original `select_lock_type` inside `LOCK TABLES`. (Bug#4047 [http://bugs.mysql.com/4047])
- The special meaning of the table names `innodb_monitor`, `innodb_lock_monitor`, `innodb_tablespace_monitor`, `innodb_table_monitor`, and `innodb_validate` in `CREATE TABLE` and `DROP TABLE` statements was accidentally removed in MySQL/InnoDB-4.1.2. The diagnostic functions attached to these special table names (see Section 14.2.12.1, “`SHOW ENGINE INNODB STATUS` and the InnoDB Monitors”) are accessible again in MySQL/InnoDB-4.1.3.

- When the private SQL parser of InnoDB was modified in MySQL/InnoDB-4.0.19 in order to allow the use of the apostrophe ( ' ') in table and column names, the fix relied on a previously unused function `mem_realloc()`, whose implementation was incorrect. As a result, InnoDB can incorrectly parse column and table names as the empty string. The InnoDB `realloc()` implementation has been corrected in MySQL/InnoDB-4.1.3.
- In a clean-up of MySQL/InnoDB-4.1.2, the code for invalidating the query cache was broken. Now the query cache should be correctly invalidated for tables affected by `ON UPDATE CASCADE` or `ON DELETE CASCADE` constraints.
- Fixed a bug: in `LIKE 'abc%'`, the `'%'` did not match the empty string if the character set was not `latin1_swedish_ci`. This bug was fixed by changing the sorting order in these character sets. See the above note about data conversion in 4.1.3.

## D.4.4. Changes in MySQL/InnoDB-4.1.2, May 30, 2004

**NOTE:** CRITICAL BUG in 4.1.2 if you specify `innodb_file_per_table` in `my.cnf` on Unix. In crash recovery InnoDB skips the crash recovery for all `.ibd` files and those tables become CORRUPT! The symptom is a message `Unable to lock ...ibd with lock 1, error: 9: fcntl: Bad file descriptor` in the `.err` log in crash recovery.

Functionality added or changed:

- Support multiple character sets. Note that tables created in other collations than `latin1_swedish_ci` cannot be accessed in MySQL/InnoDB 4.0.
- Automatically create a suitable index on a `FOREIGN KEY`, if the user does not create one. Removes most of the cases of `Error 1005 (errno 150)` in table creation.
- Do not assert in `log0log.c`, line 856 if `ib_logfiles` are too small for `innodb_thread_concurrency`. Instead, print instructions how to adjust `my.cnf` and call `exit(1)`.
- If MySQL tries to `SELECT` from an InnoDB table without setting any table locks, print a descriptive error message and assert; some subquery bugs were of this type.
- Allow a key part length in InnoDB to be up to 3,500 bytes; this is needed so that you can create an index on a column with 255 UTF-8 characters.
- All new features from InnoDB-4.0.17, InnoDB-4.0.18, InnoDB-4.0.19 and InnoDB-4.0.20.

Bugs fixed:

- If you configure `innodb_additional_mem_pool_size` so small that InnoDB memory allocation spills over from it, then every 4 billionth spill may cause memory corruption. A symptom is a printout like the one following in the `.err` log.

```
InnoDB: Error: Mem area size is 0. Possibly a memory overrun of the
InnoDB: previous allocated area!
InnoDB: Apparent memory corruption: mem dump len 500; hex
```

- Improved portability to 64-bit platforms, especially Win64.
- Fixed an assertion failure when a purge of a table was not possible because of missing `.ibd` file.

- Fixed a bug: do not retrieve all columns in a table if we only need the 'ref' of the row (usually, the `PRIMARY KEY`) to calculate an `ORDER BY`. ([Bug#1942](http://bugs.mysql.com/1942) [<http://bugs.mysql.com/1942>])
- On Unix-like systems, obtain an exclusive advisory lock on InnoDB files, to prevent corruption when multiple instances of MySQL are running on the same set of data files. The Windows version of InnoDB currently takes a mandatory lock on the files. ([Bug#3608](http://bugs.mysql.com/3608) [<http://bugs.mysql.com/3608>])
- Added a missing space to the output format of `SHOW INNODB STATUS`; reported by Jocelyn Fournier.
- All bugfixes from InnoDB-4.0.17, InnoDB-4.0.18, InnoDB-4.0.19 and InnoDB-4.0.20.

## D.4.5. Changes in MySQL/InnoDB-4.0.20, May 18, 2004

Bugs fixed:

- Apostrophe characters now are recognized by the internal InnoDB parser and can be used within quoted table and column identifiers in `FOREIGN KEY` clauses.
- Make `LOCK TABLE` aware of InnoDB row-level locks and InnoDB aware of locks set with `LOCK TABLE`. ([Bug#3299](http://bugs.mysql.com/3299) [<http://bugs.mysql.com/3299>])
- Fixed race conditions in `SHOW INNODB STATUS`. ([Bug#3596](http://bugs.mysql.com/3596) [<http://bugs.mysql.com/3596>])

## D.4.6. Changes in MySQL/InnoDB-4.0.19, May 4, 2004

Functionality added or changed:

- Better error message when the server has to crash because the buffer pool is exhausted by the lock table or the adaptive hash index.
- Print always the count of pending `pread()` and `pwrite()` calls if there is a long semaphore wait.
- Improve space utilization when rows of 1,500 to 8,000 bytes are inserted in the order of the primary key.
- Remove potential buffer overflow errors by sending diagnostic output to `stderr` or files instead of `stdout` or fixed-size memory buffers. As a side effect, the output of `SHOW INNODB STATUS` is written to a file `<datadir>/innodb.status.<pid>` every 15 seconds.

Bugs fixed:

- Fixed a bug: `DROP DATABASE` did not work if `FOREIGN KEY` references were defined within the database. ([Bug#3058](http://bugs.mysql.com/3058) [<http://bugs.mysql.com/3058>])
- Remove unnecessary files, functions and variables. Many of these were needed in the standalone version of InnoDB. Remove debug functions and variables from non-debug build.
- Add diagnostic code to analyze an assertion failure in `ha_innodb.cc` on line 2020 reported by a user. ([Bug#2903](http://bugs.mysql.com/2903) [<http://bugs.mysql.com/2903>])
- Fixed a bug: in a `FOREIGN KEY, ON UPDATE CASCADE` was not triggered if the update changed



a string to another value identical in alphabetical ordering, for example, 'abc' -> 'aBc'.

- Protect the reading of the latest foreign key error explanation buffer with a mutex; in theory, a race condition could cause `SHOW INNODB STATUS` print garbage characters after the error info.
- Fixed a bug: The row count and key cardinality estimate was grossly too small if each clustered index page only contained one record.
- Parse `CONSTRAINT FOREIGN KEY` correctly. ([Bug#3332](http://bugs.mysql.com/3332) [<http://bugs.mysql.com/3332>])
- Fixed a memory corruption bug on Windows. The bug is present in all InnoDB versions in Windows, but it depends on how the linker places a static array in `srv0srv.c`, whether the bug shows itself. 4 bytes were overwritten with a pointer to a statically allocated string 'get windows aio return value'.
- Fix a glitch reported by Philippe Lewicki on the general mailing list: do not print a warning to the `.err` log if `read_key` fails with a lock wait timeout error 146.
- Allow quotes to be embedded in strings in the private SQL parser of InnoDB, so that ' ' can be used in InnoDB table and column names. Display quotes within identifiers properly.
- Debugging: Allow `UNIV_SYNC_DEBUG` to be disabled while `UNIV_DEBUG` is enabled.
- Debugging: Handle magic numbers in a more consistent way.

## D.4.7. Changes in MySQL/InnoDB-4.0.18, February 13, 2004

- Do not allow dropping a table referenced by a `FOREIGN KEY` constraint, unless the user does `SET FOREIGN_KEY_CHECKS=0`. The error message here is somewhat misleading "Cannot delete or update a parent row...", and must be changed in a future version 4.1.x.
- Make InnoDB to remember the `CONSTRAINT` name given by a user for a `FOREIGN KEY`.
- Change the print format of `FOREIGN KEY` constraints spanning multiple databases to ``db_name`.`tbl_name``. But when parsing them, we must also accept ``db_name.tbl_name``, because that was the output format in < 4.0.18.
- An optimization in locking: If `AUTOCOMMIT=1`, then we do not need to make a plain `SELECT` set shared locks even on the `SERIALIZABLE` isolation level, because we know that the transaction is read-only. A read-only transaction can always be performed on the `REPEATABLE READ` level, and that does not endanger the serializability.
- Implement an automatic downgrade from >= 4.1.1 -> 4.0.18 if the user has not created tables in `.ibd` files or used other 4.1.x features. **Consult** the manual section on **multiple tablespaces** carefully if you want to downgrade!
- Fixed a bug: MySQL should not allow `REPLACE` to internally perform an `UPDATE` if the table is referenced by a `FOREIGN KEY`. The MySQL manual states that `REPLACE` must resolve a duplicate-key error semantically with `DELETE + INSERT`, and not by an `UPDATE`. In versions < 4.0.18 and < 4.1.2, MySQL could resolve a duplicate key conflict in `REPLACE` by doing an `UPDATE` on the existing row, and `FOREIGN KEY` checks could behave in a semantically wrong way. ([Bug#2418](http://bugs.mysql.com/2418) [<http://bugs.mysql.com/2418>])
- Fixed a bug: generate `FOREIGN KEY` constraint identifiers locally for each table, in the form `db_name/tbl_name_ibfk_number`. If the user gives the constraint name explicitly, then remember it. These changes should ensure that foreign key id's in a slave are the same as in the master,

and `DROP FOREIGN KEY` does not break replication. ([Bug#2167](http://bugs.mysql.com/2167) [<http://bugs.mysql.com/2167>])

- Fixed a bug: allow quoting of identifiers in InnoDB's `FOREIGN KEY` definitions with a backtick ( ` ) and a double quote ( " ). You can now use also spaces in table and column names, if you quote the identifiers. ([Bug#1725](http://bugs.mysql.com/1725) [<http://bugs.mysql.com/1725>], [Bug#2424](http://bugs.mysql.com/2424) [<http://bugs.mysql.com/2424>])
- Fixed a bug: `FOREIGN KEY ... ON UPDATE/DELETE NO ACTION` must check the foreign key constraint, not ignore it. Since we do not have deferred constraints in InnoDB, this bugfix makes InnoDB to check `NO ACTION` constraints immediately, like it checks `RESTRICT` constraints.
- Fixed a bug: InnoDB crashed in `RENAME TABLE` if `db_name.tbl_name` is shorter than 5 characters. ([Bug#2689](http://bugs.mysql.com/2689) [<http://bugs.mysql.com/2689>])
- Fixed a bug: in `SHOW TABLE STATUS`, InnoDB row count and index cardinality estimates wrapped around at 512 million in 32-bit computers. Note that unless MySQL is compiled with the `BIG_TABLES` option, they still wrap around at 4 billion.
- Fixed a bug: If there was a `UNIQUE` secondary index, and `NULL` values in that unique index, then with the `IS NULL` predicate, InnoDB returned only the first matching row, though there can be many. This bug was introduced in 4.0.16. ([Bug#2483](http://bugs.mysql.com/2483) [<http://bugs.mysql.com/2483>])

## D.4.8. Changes in MySQL/InnoDB-5.0.0, December 24, 2003

- **Important note:** If you upgrade to MySQL 4.1.1 or higher, it is difficult to downgrade back to 4.0 or 4.1.0! That is because, for earlier versions, InnoDB is not aware of multiple tablespaces.
- InnoDB in 5.0.0 is essentially the same as InnoDB-4.1.1 with the bugfixes of InnoDB-4.0.17 included.

## D.4.9. Changes in MySQL/InnoDB-4.0.17, December 17, 2003

- Fixed a bug: If you created a column prefix secondary index and updated it so that the last characters in the column prefix were spaces, InnoDB would assert in `row0upd.c`, line 713. The same assertion failed if you updated a column in an ordinary secondary index so that the new value was alphabetically equivalent, but had a different length. This could happen, for example, in the UTF8 character set if you updated a letter to its accented or umlaut form.
- Fixed a bug: InnoDB could think that a secondary index record was not locked though it had been updated to an alphabetically equivalent value, for example, 'abc' -> 'aBc'.
- Fixed a bug: If you updated a secondary index column to an alphabetically equivalent value, and rolled back your update, InnoDB failed to restore the field in the secondary index to its original value.
- There are still several outstanding non-critical bugs reported in the MySQL bugs database. Their fixing has been delayed, because resources were allocated to the 4.1.1 release.

## D.4.10. Changes in MySQL/InnoDB-4.1.1, December 4, 2003

- **Important note:** If you upgrade to MySQL 4.1.1 or higher, you cannot downgrade to a version

lower than 4.1.1 any more! That is because, for earlier versions, `InnoDB` is not aware of multiple tablespaces.

- Multiple tablespaces now available for `InnoDB`. You can store each `InnoDB` type table and its indexes into a separate `.ibd` file into a MySQL database directory, into the same directory where the `.frm` file is stored.
- The MySQL query cache now works for `InnoDB` tables also if `AUTOCOMMIT=0`, or the statements are enclosed inside `BEGIN ... COMMIT`.
- Reduced `InnoDB` memory consumption by a few megabytes if one sets the buffer pool size  $< 8\text{MB}$ .
- You can use raw disk partitions also in Windows.

### D.4.11. Changes in MySQL/InnoDB-4.0.16, October 22, 2003

- Fixed a bug: in contrary to what was said in the manual, in a locking read `InnoDB` set two record locks if a unique exact match search condition was used on a multi-column unique key. For a single column unique key it worked right.
- Fixed a bug: If you used the rename trick `#sql... -> rsql...` to recover a temporary table, `InnoDB` asserted in `row_mysql_lock_data_dictionary()`.
- There are several outstanding non-critical bugs reported in the MySQL bugs database. Their fixing has been delayed, because resources are allocated to the upcoming 4.1.1 release.

### D.4.12. Changes in MySQL/InnoDB-3.23.58, September 15, 2003

- Fixed a bug: `InnoDB` could make the index page directory corrupt in the first B-tree page splits after `mysqld` startup. A symptom would be an assertion failure in `page0page.c`, in function `page_dir_find_slot()`.
- Fixed a bug: `InnoDB` could in rare cases return an extraneous row if a rollback, purge, and a `SELECT` coincided.
- Fixed a possible hang over the `btr0sea.c` latch if `SELECT` was used inside `LOCK TABLES`.
- Fixed a bug: If a single `DELETE` statement first managed to delete some rows and then failed in a `FOREIGN KEY` error or a `Table is full` error, MySQL did not roll back the whole SQL statement as it should.

### D.4.13. Changes in MySQL/InnoDB-4.0.15, September 10, 2003

- Fixed a bug: If you updated a row so that the 8000 byte maximum length (without `BLOB` and `TEXT`) was exceeded, `InnoDB` simply removed the record from the clustered index. In a similar insert, `InnoDB` would leak reserved file space extents, which would only be freed at the next `mysqld` startup.

- Fixed a bug: If you used big `BLOB` values, and your log files were relatively small, `InnoDB` could in a big `BLOB` operation temporarily write over the log produced after the latest checkpoint. If `InnoDB` would crash at that moment, then the crash recovery would fail, because `InnoDB` would not be able to scan the log even up to the latest checkpoint. Starting from this version, `InnoDB` tries to ensure the latest checkpoint is young enough. If that is not possible, `InnoDB` prints a warning to the `.err` log of MySQL and advises you to make the log files bigger.
- Fixed a bug: setting `innodb_fast_shutdown=0` had no effect.
- Fixed a bug introduced in 4.0.13: If a `CREATE TABLE` ended in a comment, that could cause a memory overrun.
- Fixed a bug: If `InnoDB` printed `Operating system error number .. in a file operation` to the `.err` log in Windows, the error number explanation was wrong. Workaround: look at section 13.2 of <http://www.innodb.com/ibman.php> about Windows error numbers.
- Fixed a bug: If you created a column prefix `PRIMARY KEY` like in `t(a CHAR(200), PRIMARY KEY (a(10)))` on a fixed-length `CHAR` column, `InnoDB` would crash even in a simple `SELECT`. A `CHECK TABLE` would report the table as corrupt, also in the case where the created key was not `PRIMARY`.

## D.4.14. Changes in MySQL/InnoDB-4.0.14, July 22, 2003

- `InnoDB` now supports the `SAVEPOINT` and `ROLLBACK TO SAVEPOINT` SQL statements. See <http://www.innodb.com/ibman.php#Savepoints> for the syntax.
- You can now create column prefix keys like in `CREATE TABLE t (a BLOB, INDEX (a(10)))`.
- You can also use `O_DIRECT` as the `innodb_flush_method` on the latest versions of Linux and FreeBSD. Beware of possible bugs in those operating systems, though.
- Fixed the checksum calculation of data pages. Previously most OS file system corruption went unnoticed. Note that if you downgrade from version 4.0.14 or up to a version earlier than 4.0.14, `InnoDB` prints warnings in the first startup:

```
InnoDB: Warning: An inconsistent page in the doublewrite buffer
InnoDB: space id 2552202359 page number 8245, 127'th page in dblwr buf.
```

but that is not dangerous and can be ignored.

- Modified the buffer pool replacement algorithm so that it tries to flush modified pages if there are no replaceable pages in the last 10 % of the LRU list. This can reduce disk I/O if the workload is a mixture of reads and writes.
- The buffer pool checkpoint flush algorithm now tries to flush also close neighbors of the page at the end of the flush list. This can speed up database shutdown, and can also speed up disk writes if `InnoDB` log files are very small compared to the buffer pool size.
- In 4.0.13 we made `SHOW INNODB STATUS` to print detailed info on the latest `UNIQUE KEY` error, but storing that information could slow down `REPLACE` significantly. We no longer store or print the info.
- Fixed a bug: `SET FOREIGN_KEY_CHECKS=0` was not replicated properly in the MySQL replication. The fix will not be backported to 3.23.

- Fixed a bug: the parameter `innodb_max_dirty_pages_pct` forgot to take into account the free pages in the buffer pool. This could lead to excessive flushing even though there were lots of free pages in the buffer pool. Workaround: `SET GLOBAL innodb_max_dirty_pages_pct = 100`.
- Fixed a bug: If there were big index scans then a file read request could starve and InnoDB could assert because of a very long semaphore wait.
- Fixed a bug: If `AUTOCOMMIT=1` then inside `LOCK TABLES` MySQL failed to do the commit after an updating SQL statement if binary logging was not on, and for `SELECT` statements did not commit regardless of binary logging state.
- Fixed a bug: InnoDB could make the index page directory corrupt in the first B-tree page splits after a `mysqld` startup. A symptom would be an assertion in `page0page.c`, in function `page_dir_find_slot()`.
- Fixed a bug: If in a `FOREIGN KEY` with an `UPDATE CASCADE` clause the parent column was of a different internal storage length than the child column, then a cascaded update would make the column length wrong in the child table and corrupt the child table. Because of MySQL's 'silent column specification changes' a fixed-length `CHAR` column can change internally to a `VARCHAR` and cause this error.
- Fixed a bug: If a non-`latin1` character set was used and if in a `FOREIGN KEY` the parent column was of a different internal storage length than the child column, then all inserts to the child table would fail in a foreign key error.
- Fixed a bug: InnoDB could complain that it cannot find the clustered index record, or in rare cases return an extraneous row if a rollback, purge, and a `SELECT` coincided.
- Fixed a possible hang over the `btr0sea.c` latch if `SELECT` was used inside `LOCK TABLES`.
- Fixed a bug: contrary to what the release note of 4.0.13 said, the group commit still did not work if the MySQL binary logging was on.
- Fixed a bug: `os_event_wait()` did not work properly in Unix, which might have caused starvation in various log operations.
- Fixed a bug: If a single `DELETE` statement first managed to delete some rows and then failed in a `FOREIGN KEY` error or a `Table is full` error, MySQL did not roll back the whole SQL statement as it should, and also wrote the failed statement to the binary log, reporting there a non-zero `error_code`.
- Fixed a bug: the maximum allowed number of columns in a table is 1000, but InnoDB did not check that limit in `CREATE TABLE`, and a subsequent `INSERT` or `SELECT` from that table could cause an assertion.

## D.4.15. Changes in MySQL/InnoDB-3.23.57, June 20, 2003

- Changed the default value of `innodb_flush_log_at_trx_commit` from 0 to 1. If you have not specified it explicitly in your `my.cnf`, and your application runs much slower with this new release, it is because the value 1 causes a log flush to disk at each transaction commit.
- Fixed a bug: InnoDB forgot to call `pthread_mutex_destroy()` when a table was dropped. That could cause memory leakage on FreeBSD and other non-Linux Unixes.
- Fixed a bug: MySQL could erroneously return 'Empty set' if InnoDB estimated an index range size

to 0 records though the range was not empty; MySQL also failed to do the next-key locking in the case of an empty index range.

- Fixed a bug: `GROUP BY` and `DISTINCT` could treat `NULL` values unequal.

## D.4.16. Changes in MySQL/InnoDB-4.0.13, May 20, 2003

- `InnoDB` now supports `ALTER TABLE DROP FOREIGN KEY`. You have to use `SHOW CREATE TABLE` to find the internally generated foreign key ID when you want to drop a foreign key.
- `SHOW INNODB STATUS` now prints detailed information of the latest detected `FOREIGN KEY` and `UNIQUE KEY` errors. If you do not understand why `InnoDB` gives the error 150 from a `CREATE TABLE`, you can use this statement to study the reason.
- `ANALYZE TABLE` now works also for `InnoDB` type tables. It makes eight random dives to each of the index trees and updates index cardinality estimates accordingly. Note that because these are only estimates, repeated runs of `ANALYZE TABLE` may produce different numbers. MySQL uses index cardinality estimates only in join optimization. If some join is not optimized in the right way, you may try using `ANALYZE TABLE`.
- `InnoDB` group commit capability now works also when MySQL binary logging is switched on. There have to be > 2 client threads for the group commit to become active.
- Changed the default value of `innodb_flush_log_at_trx_commit` from 0 to 1. If you have not specified it explicitly in your `my.cnf`, and your application runs much slower with this new release, it is because the value 1 causes a log flush to disk at each transaction commit.
- Added a new global settable MySQL system variable `innodb_max_dirty_pages_pct`. It is an integer in the range 0 - 100. The default is 90. The main thread in `InnoDB` tries to flush pages from the buffer pool so that at most this many percents are not yet flushed at any time.
- If `innodb_force_recovery=6`, do not let `InnoDB` do repair of corrupt pages based on the doublewrite buffer.
- `InnoDB` startup now happens faster because it does not set the memory in the buffer pool to zero.
- Fixed a bug: The `InnoDB` parser for `FOREIGN KEY` definitions was confused by the keywords 'foreign key' inside MySQL comments.
- Fixed a bug: If you dropped a table to which there was a `FOREIGN KEY` reference, and later created the same table with non-matching data types, `InnoDB` could assert in `dict0load.c`, in function `dict_load_table()`.
- Fixed a bug: `GROUP BY` and `DISTINCT` could treat `NULL` values as not equal. MySQL also failed to do the next-key locking in the case of an empty index range.
- Fixed a bug: Do not commit the current transaction when a `MyISAM` table is updated; this also makes `CREATE TABLE` not to commit an `InnoDB` transaction, even when binary logging is enabled.
- Fixed a bug: We did not allow `ON DELETE SET NULL` to modify the same table where the delete was made; we can allow it because that cannot produce infinite loops in cascaded operations.
- Fixed a bug: Allow `HANDLER PREV` and `NEXT` also after positioning the cursor with a unique search on the primary key.
- Fixed a bug: If `MIN()` or `MAX()` resulted in a deadlock or a lock wait timeout, MySQL did not re-

turn an error, but returned `NULL` as the function value.

- Fixed a bug: `InnoDB` forgot to call `pthread_mutex_destroy()` when a table was dropped. That could cause memory leakage on FreeBSD and other non-Linux Unix systems.

### D.4.17. Changes in MySQL/InnoDB-4.1.0, April 3, 2003

- `InnoDB` now supports up to 64GB of buffer pool memory in a Windows 32-bit Intel computer. This is possible because `InnoDB` can use the AWE extension of Windows to address memory over the 4GB limit of a 32-bit process. A new startup variable `innodb_buffer_pool_awe_mem_mb` enables AWE and sets the size of the buffer pool in megabytes.
- Reduced the size of buffer headers and the lock table. `InnoDB` uses 2 % less memory.

### D.4.18. Changes in MySQL/InnoDB-3.23.56, March 17, 2003

- Fixed a major bug in `InnoDB` query optimization: queries of type `SELECT ... WHERE index-column < x` and `SELECT ... WHERE indexcolumn > x` could cause a table scan even if the selectivity would have been very good.
- Fixed a potential bug if MySQL calls `store_lock` with `TL_IGNORE` in the middle of a query.

### D.4.19. Changes in MySQL/InnoDB-4.0.12, March 18, 2003

- In crash recovery `InnoDB` now prints the progress in percents of a transaction rollback.
- Fixed a bug/feature: If your application program used `mysql_use_result()`, and used  $\geq 2$  connections to send SQL queries, it could deadlock on the adaptive hash S-latch in `btr0sea.c`. Now `mysqld` releases the S-latch whenever it passes data from a `SELECT` to the client.
- Fixed a bug: MySQL could erroneously return 'Empty set' if `InnoDB` estimated an index range size to 0 records though the range was not empty; MySQL also failed to do the next-key locking in the case of an empty index range.

### D.4.20. Changes in MySQL/InnoDB-4.0.11, February 25, 2003

- Fixed a bug introduced in 4.0.10: `SELECT ... FROM ... ORDER BY ... DESC` could hang in an infinite loop.
- An outstanding bug: `SET FOREIGN_KEY_CHECKS=0` is not replicated properly in the MySQL replication.

### D.4.21. Changes in MySQL/InnoDB-4.0.10, February 4, 2003

- In `INSERT INTO t1 SELECT ... FROM t2 WHERE ...` MySQL previously set a table level read lock on `t2`. This lock is now removed.

- Increased SHOW INNODB STATUS maximum printed length to 200KB.
- Fixed a major bug in InnoDB query optimization: queries of type SELECT ... WHERE index-column < x and SELECT ... WHERE indexcolumn > x could cause a table scan even if the selectivity would have been very good.
- Fixed a bug: purge could cause a hang in a BLOB table where the primary key index tree was of height 1. Symptom: semaphore waits caused by an X-latch set in btr\_free\_externally\_stored\_field().
- Fixed a bug: using InnoDB HANDLER commands on a fresh handle crashed mysqld in ha\_innobase::change\_active\_index().
- Fixed a bug: If MySQL estimated a query in the middle of a SELECT statement, InnoDB could hang on the adaptive hash index latch in btr0sea.c.
- Fixed a bug: InnoDB could report table corruption and assert in page\_dir\_find\_owner\_slot() if an adaptive hash index search coincided with purge or an insert.
- Fixed a bug: some filesystem snapshot tool in Windows 2000 could cause an InnoDB file write to fail with error 33 ERROR\_LOCK\_VIOLATION. In synchronous writes InnoDB now retries the write 100 times at 1 second intervals.
- Fixed a bug: REPLACE INTO t1 SELECT ... did not work if t1 has an auto-inc column.
- An outstanding bug: SET FOREIGN\_KEY\_CHECKS=0 is not replicated properly in the MySQL replication.

## D.4.22. Changes in MySQL/InnoDB-3.23.55, January 24, 2003

- In INSERT INTO t1 SELECT ... FROM t2 WHERE ... MySQL previously set a table level read lock on t2. This lock is now removed.
- Fixed a bug: If the combined size of InnoDB log files was >= 2GB in a 32-bit computer, InnoDB would write log in a wrong position. That could make crash recovery and InnoDB Hot Backup to fail in log scan.
- Fixed a bug: index cursor restoration could theoretically fail.
- Fixed a bug: an assertion in btr0sea.c, in function btr\_search\_info\_update\_slow could theoretically fail in a race of 3 threads.
- Fixed a bug: purge could cause a hang in a BLOB table where the primary key index tree was of height 1. Symptom: semaphore waits caused by an X-latch set in btr\_free\_externally\_stored\_field().
- Fixed a bug: If MySQL estimated a query in the middle of a SELECT statement, InnoDB could hang on the adaptive hash index latch in btr0sea.c.
- Fixed a bug: InnoDB could report table corruption and assert in page\_dir\_find\_owner\_slot() if an adaptive hash index search coincided with purge or an insert.
- Fixed a bug: some filesystem snapshot tool in Windows 2000 could cause an InnoDB file write to fail with error 33 ERROR\_LOCK\_VIOLATION. In synchronous writes InnoDB now retries the write 100 times at 1 second intervals.
- An outstanding bug: SET FOREIGN\_KEY\_CHECKS=0 is not replicated properly in the MySQL replication. The fix appears in 4.0.11 and probably will not be backported to 3.23.



- Fixed bug in `InnoDB page0cur.c` file in function `page_cur_search_with_match` which caused `InnoDB` to remain on the same page forever. This bug is evident only in tables with more than one page.

### D.4.23. Changes in MySQL/InnoDB-4.0.9, January 14, 2003

- Removed the warning message: 'InnoDB: Out of memory in additional memory pool.'
- Fixed a bug: If the combined size of `InnoDB` log files was  $\geq 2$ GB in a 32-bit computer, `InnoDB` would write log in a wrong position. That could make crash recovery and `InnoDB Hot Backup` to fail.
- Fixed a bug: index cursor restoration could theoretically fail.

### D.4.24. Changes in MySQL/InnoDB-4.0.8, January 7, 2003

- `InnoDB` now supports also `FOREIGN KEY (...) REFERENCES ...(...)` [`ON UPDATE CASCADE` | `ON UPDATE SET NULL` | `ON UPDATE RESTRICT` | `ON UPDATE NO ACTION`].
- Tables and indexes now reserve 4 % less space in the tablespace. Also existing tables reserve less space. By upgrading to 4.0.8 you should see more free space in "InnoDB free" in `SHOW TABLE STATUS`.
- Fixed bugs: updating the `PRIMARY KEY` of a row would generate a foreign key error on all `FOREIGN KEY`s which referenced secondary keys of the row to be updated. Also, if a referencing `FOREIGN KEY` constraint only referenced the first columns in an index, and there were more columns in that index, updating the additional columns generated a foreign key error.
- Fixed a bug: If an index contains some column twice, and that column is updated, the table becomes corrupt. From now on `InnoDB` prevents creation of such indexes.
- Fixed a bug: removed superfluous error 149 and 150 printouts from the `.err` log when a locking `SELECT` caused a deadlock or a lock wait timeout.
- Fixed a bug: an assertion in `btr0sea.c`, in function `btr_search_info_update_slow` could theoretically fail in a race of 3 threads.
- Fixed a bug: one could not switch a session transaction isolation level back to `REPEATABLE READ` after setting it to something else.

### D.4.25. Changes in MySQL/InnoDB-4.0.7, December 26, 2002

- `InnoDB` in 4.0.7 is essentially the same as in 4.0.6.

### D.4.26. Changes in MySQL/InnoDB-4.0.6, December 19, 2002

- Since `innodb_log_arch_dir` has no relevance under MySQL, there is no need to specify it any more in `my.cnf`.

- `LOAD DATA INFILE` in `AUTOCOMMIT=1` mode no longer does implicit commits for each 1MB of written binary log.
- Fixed a bug introduced in 4.0.4: `LOCK TABLES ... READ LOCAL` should not set row locks on the rows read. This caused deadlocks and lock wait timeouts in `mysqldump`.
- Fixed two bugs introduced in 4.0.4: in `AUTO_INCREMENT`, `REPLACE` could cause the counter to be left 1 too low. A deadlock or a lock wait timeout could cause the same problem.
- Fixed a bug: `TRUNCATE` on a `TEMPORARY` table crashed `InnoDB`.
- Fixed a bug introduced in 4.0.5: If binary logging was not switched on, `INSERT INTO ... SELECT ...` or `CREATE TABLE ... SELECT ...` could cause `InnoDB` to hang on a semaphore created in `btr0sea.c`, line 128. Workaround: switch binary logging on.
- Fixed a bug: in replication issuing `SLAVE STOP` in the middle of a multiple-statement transaction could cause that `SLAVE START` would only perform a part of the transaction. A similar error could occur if the slave crashed and was restarted.

## D.4.27. Changes in MySQL/InnoDB-3.23.54, December 12, 2002

- Fixed a bug: the `InnoDB` range estimator greatly exaggerated the size of a short index range if the paths to the endpoints of the range in the index tree happened to branch in the root. This could cause unnecessary table scans in SQL queries.
- Fixed a bug: `ORDER BY` could fail if you had not created a primary key to a table, but had defined several indexes of which at least one was a `UNIQUE` index with all its columns declared as `NOT NULL`.
- Fixed a bug: a lock wait timeout in connection with `ON DELETE CASCADE` could cause corruption in indexes.
- Fixed a bug: If a `SELECT` was done with a unique key from a primary index, and the search matched to a delete-marked record, `InnoDB` could erroneously return the `NEXT` record.
- Fixed a bug introduced in 3.23.53: `LOCK TABLES ... READ LOCAL` should not set row locks on the rows read. This caused deadlocks and lock wait timeouts in `mysqldump`.
- Fixed a bug: If an index contains some column twice, and that column is updated, the table becomes corrupt. From now on `InnoDB` prevents creation of such indexes.

## D.4.28. Changes in MySQL/InnoDB-4.0.5, November 18, 2002

- `InnoDB` now supports also transaction isolation levels `READ COMMITTED` and `READ UNCOMMITTED`. `READ COMMITTED` more closely emulates Oracle and makes porting of applications from Oracle to MySQL easier.
- Deadlock resolution is now selective: we try to pick as victims transactions with less modified or inserted rows.
- `FOREIGN KEY` definitions are now aware of the `lower_case_table_names` setting in `my.cnf`.

- SHOW CREATE TABLE does not output the database name to a FOREIGN KEY definition if the referred table is in the same database as the table.
- InnoDB does a consistency check to most index pages before writing them to a data file.
- If you set `innodb_force_recovery > 0`, InnoDB tries to jump over corrupt index records and pages when doing SELECT \* FROM table. This helps in dumping.
- InnoDB now again uses asynchronous unbuffered I/O in Windows 2000 and XP; only unbuffered simulated async I/O in NT, 95/98/ME.
- Fixed a bug: the InnoDB range estimator greatly exaggerated the size of a short index range if the paths to the endpoints of the range in the index tree happened to branch in the root. This could cause unnecessary table scans in SQL queries. The fix is also backported to 3.23.54.
- Fixed a bug present in 3.23.52, 4.0.3, 4.0.4: InnoDB startup could take very long or even crash on some Windows 95/98/ME computers.
- Fixed a bug: the AUTO-INC lock was held to the end of the transaction if it was granted after a lock wait. This could cause unnecessary deadlocks.
- Fixed a bug: If SHOW INNODB STATUS, innodb\_monitor, or innodb\_lock\_monitor had to print several hundred transactions in one report, and the output became truncated, InnoDB would hang, printing to the error log many waits for a mutex created at `srv0srv.c`, line 1621.
- Fixed a bug: SHOW INNODB STATUS on Unix always reported average file read size as 0 bytes.
- Fixed a potential bug in 4.0.4: InnoDB now does ORDER BY ... DESC like MyISAM.
- Fixed a bug: DROP TABLE could cause crash or a hang if there was a rollback concurrently running on the table. The fix will be backported to 3.23 only if this appears a real problem for users.
- Fixed a bug: ORDER BY could fail if you had not created a primary key to a table, but had defined several indexes of which at least one was a UNIQUE index with all its columns declared as NOT NULL.
- Fixed a bug: a lock wait timeout in connection with ON DELETE CASCADE could cause corruption in indexes.
- Fixed a bug: If a SELECT was done with a unique key from a primary index, and the search matched to a delete-marked record, InnoDB could return the NEXT record.
- Outstanding bugs: in 4.0.4 two bugs were introduced to AUTO\_INCREMENT. REPLACE can cause the counter to be left 1 too low. A deadlock or a lock wait timeout can cause the same problem. These are fixed in 4.0.6.

## D.4.29. Changes in MySQL/InnoDB-3.23.53, October 9, 2002

- We again use unbuffered disk I/O to data files in Windows. Windows XP and Windows 2000 read performance seems to be very poor with normal I/O.
- Tuned range estimator so that index range scans are preferred over full index scans.
- Allow dropping and creating a table even if `innodb_force_recovery` is set. One can use this to drop a table which would cause a crash in rollback or purge, or if a failed table import causes a runaway rollback in recovery.

- Fixed a bug present in 3.23.52, 4.0.3, 4.0.4: **InnoDB** startup could take very long or even crash on some Windows 95/98/ME computers.
- Fixed a bug: fast shutdown (which is the default) sometimes was slowed down by purge and insert buffer merge.
- Fixed a bug: doing a big SELECT from a table where no rows were visible in a consistent read could cause a very long (> 600 seconds) semaphore wait in btr0cur.c line 310.
- Fixed a bug: the AUTO-INC lock was held to the end of the transaction if it was granted after a lock wait. This could cause unnecessary deadlocks.
- Fixed a bug: If you created a temporary table inside LOCK TABLES, and used that temporary table, that caused an assertion failure in ha\_innobase.cc.
- Fixed a bug: If SHOW INNODB STATUS, innodb\_monitor, or innodb\_lock\_monitor had to print several hundred transactions in one report, and the output became truncated, **InnoDB** would hang, printing to the error log many waits for a mutex created at srv0srv.c, line 1621.
- Fixed a bug: SHOW INNODB STATUS on Unix always reported average file read size as 0 bytes.

### D.4.30. Changes in MySQL/InnoDB-4.0.4, October 2, 2002

- We again use unbuffered disk I/O in Windows. Windows XP and Windows 2000 read performance seems to be very poor with normal I/O.
- Increased the maximum key length of **InnoDB** tables from 500 to 1024 bytes.
- Increased the table comment field in SHOW TABLE STATUS so that up to 16000 characters of foreign key definitions can be printed there.
- The auto-increment counter is no longer incremented if an insert of a row immediately fails in an error.
- Allow dropping and creating a table even if innodb\_force\_recovery is set. One can use this to drop a table which would cause a crash in rollback or purge, or if a failed table import causes a runaway rollback in recovery.
- Fixed a bug: Using ORDER BY primarykey DESC in 4.0.3 causes an assertion failure in btr0pcur.c, line 203.
- Fixed a bug: fast shutdown (which is the default) sometimes was slowed down by purge and insert buffer merge.
- Fixed a bug: doing a big SELECT from a table where no rows were visible in a consistent read could cause a very long (> 600 seconds) semaphore wait in btr0cur.c line 310.
- Fixed a bug: If the MySQL query cache was used, it did not get invalidated by a modification done by ON DELETE CASCADE or ...SET NULL.
- Fixed a bug: If you created a temporary table inside LOCK TABLES, and used that temporary table, that caused an assertion failure in ha\_innobd.cc.
- Fixed a bug: If you set innodb\_flush\_log\_at\_trx\_commit to 1, SHOW VARIABLES would show its value as 16 million.

### D.4.31. Changes in MySQL/InnoDB-4.0.3, August 28, 2002

- Removed unnecessary deadlocks when inserts have to wait for a locking read, update, or delete to release its next-key lock.
- The MySQL `HANDLER` SQL commands now work also for `InnoDB` type tables. `InnoDB` does the `HANDLER` reads always as consistent reads. `HANDLER` is a direct access path to read individual indexes of tables. In some cases, `HANDLER` can be used as a substitute of server-side cursors.
- Fixed a bug in 4.0.2: even a simple insert could crash the AIX version.
- Fixed a bug: If you used in a table name characters whose code is > 127, in `DROP TABLE InnoDB` could assert on line 155 of `pars0sym.c`.
- Compilation from source now provides a working version both on HP-UX-11 and HP-UX-10.20. The source of 4.0.2 worked only on 11, and the source of 3.23.52 only on 10.20.
- Fixed a bug: If compiled on 64-bit Solaris, `InnoDB` produced a bus error at startup.

### D.4.32. Changes in MySQL/InnoDB-3.23.52, August 16, 2002

- The feature set of 3.23 is frozen from this version on. New features go the 4.0 branch, and only bug-fixes are made to the 3.23 branch.
- Many CPU-bound join queries now run faster. On Windows also many other CPU-bound queries run faster.
- A new SQL command `SHOW INNODB STATUS` returns the output of the `InnoDB` Monitor to the client. The `InnoDB` Monitor now prints detailed information on the latest detected deadlock.
- `InnoDB` made the SQL query optimizer to avoid too much index-only range scans and choose full table scans instead. This is now fixed.
- `BEGIN` and `COMMIT` are now added in the binary log around transactions. The MySQL replication now respects transaction borders: a user no longer sees half transactions in replication slaves.
- A replication slave now prints in crash recovery the last master binary log position it was able to recover to.
- A new setting `innodb_flush_log_at_trx_commit=2` makes `InnoDB` to write the log to the operating system file cache at each commit. This is almost as fast as the setting `innodb_flush_log_at_trx_commit=0`, and the setting 2 also has the nice feature that in a crash where the operating system does not crash, no committed transaction is lost. If the operating system crashes or there is a power outage, then the setting 2 is no safer than the setting 0.
- Added checksum fields to log blocks.
- `SET FOREIGN_KEY_CHECKS=0` helps in importing tables in an arbitrary order which does not respect the foreign key rules.
- `SET UNIQUE_CHECKS=0` speeds up table imports into `InnoDB` if you have `UNIQUE` constraints on secondary indexes. This flag should be used only if you are certain that the input records contain no `UNIQUE` constraint violations.
- `SHOW TABLE STATUS` now lists also possible `ON DELETE CASCADE` or `ON DELETE SET NULL` in the comment field of the table.

- When CHECK TABLE is run on any [InnoDB](#) type table, it now checks also the adaptive hash index for all tables.
- If you defined ON DELETE CASCADE or SET NULL and updated the referenced key in the parent row, [InnoDB](#) deleted or updated the child row. This is now changed to conform to standard SQL: you get the error 'Cannot delete parent row'.
- Improved the auto-increment algorithm: now the first insert or SHOW TABLE STATUS initializes the auto-increment counter for the table. This removes almost all surprising deadlocks caused by SHOW TABLE STATUS.
- Aligned some buffers used in reading and writing to data files. This allows using unbuffered raw devices as data files in Linux.
- Fixed a bug: If you updated the primary key of a table so that only the case of characters changed, that could cause assertion failures, mostly in page0page.ic line 515.
- Fixed a bug: If you delete or update a row referenced in a foreign key constraint and the foreign key check has to wait for a lock, then the check may report an erroneous result. This affects also the ON DELETE... operation.
- Fixed a bug: A deadlock or a lock wait timeout error in [InnoDB](#) causes [InnoDB](#) to roll back the whole transaction, but MySQL could still write the earlier SQL statements to the binary log, even though [InnoDB](#) rolled them back. This could, for example, cause replicated databases to get out-of-sync.
- Fixed a bug: If the database happened to crash in the middle of a commit, then the recovery might leak tablespace pages.
- Fixed a bug: If you specified a non-latin1 character set in my.cnf, then, in contrary to what is stated in the manual, in a foreign key constraint a string type column had to have the same length specification in the referencing table and the referenced table.
- Fixed a bug: DROP TABLE or DROP DATABASE could fail if there simultaneously was a CREATE TABLE running.
- Fixed a bug: If you configured the buffer pool bigger than 2GB in a 32-bit computer, [InnoDB](#) would assert in buf0buf.ic line 214.
- Fixed a bug: on 64-bit computers updating rows which contained the SQL NULL in some column could cause the undo log and the ordinary log to become corrupt.
- Fixed a bug: innodb\_log\_monitor caused a hang if it suppressed lock prints for a page.
- Fixed a bug: in the HP-UX-10.20 version mutexes would leak and cause race conditions and crashes in any part of [InnoDB](#) code.
- Fixed a bug: If you ran in the AUTOCOMMIT mode, executed a SELECT, and immediately after that a RENAME TABLE, then RENAME would fail and MySQL would complain about error 1192.
- Fixed a bug: If compiled on 64-bit Solaris, [InnoDB](#) produced a bus error at startup.

### D.4.33. Changes in MySQL/InnoDB-4.0.2, July 10, 2002

- [InnoDB](#) is essentially the same as InnoDB-3.23.51.
- If no innodb\_data\_file\_path is specified, [InnoDB](#) at the database creation now creates a 10MB auto-

extending data file `ibdata1` to the `datadir` of MySQL. In 4.0.1 the file was 64MB and not auto-extending.

## D.4.34. Changes in MySQL/InnoDB-3.23.51, June 12, 2002

- Fixed a bug: a join could result in a segmentation fault in copying of a BLOB or TEXT column if some of the BLOB or TEXT columns in the table contained SQL NULL values.
- Fixed a bug: If you added self-referential foreign key constraints with ON DELETE CASCADE to tables and a row deletion caused InnoDB to attempt the deletion of the same row twice because of a cascading delete, then you got an assertion failure.
- Fixed a bug: If you use MySQL 'user-level locks' and close a connection, then InnoDB may assert in `ha_innbase.cc`, line 302.

## D.4.35. Changes in MySQL/InnoDB-3.23.50, April 23, 2002

- InnoDB now supports an auto-extending last data file. You do not need to preallocate the whole data file at the database startup.
- Made several changes to facilitate the use of the InnoDB Hot Backup tool. It is a separate non-free tool you can use to take online backups of your database without shutting down the server or setting any locks.
- If you want to run the InnoDB Hot Backup tool on an auto-extending data file you have to upgrade it to version `ibbackup-0.35`.
- The log scan phase in crash recovery now runs much faster.
- Starting from this server version, the hot backup tool truncates unused ends in the backup InnoDB data files.
- To allow the hot backup tool to work, on Windows we no longer use unbuffered I/O or native async I/O; instead we use the same simulated async I/O as on Unix.
- You can now define the ON DELETE CASCADE or ON DELETE SET NULL clause on foreign keys.
- FOREIGN KEY constraints now survive ALTER TABLE and CREATE INDEX.
- We suppress the FOREIGN KEY check if any of the column values in the foreign key or referenced key to be checked is the SQL NULL. This is compatible with Oracle, for example.
- SHOW CREATE TABLE now lists also foreign key constraints. Also mysqldump no longer forgets about foreign keys in table definitions.
- You can now add a new foreign key constraint with ALTER TABLE ... ADD CONSTRAINT FOREIGN KEY (...) REFERENCES ... (...).
- FOREIGN KEY definitions now allow backticks around table and column names.
- MySQL command SET TRANSACTION ISOLATION LEVEL ... has now the following effect on InnoDB tables: If a transaction is defined as SERIALIZABLE then InnoDB conceptually adds LOCK IN SHARE MODE to all consistent reads. If a transaction is defined to have any other isola-

tion level, then [InnoDB](#) obeys its default locking strategy which is REPEATABLE READ.

- `SHOW TABLE STATUS` no longer sets an x-lock at the end of an auto-increment index if the auto-increment counter has been initialized. This removes in almost all cases the surprising deadlocks caused by `SHOW TABLE STATUS`.
- Fixed a bug: in a `CREATE TABLE` statement the string 'foreign' followed by a non-space character confused the FOREIGN KEY parser and caused table creation to fail with errno 150.

## D.4.36. Changes in MySQL/InnoDB-3.23.49, February 17, 2002

- Fixed a bug: If you called `DROP DATABASE` for a database on which there simultaneously were running queries, the MySQL server could crash or hang. Crashes fixed, but a full fix has to wait some changes in the MySQL layer of code.
- Fixed a bug: on Windows one had to put the database name in lowercase for `DROP DATABASE` to work. Fixed in 3.23.49: case no longer matters on Windows. On Unix, the database name remains case sensitive.
- Fixed a bug: If one defined a non-latin1 character set as the default character set, then definition of foreign key constraints could fail in an assertion failure in `dict0crea.c`, reporting an internal error 17.

## D.4.37. Changes in MySQL/InnoDB-3.23.48, February 9, 2002

- Tuned the SQL optimizer to favor more often index searches over table scans.
- Fixed a performance problem when several large `SELECT` queries are run concurrently on a multi-processor Linux computer. Large CPU-bound `SELECT` queries now also generally run faster on all platforms.
- If MySQL binary logging is used, [InnoDB](#) now prints after crash recovery the latest MySQL binary log file name and the position in that file (= byte offset) [InnoDB](#) was able to recover to. This is useful, for example, when resynchronizing a master and a slave database in replication.
- Added better error messages to help in installation problems.
- One can now recover also MySQL temporary tables which have become orphaned inside the [InnoDB](#) tablespace.
- [InnoDB](#) now prevents a FOREIGN KEY declaration where the signedness is not the same in the referencing and referenced integer columns.
- Fixed a bug: calling `SHOW CREATE TABLE` or `SHOW TABLE STATUS` could cause memory corruption and make `mysqld` to crash. Especially at risk was `mysqldump`, because it calls frequently `SHOW CREATE TABLE`.
- Fixed a bug: If on Unix you did an `ALTER TABLE` to an [InnoDB](#) table and simultaneously did queries to it, `mysqld` could crash with an assertion failure in `row0row.c`, line 474.
- Fixed a bug: If inserts to several tables containing an auto-inc column were wrapped inside one `LOCK TABLES`, [InnoDB](#) asserted in `lock0lock.c`.
- In 3.23.47 we allowed several NULLS in a UNIQUE secondary index. But `CHECK TABLE` was not relaxed: it reports the table as corrupt. `CHECK TABLE` no longer complains in this situation.



- Fixed a bug: on Sparc and other high-endian processors SHOW VARIABLES showed innodb\_flush\_log\_at\_trx\_commit and other boolean-valued startup parameters always OFF even if they were switched on.
- Fixed a bug: If you ran mysqld-max-nt as a service on Windows NT/2000, the service shutdown did not always wait long enough for the InnoDB shutdown to finish.

### D.4.38. Changes in MySQL/InnoDB-3.23.47, December 28, 2001

- Recovery happens now faster, especially in a lightly loaded system, because background checkpointing has been made more frequent.
- InnoDB allows now several similar key values in a UNIQUE secondary index if those values contain SQL NULLs. Thus the convention is now the same as in MyISAM tables.
- InnoDB gives a better row count estimate for a table which contains BLOBs.
- In a FOREIGN KEY constraint InnoDB is now case-insensitive to column names, and in Windows also to table names.
- InnoDB allows a FOREIGN KEY column of CHAR type to refer to a column of VARCHAR type, and vice versa. MySQL silently changes the type of some columns between CHAR and VARCHAR, and these silent changes do not hinder FOREIGN KEY declaration any more.
- Recovery has been made more resilient to corruption of log files.
- Unnecessary statistics calculation has been removed from queries which generate a temporary table. Some ORDER BY and DISTINCT queries now run much faster.
- MySQL now knows that the table scan of an InnoDB table is done through the primary key. This saves a sort in some ORDER BY queries.
- The maximum key length of InnoDB tables is again restricted to 500 bytes. The MySQL interpreter is not able to handle longer keys.
- The default value of innodb\_lock\_wait\_timeout was changed from infinite to 50 seconds, the default value of innodb\_file\_io\_threads from 9 to 4.

### D.4.39. Changes in MySQL/InnoDB-4.0.1, December 23, 2001

- InnoDB is the same as in 3.23.47.
- In 4.0.0 the MySQL interpreter did not know the syntax LOCK IN SHARE MODE. This has been fixed.
- In 4.0.0 multiple-table delete did not work for transactional tables. This has been fixed.

### D.4.40. Changes in MySQL/InnoDB-3.23.46, November 30, 2001

- This is the same as 3.23.45.

## D.4.41. Changes in MySQL/InnoDB-3.23.45, November 23, 2001

- This is a bugfix release.
- In versions 3.23.42-.44 when creating a table on Windows, you have to use lowercase letters in the database name to be able to access the table. Fixed in 3.23.45.
- [InnoDB](#) now flushes stdout and stderr every 10 seconds: If these are redirected to files, the file contents can be better viewed with an editor.
- Fixed an assertion failure in .44, in `trx0trx.c`, line 178 when you drop a table which has the `.frm` file but does not exist inside [InnoDB](#).
- Fixed a bug in the insert buffer. The insert buffer tree could get into an inconsistent state, causing a crash, and also crashing the recovery. This bug could appear especially in large table imports or alterations.
- Fixed a bug in recovery: [InnoDB](#) could go into an infinite loop constantly printing a warning message that it cannot find free blocks from the buffer pool.
- Fixed a bug: when you created a temporary table of the [InnoDB](#) type, and then used `ALTER TABLE` to it, the MySQL server could crash.
- Prevented creation of MySQL system tables 'mysql.user', 'mysql.host', or 'mysql.db', in the [InnoDB](#) type.
- Fixed a bug which can cause an assertion failure in 3.23.44 in `srv0srv.c`, line 1728.

## D.4.42. Changes in MySQL/InnoDB-3.23.44, November 2, 2001

- You can define foreign key constraints on [InnoDB](#) tables. An example: `FOREIGN KEY (col1) REFERENCES table2(col2)`.
- You can create data files larger than 4GB in those filesystems that allow it.
- Improved [InnoDB](#) monitors, including a new `innodb_table_monitor` which allows you to print the contents of the [InnoDB](#) internal data dictionary.
- `DROP DATABASE` now works also for [InnoDB](#) tables.
- Accent characters in the default character set `latin1` are ordered according to the MySQL ordering.  
  
NOTE: If you are using `latin1` and have inserted characters whose code is  $> 127$  to an indexed `CHAR` column, you should run `CHECK TABLE` on your table when you upgrade to 3.23.43, and drop and reimport the table if `CHECK TABLE` reports an error!
- [InnoDB](#) calculates better table cardinality estimates.
- Change in deadlock resolution: in .43 a deadlock rolls back only the SQL statement, in .44 it rolls back the whole transaction.

- Deadlock, lock wait timeout, and foreign key constraint violations (no parent row, child rows exist) now return native MySQL error codes 1213, 1205, 1216, 1217, respectively.
- A new my.cnf parameter `innodb_thread_concurrency` helps in performance tuning in high concurrency environments.
- A new my.cnf option `innodb_force_recovery` helps you in dumping tables from a corrupted database.
- A new my.cnf option `innodb_fast_shutdown` speeds up shutdown. Normally `InnoDB` does a full purge and an insert buffer merge at shutdown.
- Raised maximum key length to 7000 bytes from a previous limit of 500 bytes.
- Fixed a bug in replication of auto-inc columns with multiline inserts.
- Fixed a bug when the case of letters changes in an update of an indexed secondary column.
- Fixed a hang when there are more than 24 data files.
- Fixed a crash when `MAX(col)` is selected from an empty table, and `col` is a not the first column in a multi-column index.
- Fixed a bug in purge which could cause crashes.

#### **D.4.43. Changes in MySQL/InnoDB-3.23.43, October 4, 2001**

- This is essentially the same as InnoDB-3.23.42.

#### **D.4.44. Changes in MySQL/InnoDB-3.23.42, September 9, 2001**

- Fixed a bug which corrupted the table if the primary key of a > 8000-byte row was updated.
- There are now 3 types of `InnoDB` Monitors: `innodb_monitor`, `innodb_lock_monitor`, and `innodb_tablespace_monitor`. `innodb_monitor` now prints also buffer pool hit rate and the total number of rows inserted, updated, deleted, read.
- Fixed a bug in `RENAME TABLE`.
- Fixed a bug in replication with an auto-increment column.

#### **D.4.45. Changes in MySQL/InnoDB-3.23.41, August 13, 2001**

- Support for < 4GB rows. The previous limit was 8000 bytes.
- Use the doublewrite file flush method.
- Raw disk partitions supported as data files.
- `InnoDB` Monitor.

- Several hang bugs fixed and an `ORDER BY` bug (“Sort aborted”) fixed.

## D.4.46. Changes in MySQL/InnoDB-3.23.40, July 16, 2001

- Only a few rare bugs fixed.

## D.4.47. Changes in MySQL/InnoDB-3.23.39, June 13, 2001

- `CHECK TABLE` now works for InnoDB tables.
- A new `my.cnf` parameter `innodb_unix_file_flush_method` introduced. It can be used to tune disk write performance.
- An auto-increment column now gets new values past the transaction mechanism. This saves CPU time and eliminates transaction deadlocks in new value assignment.
- Several bugfixes, most notably the rollback bug in 3.23.38.

## D.4.48. Changes in MySQL/InnoDB-3.23.38, May 12, 2001

- The new syntax `SELECT ... LOCK IN SHARE MODE` is introduced.
- InnoDB now calls `fsync()` after every disk write and calculates a checksum for every database page it writes or reads, which reveals disk defects.
- Several bugfixes.

## D.5. Changes in MySQL Cluster

Starting from 4.1.13 and 5.0.7, all Cluster changes are included in the MySQL Change History, and this manual section is no longer separately maintained.

### D.5.1. Changes in MySQL Cluster-5.0.7 (10 June 2005)

**Note:** Starting with version 5.0.8, changes for MySQL Cluster can be found in the combined MySQL Change History.

Functionality added or changed:

Bugs fixed:

- ([Bug#11019](http://bugs.mysql.com/11019) [http://bugs.mysql.com/11019]) `mgmapi start backup` in some cases returns wrong `backupid`
- ([Bug#10190](http://bugs.mysql.com/10190) [http://bugs.mysql.com/10190]) Backup from cluster with `NoOfReplica=1` is corrupt
- ([Bug#9246](http://bugs.mysql.com/9246) [http://bugs.mysql.com/9246]) Condition pushdown and left join, wrong result
- ([Bug#10956](http://bugs.mysql.com/10956) [http://bugs.mysql.com/10956]) More than 7 node restarts with `--initial` caused

cluster to fail.

- ([Bug#9945](http://bugs.mysql.com/9945) [<http://bugs.mysql.com/9945>]) `ALTER TABLE` caused server crash. (Linux/390)
- ([Bug#9826](http://bugs.mysql.com/9826) [<http://bugs.mysql.com/9826>]) ([Bug#10948](http://bugs.mysql.com/10948) [<http://bugs.mysql.com/10948>]) Schema change (`DROP TABLE`, `ALTER TABLE`) crashed HPUX and PPC32.
- ([Bug#10711](http://bugs.mysql.com/10711) [<http://bugs.mysql.com/10711>]) ([Bug#9363](http://bugs.mysql.com/9363) [<http://bugs.mysql.com/9363>]) ([Bug#8918](http://bugs.mysql.com/8918) [<http://bugs.mysql.com/8918>]) ([Bug#10058](http://bugs.mysql.com/10058) [<http://bugs.mysql.com/10058>]) ([Bug#9025](http://bugs.mysql.com/9025) [<http://bugs.mysql.com/9025>]) Cluster would time out and crash after first query; setting DataMemory to more than 2GB prevented cluster from starting; calling `ndb_select_count()` crashed the cluster. (64-bit Unix OSes)

## D.5.2. Changes in MySQL Cluster-5.0.6 (26 May 2005)

Functionality added or changed:

- Limit on number of metadata objects (number of tables, indexes and BLOBs) now increased to 20,320

Bugs fixed:

- The server would hang on successive calls to an `INSERT ... ON DUPLICATE KEY UPDATE` query. ([Bug#9725](http://bugs.mysql.com/9725) [<http://bugs.mysql.com/9725>])
- ([Bug#10193](http://bugs.mysql.com/10193) [<http://bugs.mysql.com/10193>]) Invalid DataDir in config causes ndbd segmentation fault
- ([Bug#10813](http://bugs.mysql.com/10813) [<http://bugs.mysql.com/10813>]) Build with SCI Transporter fails
- ([Bug#10831](http://bugs.mysql.com/10831) [<http://bugs.mysql.com/10831>]) ndb mgmd LogDestination maxfiles does not rotate logs properly

## D.5.3. Changes in MySQL Cluster-5.0.5 (Not released)

Functionality added or changed:

- Decreased IndexMemory Usage
- Parallel key lookup (read-multi-range) for queries like `SELECT * FROM t1 WHERE primary_key IN (1,2,3,4,5,6,7,8,9,10);`

Bugs fixed:

Patches merged from versions 4.1.11 and 4.1.12

- ([Bug#8315](http://bugs.mysql.com/8315) [<http://bugs.mysql.com/8315>]) NdbScanFilter cmp method only works for strings of exact word boundary length
- ([Bug#8103](http://bugs.mysql.com/8103) [<http://bugs.mysql.com/8103>]) Configuration handling error

- ([Bug#8035](http://bugs.mysql.com/8035) [http://bugs.mysql.com/8035]) mysqld signal 10 when ndbd is shutdown
- ([Bug#7631](http://bugs.mysql.com/7631) [http://bugs.mysql.com/7631]) NDB\$EVENT contains unreadable event and table names
- ([Bug#7628](http://bugs.mysql.com/7628) [http://bugs.mysql.com/7628]) Filtered event types are ignored
- ([Bug#7627](http://bugs.mysql.com/7627) [http://bugs.mysql.com/7627]) Drop Event operation fails
- ([Bug#7424](http://bugs.mysql.com/7424) [http://bugs.mysql.com/7424]) create index on datetime fails

## D.5.4. Changes in MySQL Cluster-5.0.4 (16 April 2005)

Functionality added or changed:

- Condition pushdown to storage engine now works for update and delete as well

Bugs fixed:

- ([Bug#9675](http://bugs.mysql.com/9675) [http://bugs.mysql.com/9675]) Auto-increment not working with INSERT..SELECT and NDB storage
- ([Bug#9517](http://bugs.mysql.com/9517) [http://bugs.mysql.com/9517]) Condition pushdown to storage engine does not work for update/delete
- ([Bug#9282](http://bugs.mysql.com/9282) [http://bugs.mysql.com/9282]) API Node Crashes/Reloads on 'DELETE FROM'
- ([Bug#9280](http://bugs.mysql.com/9280) [http://bugs.mysql.com/9280]) Memory leak in cluster when dependent sub-queries are used
- ([Bug#8585](http://bugs.mysql.com/8585) [http://bugs.mysql.com/8585]) ndb\_cache2 fails on aix52

## D.5.5. Changes in MySQL Cluster-5.0.3 (23 March 2005: Beta)

Functionality added or changed:

- Condition pushdown to storage engine
- Query cache enabled for cluster

Bugs fixed:

- Patches merged from version 4.1.10

## D.5.6. Changes in MySQL Cluster-5.0.1 (27 July 2004)

Functionality added or changed:

- This was the first MySQL Cluster release in the 5.0 series. As nearly all attention was still focused

on getting 4.1 stable, it is not recommended to use MySQL 5.0.1 for MySQL Cluster.

Bugs fixed:

- N/A

## D.5.7. Changes in MySQL Cluster-4.1.13 (15 July 2005)

Functionality added or changed:

Bugs fixed:

- ([Bug#11132](http://bugs.mysql.com/11132) [<http://bugs.mysql.com/11132>]) Connections between data nodes and management nodes were not being closed following shutdown of `ndb_mgmd`.
- ([Bug#11050](http://bugs.mysql.com/11050) [<http://bugs.mysql.com/11050>]) `ndb_mgm> show` printed incorrectly after master data node failure.
- ([Bug#10956](http://bugs.mysql.com/10956) [<http://bugs.mysql.com/10956>]) More than 7 node restarts with `--initial` caused cluster to fail.
- ([Bug#9826](http://bugs.mysql.com/9826) [<http://bugs.mysql.com/9826>]) ([Bug#10948](http://bugs.mysql.com/10948) [<http://bugs.mysql.com/10948>]) Schema change (`DROP TABLE`, `ALTER TABLE`) crashed HP-UX and PPC32.
- ([Bug#9025](http://bugs.mysql.com/9025) [<http://bugs.mysql.com/9025>]) Data nodes failed to restart on 64-bit Solaris.
- ([Bug#11166](http://bugs.mysql.com/11166) [<http://bugs.mysql.com/11166>]) Insert records were incorrectly applied by `ndb_restore`, thus making restoration from backup inconsistent if the binlog contained inserts.
- ([Bug#8918](http://bugs.mysql.com/8918) [<http://bugs.mysql.com/8918>]) ([Bug#9363](http://bugs.mysql.com/9363) [<http://bugs.mysql.com/9363>]) ([Bug#10711](http://bugs.mysql.com/10711) [<http://bugs.mysql.com/10711>]) ([Bug#10058](http://bugs.mysql.com/10058) [<http://bugs.mysql.com/10058>]) ([Bug#9025](http://bugs.mysql.com/9025) [<http://bugs.mysql.com/9025>]) Cluster would time out and crash after first query; setting `DataMemory` to more than 2GB prevented cluster from starting; calling `ndb_select_count()` crashed the cluster. (64-bit Unix OSes)
- ([Bug#10190](http://bugs.mysql.com/10190) [<http://bugs.mysql.com/10190>]) When making a backup of a cluster where `NumberOfReplicas` was equal to 1, the backup's metadata was corrupted. (Linux)
- ([Bug#9945](http://bugs.mysql.com/9945) [<http://bugs.mysql.com/9945>]) `ALTER TABLE` caused server crash. (Linux/390)
- ([Bug#11133](http://bugs.mysql.com/11133) [<http://bugs.mysql.com/11133>]) A delete operation performed as part of a transaction caused an erroneous result.
- ([Bug#10294](http://bugs.mysql.com/10294) [<http://bugs.mysql.com/10294>]) Not allowing sufficient parallelism in cluster configuration (for example, `NoOfTransactions` too small) caused `ndb_restore` to fail without generating any error messages.
- ([Bug#11290](http://bugs.mysql.com/11290) [<http://bugs.mysql.com/11290>]) Setting `TransactionInactiveTimeout=0` did not result in an infinite timeout.

## D.5.8. Changes in MySQL Cluster-4.1.12 (13 May 2005)

Functionality added or changed:

Bugs fixed:

- ([Bug#10471](http://bugs.mysql.com/10471) [http://bugs.mysql.com/10471]) Backup can become inconsistent with certain combinations of multiple-row updates
- ([Bug#10287](http://bugs.mysql.com/10287) [http://bugs.mysql.com/10287]) `ndb_select_all "delimiter"` option non functional
- ([Bug#10142](http://bugs.mysql.com/10142) [http://bugs.mysql.com/10142]) Unhandled resource shortage in UNIQUE index code
- ([Bug#10029](http://bugs.mysql.com/10029) [http://bugs.mysql.com/10029]) crash in ordered index scan after db full
- ([Bug#10001](http://bugs.mysql.com/10001) [http://bugs.mysql.com/10001]) 2 NDB nodes get signal 6 (abort) in DBTC
- ([Bug#9969](http://bugs.mysql.com/9969) [http://bugs.mysql.com/9969]) 4012 - has misleading error message
- ([Bug#9960](http://bugs.mysql.com/9960) [http://bugs.mysql.com/9960]) START BACKUP reports failure albeit succeeding
- ([Bug#9924](http://bugs.mysql.com/9924) [http://bugs.mysql.com/9924]) ABORT BACKUP 1 crashes 4 node cluster
- ([Bug#9892](http://bugs.mysql.com/9892) [http://bugs.mysql.com/9892]) Index activation file during node recovery
- ([Bug#9891](http://bugs.mysql.com/9891) [http://bugs.mysql.com/9891]) Crash in DBACC (line 7004) during commit
- ([Bug#9865](http://bugs.mysql.com/9865) [http://bugs.mysql.com/9865]) SELECT does not function properly
- ([Bug#9839](http://bugs.mysql.com/9839) [http://bugs.mysql.com/9839]) Column with AUTOINC contains -1 Value on node stop
- ([Bug#9757](http://bugs.mysql.com/9757) [http://bugs.mysql.com/9757]) Uncompleted node failure after gracefully stopping node
- ([Bug#9749](http://bugs.mysql.com/9749) [http://bugs.mysql.com/9749]) Transactions causes deadlock in ACC
- ([Bug#9724](http://bugs.mysql.com/9724) [http://bugs.mysql.com/9724]) Node fails to start: Message: File has already been opened
- ([Bug#9691](http://bugs.mysql.com/9691) [http://bugs.mysql.com/9691]) UPDATE fails on attempt to update primary key
- ([Bug#9675](http://bugs.mysql.com/9675) [http://bugs.mysql.com/9675]) Auto-increment not working with INSERT..SELECT and NDB storage
- ([Bug#9318](http://bugs.mysql.com/9318) [http://bugs.mysql.com/9318]) drop database does not drop ndb tables
- ([Bug#9280](http://bugs.mysql.com/9280) [http://bugs.mysql.com/9280]) Memory leak in cluster when dependent sub-queries are used
- ([Bug#8928](http://bugs.mysql.com/8928) [http://bugs.mysql.com/8928]) create table with keys will shutdown the cluster
- Creating a table did not work for a cluster with 6 nodes. ([Bug#8928](http://bugs.mysql.com/8928) [http://bugs.mysql.com/8928]) Databases with 1, 2, 4, 8, ... ( $2^n$  nodes) did not have the problem. After a rolling upgrade, restart each node manually by restarting it with the `--initial` option. Otherwise, use dump and restore after an upgrade.

## D.5.9. Changes in MySQL Cluster-4.1.11 (01 April 2005)

Functionality added or changed:

Bugs fixed:



- ([Bug#9916](http://bugs.mysql.com/9916) [http://bugs.mysql.com/9916]) DbaccMain.cpp / DBACC (Line: 4876) / Pointer too large
- ([Bug#9435](http://bugs.mysql.com/9435) [http://bugs.mysql.com/9435]) TIMESTAMP columns don't update
- ([Bug#9052](http://bugs.mysql.com/9052) [http://bugs.mysql.com/9052]) Uninitialized data during unique index build, potential cluster crash
- ([Bug#8876](http://bugs.mysql.com/8876) [http://bugs.mysql.com/8876]) Timeout when committing aborted transaction after node failure
- ([Bug#8786](http://bugs.mysql.com/8786) [http://bugs.mysql.com/8786]) ndb\_autodiscover, drop index can fail, wait 2 minutes timeout
- ([Bug#8853](http://bugs.mysql.com/8853) [http://bugs.mysql.com/8853]) Transaction aborted after long time during node failure (4012)
- ([Bug#8753](http://bugs.mysql.com/8753) [http://bugs.mysql.com/8753]) Invalid schema object version after dropping index (crash fixed, currently retry required)
- ([Bug#8645](http://bugs.mysql.com/8645) [http://bugs.mysql.com/8645]) Assertion failure with multiple management servers
- ([Bug#8557](http://bugs.mysql.com/8557) [http://bugs.mysql.com/8557]) ndbd does not get same nodeid on restart
- ([Bug#8556](http://bugs.mysql.com/8556) [http://bugs.mysql.com/8556]) corrupt ndb\_mgm show printout for certain configurations
- ([Bug#8167](http://bugs.mysql.com/8167) [http://bugs.mysql.com/8167]) cluster shared memory and mysqld signal usage clash

## D.5.10. Changes in MySQL Cluster-4.1.10 (12 February 2005)

Bugs fixed:

- ([Bug#8284](http://bugs.mysql.com/8284) [http://bugs.mysql.com/8284]) Out of fragment memory in DBACC
- ([Bug#8262](http://bugs.mysql.com/8262) [http://bugs.mysql.com/8262]) Node crash due to bug in DBLQH
- ([Bug#8208](http://bugs.mysql.com/8208) [http://bugs.mysql.com/8208]) node restart fails on Aix 5.2
- ([Bug#8167](http://bugs.mysql.com/8167) [http://bugs.mysql.com/8167]) cluster shared memory and mysqld signal usage clash
- ([Bug#8101](http://bugs.mysql.com/8101) [http://bugs.mysql.com/8101]) unique index and error 4209 while selecting
- ([Bug#8070](http://bugs.mysql.com/8070) [http://bugs.mysql.com/8070]) ([Bug#7937](http://bugs.mysql.com/7937) [http://bugs.mysql.com/7937]) ([Bug#6716](http://bugs.mysql.com/6716) [http://bugs.mysql.com/6716]) various ndb\_restore core dumps on HP-UX
- ([Bug#8010](http://bugs.mysql.com/8010) [http://bugs.mysql.com/8010]) 4006 forces MySQL Node Restart
- ([Bug#7928](http://bugs.mysql.com/7928) [http://bugs.mysql.com/7928]) out of connection objects
- ([Bug#7898](http://bugs.mysql.com/7898) [http://bugs.mysql.com/7898]) mysqld crash with ndb (solaris)
- ([Bug#7864](http://bugs.mysql.com/7864) [http://bugs.mysql.com/7864]) Not possible to have more than 4.5G data memory

## D.5.11. Changes in MySQL Cluster-4.1.9 (13 January 2005)

Functionality added or changed:

- New implementation of shared memory transporter.
- Cluster automatically configures shared memory transporter if possible.
- Cluster prioritizes usage of transporters with shared memory and localhost TCP
- Added switches to control the above functions, `ndb-shm` and `ndb-optimized-node-selection`.

Bugs fixed:

- ([Bug#7805](http://bugs.mysql.com/7805) [http://bugs.mysql.com/7805]) config.ini parsing error
- ([Bug#7798](http://bugs.mysql.com/7798) [http://bugs.mysql.com/7798]) Running range scan after alter table in different thread causes node failure
- ([Bug#7761](http://bugs.mysql.com/7761) [http://bugs.mysql.com/7761]) Alter table does not autocommit
- ([Bug#7725](http://bugs.mysql.com/7725) [http://bugs.mysql.com/7725]) Indexed DATETIME Columns Return Random Results
- ([Bug#7660](http://bugs.mysql.com/7660) [http://bugs.mysql.com/7660]) START BACKUP does not increment BACKUP-ID (Big Endian machines)
- ([Bug#7593](http://bugs.mysql.com/7593) [http://bugs.mysql.com/7593]) Cannot Create A Large NDB Data Warehouse
- ([Bug#7480](http://bugs.mysql.com/7480) [http://bugs.mysql.com/7480]) Mysqld crash in ha\_ndbcluster using Query Browser
- ([Bug#7470](http://bugs.mysql.com/7470) [http://bugs.mysql.com/7470]) shared memory transporter does not connect
- ([Bug#7396](http://bugs.mysql.com/7396) [http://bugs.mysql.com/7396]) Primary Key not working in NDB Mysql Clustered table (solaris)
- ([Bug#7379](http://bugs.mysql.com/7379) [http://bugs.mysql.com/7379]) ndb restore fails to handle blobs and multiple databases
- ([Bug#7346](http://bugs.mysql.com/7346) [http://bugs.mysql.com/7346]) ndb\_restore enters infinite loop
- ([Bug#7340](http://bugs.mysql.com/7340) [http://bugs.mysql.com/7340]) Problem for inserting data into the Text field on utf8
- ([Bug#7124](http://bugs.mysql.com/7124) [http://bugs.mysql.com/7124]) ndb\_mgmd is aborted on startup when using SHM connection

## D.5.12. Changes in MySQL Cluster-4.1.8 (14 December 2004)

Functionality added or changed:

- Default port for `ndb_mgmd` was changed to 1186 (from 2200) as this port number was officially assigned to MySQL Cluster by IANA.
- New command in `ndb_mgm`, PURGE STALE SESSIONS, as a workaround for cases where nodes fail to allocate a node id even if it is free to use.
- New command in `ndb_mgm`, CONNECT.

- The `ndb` executables have been changed to make use of the regular MySQL command line option parsing features. See [Section 15.6.5, “Command Options for MySQL Cluster Processes”](#), for notes on changes.
- As bonus of the above you can now specify all command line options in `my.cnf` using the executable names as sections, that is, `[ndbd]`, `[ndb_mgmd]`, `[ndb_mgm]`, `[ndb_restore]`, and so forth.

```
[ndbd]
ndb-connectstring=myhost.domain.com:1234
[ndb_mgm]
ndb-connectstring=myhost.domain.com:1234
```

- Added use of section `[mysql_cluster]` in `my.cnf`. All cluster executables, including `mysqld`, parse this section. For example, this is a convenient place to put `ndb-connectstring` so that it need be specified only once.
- Added cluster log info events on allocation and deallocation of nodeid's.
- Added cluster log info events on connection refuse as a result of version mismatch.
- Extended connectstring syntax to allow for leaving the port number out. For example, `ndb-connectstring|connect-string=myhost1,myhost2,myhost3` is a valid connectstring and connect occurs on default port 1186.
- Clear text `ndb` error messages provided also for error codes that are mapped to corresponding `mysql` error codes, by executing `SHOW WARNINGS` after an error has occurred which relates to the `ndb` storage engine.
- Significant performance improvements done for read performance, especially for blobs.
- Added some variables for performance tuning, `ndb_force_send` and `ndb_use_exact_count`. Do `show variables like 'ndb%';` in `mysql` client for listing. Use `set` command to alter variables.
- Added variables to set some options, `ndb_use_transactions` and `ndb_autoincrement_prefetch_sz`.

Bugs fixed:

- ([Bug#7303](http://bugs.mysql.com/7303) [<http://bugs.mysql.com/7303>]) `ndb_mgm`: Trying to set `CLUSTERLOG` for a specific node id core dumps
- ([Bug#7193](http://bugs.mysql.com/7193) [<http://bugs.mysql.com/7193>]) `start backup` gives false error printout
- ([Bug#7153](http://bugs.mysql.com/7153) [<http://bugs.mysql.com/7153>]) Cluster nodes don't report error on endianness mismatch
- ([Bug#7152](http://bugs.mysql.com/7152) [<http://bugs.mysql.com/7152>]) `ndb_mgmd` segmentation fault on incorrect `HostName` in configuration
- ([Bug#7104](http://bugs.mysql.com/7104) [<http://bugs.mysql.com/7104>]) `clusterlog` filtering and level setting broken
- ([Bug#6995](http://bugs.mysql.com/6995) [<http://bugs.mysql.com/6995>]) `ndb_recover` on `varchar` fields results in changing case of data
- ([Bug#6919](http://bugs.mysql.com/6919) [<http://bugs.mysql.com/6919>]) all status only shows 2 nodes on a 8-node cluster
- ([Bug#6871](http://bugs.mysql.com/6871) [<http://bugs.mysql.com/6871>]) `DBD` execute failed: Got error 897 'Unknown error code'

from ndbcluster

- ([Bug#6794](http://bugs.mysql.com/6794) [http://bugs.mysql.com/6794]) Wrong outcome of update operation of ndb table
- ([Bug#6791](http://bugs.mysql.com/6791) [http://bugs.mysql.com/6791]) Segmentation fault when config.ini is not correctly set
- ([Bug#6775](http://bugs.mysql.com/6775) [http://bugs.mysql.com/6775]) failure in acc when running many mysql clients
- ([Bug#6696](http://bugs.mysql.com/6696) [http://bugs.mysql.com/6696]) ndb\_mgm command line options inconsistent with behavior
- ([Bug#6684](http://bugs.mysql.com/6684) [http://bugs.mysql.com/6684]) ndb\_restore doesn't give error messages if improper command given
- ([Bug#6677](http://bugs.mysql.com/6677) [http://bugs.mysql.com/6677]) ndb\_mgm can crash on "ALL CLUSTERLOG"
- ([Bug#6538](http://bugs.mysql.com/6538) [http://bugs.mysql.com/6538]) Error code returned when select max() on empty table with index
- ([Bug#6451](http://bugs.mysql.com/6451) [http://bugs.mysql.com/6451]) failing create table gives "ghost" tables which are impossible to remove
- ([Bug#6435](http://bugs.mysql.com/6435) [http://bugs.mysql.com/6435]) strange behavior of left join
- ([Bug#6426](http://bugs.mysql.com/6426) [http://bugs.mysql.com/6426]) update with long pk fails
- ([Bug#6398](http://bugs.mysql.com/6398) [http://bugs.mysql.com/6398]) update of primary key fails
- ([Bug#6354](http://bugs.mysql.com/6354) [http://bugs.mysql.com/6354]) mysql does not complain about --ndbcluster option when NDB is not compiled in
- ([Bug#6331](http://bugs.mysql.com/6331) [http://bugs.mysql.com/6331]) INSERT IGNORE .. SELECT breaks subsequent inserts
- ([Bug#6288](http://bugs.mysql.com/6288) [http://bugs.mysql.com/6288]) cluster nodes crash on data import
- ([Bug#6031](http://bugs.mysql.com/6031) [http://bugs.mysql.com/6031]) To drop database you have to execute DROP DATABASE command twice
- ([Bug#6020](http://bugs.mysql.com/6020) [http://bugs.mysql.com/6020]) LOCK TABLE + delete returns error 208
- ([Bug#6018](http://bugs.mysql.com/6018) [http://bugs.mysql.com/6018]) REPLACE does not work for BLOBs + NDB
- ([Bug#6016](http://bugs.mysql.com/6016) [http://bugs.mysql.com/6016]) Strange crash with blobs + different DATABASES
- ([Bug#5973](http://bugs.mysql.com/5973) [http://bugs.mysql.com/5973]) ndb table belonging to different database shows up in show tables
- ([Bug#5872](http://bugs.mysql.com/5872) [http://bugs.mysql.com/5872]) ALTER TABLE with blob from ndb table to myisam fails
- ([Bug#5844](http://bugs.mysql.com/5844) [http://bugs.mysql.com/5844]) Failing mysql-test-run leaves stray NDB processes behind
- ([Bug#5824](http://bugs.mysql.com/5824) [http://bugs.mysql.com/5824]) HELP text messed up in ndb\_mgm
- ([Bug#5786](http://bugs.mysql.com/5786) [http://bugs.mysql.com/5786]) Duplicate key error after restore
- ([Bug#5785](http://bugs.mysql.com/5785) [http://bugs.mysql.com/5785]) lock timeout during concurrent update
- ([Bug#5782](http://bugs.mysql.com/5782) [http://bugs.mysql.com/5782]) Unknown error when using LIMIT with ndb table

- ([Bug#5756](http://bugs.mysql.com/5756) [<http://bugs.mysql.com/5756>]) RESTART node from ndb\_mgm fails
- A few more not reported bugs fixed

### D.5.13. Changes in MySQL Cluster-4.1.7 (23 October 2004)

Functionality added or changed:

- Optimization 1: Improved performance on index scans. Measured 30% performance increase on query which do large amounts of index scans.
- Optimization 2: Improved performance on primary key lookups. Around double performance for autocommitted primary key lookups.
- Optimization 3: Improved performance when using blobs by avoiding usage of exclusive locks for blobs.

Bugs fixed:

- A few bugs fixed.

### D.5.14. Changes in MySQL Cluster-4.1.6 (10 October 2004)

Functionality added or changed:

- Limited character set support for storage engine NDBCLUSTER:

Char set	Collation
big5	big5_chinese_ci
	big5_bin
binary	binary
euckr	euckr_korean_ci
	euckr_bin
gb2312	gb2312_chinese_ci
	gb2312_bin
gbk	gbk_chinese_ci
	gbk_bin
latin1	latin1_swedish_ci
	latin1_bin
sjis	sjis_japanese_ci
	sjis_bin
tis620	tis620_bin
ucs2	ucs2_general_ci
	ucs2_bin
ujis	ujis_japanese_ci

	ujis_bin
utf8	utf8_general_ci
	utf8_bin

- The SCI Transporter has been brought up-to-date with all changes and now works and has been documented as well.
- Optimizations when several clients to a MySQL Server access ndb tables.
- Added more checks and warnings for erroneous and inappropriate cluster configurations.
- [SHOW TABLES](#) now directly shows ndb tables created on a different MySQL server, that is, without a prior table access.
- Enhanced support for starting MySQL Server independently of ndbd and ndb\_mgmd.
- Clear text ndb error messages provided by executing [SHOW WARNINGS](#) after an error has occurred which relates to the ndb storage engine.

Bugs fixed:

- Quite a few bugs fixed.

## D.5.15. Changes in MySQL Cluster-4.1.5 (16 September 2004)

Functionality added or changed:

- Many queries in MySQL Cluster are executed as range scans or full table scans. All queries that don't use a unique hash index or the primary hash index use this access method. In a distributed system it is crucial that batching is properly performed.

In previous versions, the batch size was fixed to 16 per data node. In this version it is configurable per MySQL Server. So for queries using lots of large scans it is appropriate to set this parameter rather large and for queries using many small scans only fetching a small amount of records it is appropriate to set it low.

The performance of queries can easily change as much as 40% based on how this variable is set.

In future versions more logic will be implemented for assessing the batch size on a per-query basis. Thus, the semantics of the new configuration variable [ScanBatchSize](#) are likely to change.

- The fixed size overhead of the ndbd process has been greatly decreased. This is also true for the overhead per operation record as well as overhead per table and index.

A number of new configuration variables have been introduced to enable configuration of system buffers. Configuration variables for specifying the numbers of tables, unique hash indexes, and ordered indexes have also been introduced.

New configuration variables: [MaxNoOfOrderedIndexes](#), [MaxNoOfUniqueHashIndexes](#)

Configuration variables no longer used: [MaxNoOfIndexes](#) (split into the two above).

- In previous versions `ALTER TABLE`, `TRUNCATE TABLE`, and `LOAD DATA` were performed as one big transaction. In this version, all of these statements are automatically separated into several distinct transactions.  
  
This removes the limitation that one could not change very large tables due to the `MaxNoOfConcurrentOperations` parameter.
- MySQL Cluster's online backup feature now backs up indexes so that both data and indexes are restored.
- In previous versions it was not possible to use `NULL` in indexes. This is now possible for all supported index types.
- Much work has been put onto making `AUTO_INCREMENT` features work as for other table handlers. Autoincrements as a partial key is still only supported by `MyISAM`.
- In earlier versions, `mysqld` would crash if the cluster wasn't started with the `--ndbcluster` option. Now `mysqld` handles cluster crashes and starts without crashing.
- The `-i` option for initial startup of `ndbd` has been removed. Initial startup still can be specified by using the `--initial` option. The reason for this is to ensure that it is clear what takes place when using `--initial`: this option completely removes all data from the disk and should only be used at initial start, in certain software upgrade cases, and in some cases as a workaround when nodes cannot be restarted successfully.
- The management client (`ndb_mgm`) now has additional commands and more information is printed for some commands such as `show`.
- In previous versions, the files were called `ndb_0..` when it wasn't possible to allocate a node ID when starting the node. To ensure that files are not so easily overwritten, these files are now named `ndb_pid..`, where `pid` is the process ID assigned by the OS.
- The default parameters have changed for `ndb_mgmd` and `ndbd`. In particular, they are now started as daemons by default. The `-n` option has been removed since it could cause confusion as to its meaning (nstart or nodaemon).
- In the configuration file, you can now use `[NDBD]` as an alias for `[DB]`, `[MYSQLD]` as an alias for `[API]`, and `[NDB_MGMD]` as an alias for `[MGM]`. **Note:** In fact, `[NDBD]`, `[MYSQLD]`, and `[NDB_MGMD]` are now the preferred designations, although the older ones will continue to be supported for some time to come in order to maintain backward compatibility.
- Many more checks for consistency in configuration have been introduced to in order to provide quicker feedback on configuration errors.
- In the connect string, it is now possible to use both `;` and `,` as the separator between entries. Thus, `"nodeid=2,host=localhost:2200"` is equivalent to `"nodeid=2;host=localhost:2200"`.  
  
In the configuration file, it is also possible to use `:` or `=` for assignment values. For example, `MaxNoOfOrderedIndexes : 128` and `MaxNoOfOrderedIndexes = 128` are equivalent expressions.
- The configuration variable names are now case insensitive, so `MaxNoOfOrderedIndexes : 128` is equivalent to `MAXNOOFORDEREDINDEXES = 128`.
- It is possible now to set the backup directory separately from the `FileSystemPath` by using the `BackupDir` configuration variable.

Log files and trace files can now be placed in any directory by setting the `DataDir` configuration variable.

`FileSystemPath` is no longer mandatory and defaults to `DataDir`.

- Queries involving tables from different databases are now supported.
- It is now possible to update the primary key.
- The performance of ordered indexes has been greatly improved, particularly the maintenance of indexes on updates, inserts and deletes.

Bugs fixed:

- Quite a few bugs fixed.

## D.5.16. Changes in MySQL Cluster-4.1.4 (31 August 2004)

Functionality added or changed:

- The names of the log files and trace files created by the `ndbd` and `ndb_mgmd` processes have changed.
- Support for the many `BLOB` data types was introduced in this version.

Bugs fixed:

- Quite a few bugs were fixed in the 4.1.4 release.

## D.5.17. Changes in MySQL Cluster-4.1.3 (28 June 2004)

Functionality added or changed:

- This was the first MySQL Cluster release so all functionality was new.

Bugs fixed:

- Various bugs fixed in the development process leading up to 4.1.3.

## D.6. MySQL Connector/ODBC (MyODBC) Change History

### D.6.1. Changes in MyODBC 3.51.13

Functionality added or changed:

- N/A



Bugs fixed:

- Using Connector/ODBC, with `SQLBindCol` and binding the length to the return value from `SQL_LEN_DATA_AT_EXEC` fails with a memory allocation error. ([Bug#20547](#) [<http://bugs.mysql.com/20547>])
- The `SQLDriverConnect()` ODBC method did not work with recent MyODBC releases. ([Bug#12393](#) [<http://bugs.mysql.com/12393>])

## D.6.2. Changes in MyODBC 3.51.12

Functionality added or changed:

- N/A

Bugs fixed:

- File DSNs could not be saved. ([Bug#12019](#) [<http://bugs.mysql.com/12019>])
- `SQLColumns()` returned no information for tables that had a column named using a reserved word. ([Bug#9539](#) [<http://bugs.mysql.com/9539>])

## D.6.3. Changes in MyODBC 3.51.11

Functionality added or changed: No changes.

Bugs fixed:

- `mysql_list_dbcolumns()` and `insert_fields()` were retrieving all rows from a table. Fixed the queries generated by these functions to return no rows. ([Bug#8198](#) [<http://bugs.mysql.com/8198>])
- `SQLGetTypeInfo()` returned `tinyblob` for `SQL_VARBINARY` and nothing for `SQL_BINARY`. Fixed to return `varbinary` for `SQL_VARBINARY`, `binary` for `SQL_BINARY`, and `longblob` for `SQL_LONGVARBINARY`. ([Bug#8138](#) [<http://bugs.mysql.com/8138>])

## D.7. MySQL Connector/NET Change History

### D.7.1. Version 5.0.1 (not yet released)

Functionality added or changed:

- You can now install the Connector/NET MSI package from the command line using the `/passive`, `/quiet`, `/q` options. ([Bug#19994](#) [<http://bugs.mysql.com/19994>])
- The `MySqlException` class is now derived from the `DbException` class. ([Bug#21874](#) [<http://bugs.mysql.com/21874>])

Bugs fixed:

- Executing multiple queries as part of a transaction returns `There is already an openDataReader associated with this Connection which must be closed first.` ([Bug#7248](http://bugs.mysql.com/7248) [<http://bugs.mysql.com/7248>])
- The `#` would not be accepted within column/table names, even though it was valid. ([Bug#21521](http://bugs.mysql.com/21521) [<http://bugs.mysql.com/21521>])

## D.7.2. Version 5.0.0

This is a new Alpha development release, fixing recently discovered bugs.

**NOTE:** This Alpha release, as any other pre-production release, should not be installed on *production* level systems or systems with critical data. It is good practice to back up your data before installing any new version of software. Although MySQL has worked very hard to ensure a high level of quality, protect your data by making a backup as you would for any software beta release. Please refer to our bug database at <http://bugs.mysql.com/> for more details about the individual bugs fixed in this version.

Bugs fixed:

- `CommandText`: Question mark in comment line is being parsed as a parameter. ([Bug#6214](http://bugs.mysql.com/6214) [<http://bugs.mysql.com/6214>])

Functionality added or changed:

- Implemented Usage Advisor.
- Added Async query methods.
- Reimplemented `PacketReader/PacketWriter` support into `MySQLStream` class.
- Added internal implementation of SHA1 so we don't have to distribute the OpenNetCF on mobile devices.
- Added usage advisor warnings for requesting column values by the wrong type.
- Reworked connection string classes to be simpler and faster.
- Added procedure metadata caching.
- Added perfmon hooks for stored procedure cache hits and misses.
- Implemented `MySQLConnectionBuilder` class.
- Implemented `MySQLClientFactory` class.
- Implemented classes and interfaces for ADO.Net 2.0 support.
- Replaced use of `ICSharpCode` with .NET 2.0 internal deflate support.
- Refactored test suite to test all protocols in a single pass.
- Completely refactored how column values are handled to avoid boxing in some cases.

### D.7.3. Version 1.0.8 (not yet released)

Bugs fixed:

- You can now install the Connector/NET MSI package from the command line using the `/passive`, `/quiet`, `/q` options. ([Bug#19994](http://bugs.mysql.com/19994) [<http://bugs.mysql.com/19994>])
- Executing multiple queries as part of a transaction returns `There is already an openDataReader associated with this Connection which must be closed first`. ([Bug#7248](http://bugs.mysql.com/7248) [<http://bugs.mysql.com/7248>])
- Calling `MySqlCommandBuilder.DeriveParameters` for a stored procedure that has no parameters would cause an application crash. ([Bug#15077](http://bugs.mysql.com/15077) [<http://bugs.mysql.com/15077>])
- A `SELECT` query on a table with a date with a value of `'0000-00-00'` would hang the application. ([Bug#17736](http://bugs.mysql.com/17736) [<http://bugs.mysql.com/17736>])
- The `#` would not be accepted within column/table names, even though it was valid. ([Bug#21521](http://bugs.mysql.com/21521) [<http://bugs.mysql.com/21521>])
- Calling `Close` on a connection after calling a stored procedure would trigger a `NullReferenceException`. ([Bug#20581](http://bugs.mysql.com/20581) [<http://bugs.mysql.com/20581>])
- `IDataRecord.GetString` would raise `NullPointerException` for null values in returned rows. Method now throws `SqlNullValueException`. ([Bug#19294](http://bugs.mysql.com/19294) [<http://bugs.mysql.com/19294>])
- An exception would be raised when using an output parameter to a `System.String` value. ([Bug#17814](http://bugs.mysql.com/17814) [<http://bugs.mysql.com/17814>])
- The `DiscoverParameters` function would fail when a stored procedure used a `NUMERIC` parameter type. ([Bug#19515](http://bugs.mysql.com/19515) [<http://bugs.mysql.com/19515>])
- When running a query that included a date comparison, a `DateReader` error would be raised. ([Bug#19481](http://bugs.mysql.com/19481) [<http://bugs.mysql.com/19481>])
- Parameter substitution in queries where the order of parameters and table fields did not match would substitute incorrect values. ([Bug#19261](http://bugs.mysql.com/19261) [<http://bugs.mysql.com/19261>])
- When working with multiple threads, character set initialization would generate errors. ([Bug#17106](http://bugs.mysql.com/17106) [<http://bugs.mysql.com/17106>])
- When using an unsigned 64-bit integer in a stored procedure, the unsigned bit would be lost stored. ([Bug#16934](http://bugs.mysql.com/16934) [<http://bugs.mysql.com/16934>])
- The connection string parser did not allow single or double quotes in the password. ([Bug#16659](http://bugs.mysql.com/16659) [<http://bugs.mysql.com/16659>])
- The `CommandBuilder` ignored `Unsigned` flag at `Parameter` creation. ([Bug#17375](http://bugs.mysql.com/17375) [<http://bugs.mysql.com/17375>])
- `CHAR` type added to `MySqlDbType`. ([Bug#17749](http://bugs.mysql.com/17749) [<http://bugs.mysql.com/17749>])
- Unsigned data types were not properly supported. ([Bug#16788](http://bugs.mysql.com/16788) [<http://bugs.mysql.com/16788>])

Functionality added or changed:

- Stored procedures are now cached.
- The method for retrieving stored procedure metadata has been changed so that users without `SELECT` privileges on the `mysql.proc` table can use a stored procedure.

## D.7.4. Version 1.0.7

- Unsigned `tinyint` (NET byte) would lead to and incorrectly determined parameter type from the parameter value. (Bug#18570 [<http://bugs.mysql.com/18570>])
- The parameter collection object's `Add()` method added parameters to the list without first checking to see whether they already existed. Now it updates the value of the existing parameter object if it exists. (Bug#13927 [<http://bugs.mysql.com/13927>])
- A `#42000Query was empty` exception occurred when executing a query built with `MySQLCommandBuilder`, if the query string ended with a semicolon. (Bug#14631 [<http://bugs.mysql.com/14631>])
- Implemented the `MySQLCommandBuilder.DeriveParameters` method that is used to discover the parameters for a stored procedure. (Bug#13632 [<http://bugs.mysql.com/13632>])
- Added support for the `cp932` character set. (Bug#13806 [<http://bugs.mysql.com/13806>])
- Calling a stored procedure where a parameter contained special characters (such as `'@'`) would produce an exception. Note that `ANSI_QUOTES` had to be enabled to make this possible. (Bug#13753 [<http://bugs.mysql.com/13753>])
- A statement that contained multiple references to the same parameter could not be prepared. (Bug#13541 [<http://bugs.mysql.com/13541>])
- The `Ping()` method did not update the `State` property of the `Connection` object. (Bug#13658 [<http://bugs.mysql.com/13658>])

## D.7.5. Version 1.0.6

- The `nant` build sequence had problems. (Bug#12978 [<http://bugs.mysql.com/12978>])
- Serializing a parameter failed if the first value passed in was `NULL`. (Bug#13276 [<http://bugs.mysql.com/13276>])
- Field names that contained the following characters caused errors: `( ) % < > /` (Bug#13036 [<http://bugs.mysql.com/13036>])
- The MySQL Connector/NET 1.0.5 installer would not install alongside MySQL Connector/NET 1.0.4. (Bug#12835 [<http://bugs.mysql.com/12835>])
- MySQL Connector/NET 1.0.5 could not connect on Mono. (Bug#13345 [<http://bugs.mysql.com/13345>])

## D.7.6. Version 1.0.5

- With multiple hosts in the connection string, MySQL Connector/NET would not connect to the last

- host in the list. ([Bug#12628](http://bugs.mysql.com/12628) [<http://bugs.mysql.com/12628>])
- MySQL Connector/NET interpreted the new decimal data type as a byte array. ([Bug#11294](http://bugs.mysql.com/11294) [<http://bugs.mysql.com/11294>])
  - The `cp1250` character set was not supported. ([Bug#11621](http://bugs.mysql.com/11621) [<http://bugs.mysql.com/11621>])
  - Connection could fail when .NET thread pool had no available worker threads. ([Bug#10637](http://bugs.mysql.com/10637) [<http://bugs.mysql.com/10637>])
  - Decimal parameters caused syntax errors. ([Bug#11550](http://bugs.mysql.com/11550) [<http://bugs.mysql.com/11550>], [Bug#10486](http://bugs.mysql.com/10486) [<http://bugs.mysql.com/10486>], [Bug#10152](http://bugs.mysql.com/10152) [<http://bugs.mysql.com/10152>])
  - A call to a stored procedure caused an exception if the stored procedure had no parameters. ([Bug#11542](http://bugs.mysql.com/11542) [<http://bugs.mysql.com/11542>])
  - Certain malformed queries would trigger a `Connection must be valid and open` error message. ([Bug#11490](http://bugs.mysql.com/11490) [<http://bugs.mysql.com/11490>])
  - The `MySqlCommandBuilder` class could not handle queries that referenced tables in a database other than the default database. ([Bug#8382](http://bugs.mysql.com/8382) [<http://bugs.mysql.com/8382>])
  - MySQL Connector/NET could not work properly with certain regional settings. (WL#8228)
  - Trying to use a stored procedure when `Connection.Database` was not populated generated an exception. ([Bug#11450](http://bugs.mysql.com/11450) [<http://bugs.mysql.com/11450>])
  - Trying to read a `TIMESTAMP` column generated an exception. ([Bug#7951](http://bugs.mysql.com/7951) [<http://bugs.mysql.com/7951>])
  - Parameters were not recognized when they were separated by linefeeds. ([Bug#9722](http://bugs.mysql.com/9722) [<http://bugs.mysql.com/9722>])
  - Calling `MySqlConnection.clone` when a connection string had not yet been set on the original connection would generate an error. ([Bug#10281](http://bugs.mysql.com/10281) [<http://bugs.mysql.com/10281>])
  - Added support to call a stored function from MySQL Connector/NET. ([Bug#10644](http://bugs.mysql.com/10644) [<http://bugs.mysql.com/10644>])
  - MySQL Connector/NET could not connect to MySQL 4.1.14. ([Bug#12771](http://bugs.mysql.com/12771) [<http://bugs.mysql.com/12771>])
  - The `ConnectionString` property could not be set when a `MySqlConnection` object was added with the designer. ([Bug#12551](http://bugs.mysql.com/12551) [<http://bugs.mysql.com/12551>], [Bug#8724](http://bugs.mysql.com/8724) [<http://bugs.mysql.com/8724>])

### D.7.7. Version 1.0.4 1-20-05

- [Bug#7243](http://bugs.mysql.com/7243) [<http://bugs.mysql.com/7243>] calling prepare causing exception [fixed]
- Fixed another small problem with prepared statements
- [Bug#7258](http://bugs.mysql.com/7258) [<http://bugs.mysql.com/7258>] MySqlCommand.Connection returns an IDbConnection [fixed]
- [Bug#7345](http://bugs.mysql.com/7345) [<http://bugs.mysql.com/7345>] MySqlAdapter.Fill method throws Error message : Non-negative number required [fixed]

- [Bug#7478](http://bugs.mysql.com/7478) [http://bugs.mysql.com/7478] Clone method bug in MySqlCommand [fixed]
- [Bug#7612](http://bugs.mysql.com/7612) [http://bugs.mysql.com/7612] MySqlDataReader.GetString(index) returns non-Null value when field is Null [fixed]
- [Bug#7755](http://bugs.mysql.com/7755) [http://bugs.mysql.com/7755] MySqlReader.GetInt32 throws exception if column is unsigned [fixed]
- [Bug#7704](http://bugs.mysql.com/7704) [http://bugs.mysql.com/7704] GetBytes is working no more [fixed]
- [Bug#7724](http://bugs.mysql.com/7724) [http://bugs.mysql.com/7724] Quote character \222 not quoted in EscapeString [fixed]
- Fixed problem that causes named pipes to not work with some blob functionality
- Fixed problem with shared memory connections
- [Bug#7436](http://bugs.mysql.com/7436) [http://bugs.mysql.com/7436] Problem with Multiple resultsets... [fixed]
- Added or filled out several more topics in the API reference documentation

## D.7.8. Version 1.0.3-gamma 12-10-04

- Made MySQL the default named pipe name
- Now SHOW COLLATION is used upon connection to retrieve the full list of charset ids
- Fixed Invalid character set index: 200 ([Bug#6547](http://bugs.mysql.com/6547) [http://bugs.mysql.com/6547])
- Installer now includes options to install into GAC and create Start Menu items
- [Bug#6863](http://bugs.mysql.com/6863) [http://bugs.mysql.com/6863] - Int64 Support in MySqlCommand Parameters [fixed]
- Connections now do not have to give a database on the connection string
- [Bug#6770](http://bugs.mysql.com/6770) [http://bugs.mysql.com/6770] - MySqlDataReader.GetChar(int i) throws IndexOutOfRangeException Exception [fixed]
- Fixed problem where multiple resultsets having different numbers of columns would cause a problem
- [Bug#6983](http://bugs.mysql.com/6983) [http://bugs.mysql.com/6983] Exception stack trace lost when re-throwing exceptions [fixed]
- Fixed major problem with detecting null values when using prepared statements
- [Bug#6902](http://bugs.mysql.com/6902) [http://bugs.mysql.com/6902] Errors in parsing stored procedure parameters [fixed]
- [Bug#6668](http://bugs.mysql.com/6668) [http://bugs.mysql.com/6668] Integer "out" parameter from stored procedure returned as string [fixed]
- [Bug#7032](http://bugs.mysql.com/7032) [http://bugs.mysql.com/7032] MySqlDateTime in Datatables sorting by Text, not Date. [fixed]
- [Bug#7133](http://bugs.mysql.com/7133) [http://bugs.mysql.com/7133] Invalid query string when using inout parameters [fixed]
- [Bug#6831](http://bugs.mysql.com/6831) [http://bugs.mysql.com/6831] Test suite fails with MySQL 4.0 because of case sensitivity of table names [fixed]

- [Bug#7132](http://bugs.mysql.com/7132) [http://bugs.mysql.com/7132] Inserting DateTime causes System.InvalidCastException to be thrown [fixed]
- [Bug#6879](http://bugs.mysql.com/6879) [http://bugs.mysql.com/6879] InvalidCast when using DATE\_ADD-function [fixed]
- [Bug#6634](http://bugs.mysql.com/6634) [http://bugs.mysql.com/6634] An Open Connection has been Closed by the Host System [fixed]
- Added ServerThread property to MySqlConnection to expose server thread id
- Added Ping method to MySqlConnection
- Changed the name of the test suite to MySql.Data.Tests.dll

### D.7.9. Version 1.0.2-gamma 04-11-15

- Fixed problem with MySqlBinary where string values could not be used to update extended text columns
- Fixed Installation directory ignored using custom installation ([Bug#6329](http://bugs.mysql.com/6329) [http://bugs.mysql.com/6329])
- Fixed problem where setting command text leaves the command in a prepared state
- Fixed double type handling in MySqlParameter(string parameterName, object value) ([Bug#6428](http://bugs.mysql.com/6428) [http://bugs.mysql.com/6428])
- Fixed Zero date "0000-00-00" is returned wrong when filling Dataset ([Bug#6429](http://bugs.mysql.com/6429) [http://bugs.mysql.com/6429])
- Fixed problem where calling stored procedures might cause an "Illegal mix of collations" problem.
- Added charset connection string option
- Fixed #HY000 Illegal mix of collations (latin1\_swedish\_ci,IMPLICIT) and (utf8\_general\_ ([Bug#6322](http://bugs.mysql.com/6322) [http://bugs.mysql.com/6322])
- Added the TableEditor CS and VB sample
- Fixed Charset-map for UCS-2 ([Bug#6541](http://bugs.mysql.com/6541) [http://bugs.mysql.com/6541])
- Updated the installer to include the new samples
- Fixed Long inserts take very long time (Bu #5453)
- Fixed Objects not being disposed ([Bug#6649](http://bugs.mysql.com/6649) [http://bugs.mysql.com/6649])
- Provider is now using character set specified by server as default

### D.7.10. Version 1.0.1-beta2 04-10-27

- Fixed [Bug#5602](http://bugs.mysql.com/5602) [http://bugs.mysql.com/5602] Possible bug in MySqlParameter(string, object) constructor
- Fixed [Bug#5458](http://bugs.mysql.com/5458) [http://bugs.mysql.com/5458] Calling GetChars on a longtext column throws an ex-

ception

- Fixed [Bug#5474](http://bugs.mysql.com/5474) [http://bugs.mysql.com/5474] cannot run a stored procedure populating mysqlcommand.parameters
- Fixed [Bug#5469](http://bugs.mysql.com/5469) [http://bugs.mysql.com/5469] Setting DbType throws NullReferenceException
- Fixed problem where connector was not issuing a CMD\_QUIT before closing the socket
- Fixed [Bug#5392](http://bugs.mysql.com/5392) [http://bugs.mysql.com/5392] MySqlCommand sees "?" as parameters in string literals
- Fixed problem with ConnectionInternal where a key might be added more than once
- CP1252 is now used for Latin1 only when the server is 4.1.2 and later
- Fixed [Bug#5388](http://bugs.mysql.com/5388) [http://bugs.mysql.com/5388] DataReader reports all rows as NULL if one row is NULL
- Virtualized driver subsystem so future releases could easily support client or embedded server support
- Field buffers being reused to decrease memory allocations and increase speed
- Fixed problem where using old syntax while using the interfaces caused problems
- Using PacketWriter instead of Packet for writing to streams
- Refactored compression code into CompressedStream to clean up NativeDriver
- Added test case for resetting the command text on a prepared command
- Fixed problem where MySqlParameterCollection.Add() would throw unclear exception when given a null value ([Bug#5621](http://bugs.mysql.com/5621) [http://bugs.mysql.com/5621])
- Fixed constructor initialize problems in MySqlCommand() ([Bug#5613](http://bugs.mysql.com/5613) [http://bugs.mysql.com/5613])
- Fixed Parsing the ';' char ([Bug#5876](http://bugs.mysql.com/5876) [http://bugs.mysql.com/5876])
- Fixed missing Reference in DbType setter ([Bug#5897](http://bugs.mysql.com/5897) [http://bugs.mysql.com/5897])
- Fixed System.OverflowException when using YEAR datatype ([Bug#6036](http://bugs.mysql.com/6036) [http://bugs.mysql.com/6036])
- Added Aggregate function test (wasn't really a bug)
- Fixed serializing of floating point parameters (double, numeric, single, decimal) ([Bug#5900](http://bugs.mysql.com/5900) [http://bugs.mysql.com/5900])
- IsNullable error ([Bug#5796](http://bugs.mysql.com/5796) [http://bugs.mysql.com/5796])
- Fixed problem where connection lifetime on the connect string was not being respected
- Fixed problem where Min Pool Size was not being respected
- Fixed MySqlDataReader and 'show tables from ...' behavior ([Bug#5256](http://bugs.mysql.com/5256) [http://bugs.mysql.com/5256])
- Implemented SequentialAccess



- Fixed MySqlDateTime sets IsZero property on all subseq.records after first zero found ([Bug#6006](http://bugs.mysql.com/6006) [<http://bugs.mysql.com/6006>])
- Fixed Can't display Chinese correctly ([Bug#5288](http://bugs.mysql.com/5288) [<http://bugs.mysql.com/5288>])
- Fixed Russian character support as well
- Fixed Method TokenizeSql() uses only a limited set of valid characters for parameters ([Bug#6217](http://bugs.mysql.com/6217) [<http://bugs.mysql.com/6217>])
- Fixed NET Connector source missing resx files ([Bug#6216](http://bugs.mysql.com/6216) [<http://bugs.mysql.com/6216>])
- Fixed DBNull Values causing problems with retrieving/updating queries. ([Bug#5798](http://bugs.mysql.com/5798) [<http://bugs.mysql.com/5798>])
- Fixed Yet Another "object reference not set to an instance of an object" ([Bug#5496](http://bugs.mysql.com/5496) [<http://bugs.mysql.com/5496>])
- Fixed problem in PacketReader where it could try to allocate the wrong buffer size in EnsureCapacity
- Fixed GetBoolean returns wrong values ([Bug#6227](http://bugs.mysql.com/6227) [<http://bugs.mysql.com/6227>])
- Fixed IndexOutOfBounds when reading BLOB with DataReader with GetString(index) ([Bug#6230](http://bugs.mysql.com/6230) [<http://bugs.mysql.com/6230>])

### D.7.11. Version 1.0.0 04-09-01

- Fixed BUG# 3889 Thai encoding not correctly supported
- Updated many of the test cases
- Fixed problem with using compression
- Bumped version number to 1.0.0 for beta 1 release
- Added COPYING.rtf file for use in installer
- Removed all of the XML comment warnings (I'll clean them up better later)
- Removed some last references to ByteFX

### D.7.12. Version 0.9.0 04-08-30

- Added test fixture for prepared statements
- All type classes now implement a SerializeBinary method for sending their data to a PacketWriter
- Added PacketWriter class that will enable future low-memory large object handling
- Fixed many small bugs in running prepared statements and stored procedures
- Changed command so that an exception will not be throw in executing a stored procedure with parameters in old syntax mode

- SingleRow behavior now working right even with limit
- GetBytes now only works on binary columns
- Logger now truncates long sql commands so blob columns don't blow out our log
- host and database now have a default value of "" unless otherwise set
- FIXED BUG# 5214 Connection Timeout seems to be ignored
- Added test case for bug# 5051: GetSchema not working correctly
- Fixed problem where GetSchema would return false for IsUnique when the column is key
- MySqlDataReader GetXXX methods now using the field level MySqlValue object and not performing conversions
- FIXED BUG# 5097: DataReader returning NULL for time column
- Added test case for LOAD DATA LOCAL INFILE
- Added replacetext custom nant task
- Added CommandBuilderTest fixture
- Added Last One Wins feature to CommandBuilder
- Fixed persist security info case problem
- Fixed GetBool so that 1, true, "true", and "yes" all count as trueWL# 2024 Make parameter mark configurable
- Added the "old syntax" connection string parameter to allow use of @ parameter marker
- Fixed [Bug#4658](http://bugs.mysql.com/4658) [http://bugs.mysql.com/4658] MySqlCommandBuilder
- Fixed [Bug#4864](http://bugs.mysql.com/4864) [http://bugs.mysql.com/4864] ByteFX.MySqlClient caches passwords if 'Persist Security Info' is false
- Updated license banner in all source files to include FLOSS exception
- Added new .Types namespace and implementations for most current MySql types
- Added MySqlField41 as a subclass of MySqlField
- Changed many classes to now use the new .Types types
- Changed type enum int to Int32, short to Int16, and bigint to Int64
- Added dummy types UInt16, UInt32, and UInt64 to allow an unsigned parameter to be made
- Connections are now reset when they are pulled from the connection pool
- Refactored auth code in driver so it can be used for both auth and reset
- Added UserReset test in PoolingTests.cs
- Connections are now reset using COM\_CHANGE\_USER when pulled from the pool
- Implemented SingleResultSet behavior

- Implemented support of unicode
- Added char set mappings for utf-8 and ucs-2
- fixed [Bug#4520](http://bugs.mysql.com/4520) [http://bugs.mysql.com/4520] time fields overflow using bytefx .net mysql driver
- Modified time test in data type test fixture to check for time spans where hours > 24
- Fixed [Bug#4505](http://bugs.mysql.com/4505) [http://bugs.mysql.com/4505] Wrong string with backslash escaping in Byte-Fx.Data.MySqlClient.MySqlParameter
- Added code to Parameter test case TestQuoting to test for backslashes
- Fixed [Bug#4486](http://bugs.mysql.com/4486) [http://bugs.mysql.com/4486] mysqlcommandbuilder fails with multi-word column names
- Fixed bug in TokenizeSql where underscore would terminate character capture in parameter name
- Added test case for spaces in column names
- Fixed bug# 4324 - MySqlDataReader.GetBytes don't works correctly
- Added GetBytes() test case to DataReader test fixture
- Now reading all server variables in InternalConnection.Configure into Hashtable
- Now using string[] for index map in CharSetMap
- Added CRInSQL test case for carriage returns in SQL
- setting maxPacketSize to default value in Driver.ctor
- Fixed [Bug#4442](http://bugs.mysql.com/4442) [http://bugs.mysql.com/4442] - Setting MySqlDbType on a parameter doesn't set generic type
- Removed obsolete data types Long and LongLong
- Fixed bug# 4071 - Overflow exception thrown when using "use pipe" on connection string
- Changed "use pipe" keyword to "pipe name" or just "pipe"
- Allow reading multiple resultsets from a single query
- Added flags attribute to ServerStatusFlags enum
- Changed name of ServerStatus enum to ServerStatusFlags
- Fixed [Bug#4386](http://bugs.mysql.com/4386) [http://bugs.mysql.com/4386] - Inserted data row doesn't update properly
- Fixed [Bug#4074](http://bugs.mysql.com/4074) [http://bugs.mysql.com/4074] - Error processing show create table
- Change Packet.ReadLenInteger to ReadPackedLong and added packet.ReadPackedInteger that always reads integers packed with 2,3,4
- Added syntax.cs test fixture to test various SQL syntax bugs
- Fixed bug# 4149 Improper handling of time values. Now time value of 00:00:00 is not treated as null.
- Moved all test suite files into TestSuite folder

- Fixed bug where null column would move the result packet pointer backward
- Added new nant build script
- Fixed [Bug#3917](http://bugs.mysql.com/3917) [http://bugs.mysql.com/3917] - clear tablename so it will be regen'ed properly during the next GenerateSchema.
- Fixed [Bug#3915](http://bugs.mysql.com/3915) [http://bugs.mysql.com/3915] - GetValues was always returning zero and was also always trying to copy all fields rather than respecting the size of the array passed in.
- Implemented shared memory access protocol
- Implemented prepared statements for MySQL 4.1
- Implemented stored procedures for MySQL 5.0
- Renamed MySqlInternalConnection to InternalConnection
- SQL is now parsed as chars, fixes problems with other languages
- Added logging and allow batch connection string options
- Fixed [Bug#3888](http://bugs.mysql.com/3888) [http://bugs.mysql.com/3888] - RowUpdating event not set when setting the DataAdapter property
- Fixed bug in char set mapping
- Implemented 4.1 authentication
- Improved open/auth code in driver
- Improved how connection bits are set during connection
- Database name is now passed to server during initial handshake
- Changed namespace for client to MySql.Data.MySqlClient
- Changed assembly name of client to MySql.Data.dll
- Changed license text in all source files to GPL
- Added the MySqlClient.build Nant file
- Removed the mono batch files
- Moved some of the unused files into notused folder so nant build file can use wildcards
- Implemented shared memory accesss
- Major revamp in code structure
- Prepared statements now working for MySql 4.1.1 and later
- Finished implementing auth for 4.0, 4.1.0, and 4.1.1
- Changed namespace from MySQL.Data.MySQLClient back to MySql.Data.MySqlClient
- Fixed bug in CharSetMapping where it was trying to use text names as ints
- Changed namespace to MySQL.Data.MySQLClient

- Integrated auth changes from UC2004
- Fixed bug where calling any of the GetXXX methods on a datareader before or after reading data would not throw the appropriate exception (thanks Luca Morelli <morelli.luca@iol.it>)
- Added TimeSpan code in parameter.cs to properly serialize a timespan object to mysql time format (thanks Gianluca Colombo <g.colombo@alfi.it>)
- Added TimeStamp to parameter serialization code. Prevented DataAdapter updates from working right (thanks Michael King)
- Fixed a misspelling in MySqlHelper.cs (thanks Patrick Kristiansen)

### D.7.13. Version 0.76

- Driver now using charset number given in handshake to create encoding
- Changed command editor to point to MySqlConnection.Design
- Fixed bug in Version.isAtLeast
- Changed DBConnectionString to support changes done to MySqlConnectionString
- Removed SqlCommandEditor and DataAdapterPreviewDialog
- Using new long return values in many places
- Integrated new CompressedStream class
- Changed ConnectionString and added attributes to allow it to be used in MySqlConnection.Design
- Changed packet.cs to support newer lengths in ReadLenInteger
- changed other classes to use new properties and fields of MySqlConnectionString
- ConnectionInternal is now using PING to see whether the server is alive
- Moved toolbox bitmaps into resource/
- Changed field.cs to allow values to come directly from row buffer
- Changed to use the new driver.Send syntax
- Using a new packet queueing system
- started work handling the "broken" compression packet handling
- Fixed bug in StreamCreator where failure to connect to a host would continue to loop infinitely (thanks Kevin Casella)
- Improved connectstring handling
- Moved designers into Pro product
- Removed some old commented out code from command.cs
- Fixed a problem with compression

- Fixed connection object where an exception throw prior to the connection opening would not leave the connection in the connecting state (thanks Chris Cline )
- Added GUID support
- Fixed sequence out of order bug (thanks Mark Reay)

## D.7.14. Version 0.75

- Enum values now supported as parameter values (thanks Philipp Sumi)
- Year datatype now supported
- fixed compression
- Fixed bug where a parameter with a TimeSpan as the value would not serialize properly
- Fixed bug where default ctor would not set default connection string values
- Added some XML comments to some members
- Work to fix/improve compression handling
- Improved ConnectionString handling so that it better matches the standard set by SqlConnection.
- A MySqlConnection is now thrown if a username is not included in the connection string
- Localhost is now used as the default if not specified on the connection string
- An exception is now thrown if an attempt is made to set the connection string while the connection is open
- Small changes to ConnectionString docs
- Removed MultiHostStream and MySqlStream. Replaced it with Common/StreamCreator
- Added support for Use Pipe connection string value
- Added Platform class for easier access to platform utility functions
- Fixed small pooling bug where new connection was not getting created after IsAlive fails
- Added Platform.cs and StreamCreator.cs
- Fixed Field.cs to properly handle 4.1 style timestamps
- Changed Common.Version to Common.DBVersion to avoid name conflict
- Fixed field.cs so that text columns return the right field type (thanks beni27@gmx.net)
- Added MySqlConnection class to provide some reference for error codes (thanks Geert Veenstra)

## D.7.15. Version 0.74

- Added Unix socket support (thanks Mohammad DAMT [md@mt.web.id])

- only calling Thread.Sleep when no data is available
- improved escaping of quote characters in parameter data
- removed misleading comments from parameter.cs
- fixed pooling bug
- same pooling bug fixed again!! ;-)
- Fixed ConnectionString editor dialog (thanks marco p (pomarc))
- UserId now supported in connection strings (thanks Jeff Neeley)
- Attempting to create a parameter that is not input throws an exception (thanks Ryan Gregg)
- Added much documentation
- checked in new MultiHostStream capability. Big thanks to Dan Guisinger for this. he originally submitted the code and idea of supporting multiple machines on the connect string.
- Added alot of documentation. Still alot to do.
- Fixed speed issue with 0.73
- changed to Thread.Sleep(0) in MySqlDataStream to help optimize the case where it doesn't need to wait (thanks Todd German)
- Prepopulating the idlepools to MinPoolSize
- Fixed MySqlPool deadlock condition as well as stupid bug where CreateNewPooledConnection was not ever adding new connections to the pool. Also fixed MySqlStream.ReadBytes and ReadByte to not use TicksPerSecond which does not appear to always be right. (thanks Matthew J. Peddlesden)
- Fix for precision and scale (thanks Matthew J. Peddlesden)
- Added Thread.Sleep(1) to stream reading methods to be more cpu friendly (thanks Sean McGinnis)
- Fixed problem where ExecuteReader would sometime return null (thanks Lloyd Dupont )
- Fixed major bug with null field handling (thanks Naucki)
- enclosed queries for max\_allowed\_packet and charset inside try catch (and set defaults)
- fixed problem where socket was not getting closed properly (thanks Steve!)
- Fixed problem where ExecuteNonQuery was not always returning the right value
- Fixed InternalConnection to not use @@session.max\_allowed\_packet but use @@max\_allowed\_packet. (Thanks Miguel)
- Added many new XML doc lines
- Fixed sql parsing to not send empty queries (thanks Rory)
- Fixed problem where the reader was not unpeeking the packet on close
- Fixed problem where user variables were not being handled (thanks Sami Vaaraniemi)
- Fixed loop checking in the MySqlPool (thanks Steve M. Brown)

- Fixed ParameterCollection.Add method to match SqlClient (thanks Joshua Mouch)
- Fixed ConnectionString parsing to handle no and yes for boolean and not lowercase values (thanks Naucki)
- Added InternalConnection class, changes to pooling
- Implemented Persist Security Info
- Added security.cs and version.cs to project
- Fixed DateTime handling in Parameter.cs (thanks Burkhard Perens-Golomb)
- Fixed parameter serialization where some types would throw a cast exception
- Fixed DataReader to convert all returned values to prevent casting errors (thanks Keith Murray)
- Added code to Command.ExecuteReader to return null if the initial SQL command throws an exception (thanks Burkhard Perens-Golomb)
- Fixed ExecuteScalar bug introduced with restructure
- Restructure to allow for LOCAL DATA INFILE and better sequencing of packets
- Fixed several bugs related to restructure.
- Early work done to support more secure passwords in Mysql 4.1. Old passwords in 4.1 not supported yet
- Parameters appearing after system parameters are now handled correctly (Adam M. (adammil))
- strings can now be assigned directly to blob fields (Adam M.)
- Fixed float parameters (thanks Pent)
- Improved Parameter ctor and ParameterCollection.Add methods to better match SqlClient (thx Joshua Mouch )
- Corrected Connection.CreateCommand to return a MySqlCommand type
- Fixed connection string designer dialog box problem (thanks Abraham Guyt)
- Fixed problem with sending commands not always reading the response packet (thanks Joshua Mouch )
- Fixed parameter serialization where some blobs types were not being handled (thanks Sean McGinnis )
- Removed spurious MessageBox.show from DataReader code (thanks Joshua Mouch )
- Fixed a nasty bug in the split sql code (thanks everyone! :-) )

## D.7.16. Version 0.71

- Fixed bug in MySqlStream where too much data could attempt to be read (thanks Peter Belbin)
- Implemented HasRows (thanks Nash Pherson)



- Fixed bug where tables with more than 252 columns cause an exception ( thanks Joshua Kessler )
- Fixed bug where SQL statements ending in ; would cause a problem ( thanks Shane Krueger )
- Fixed bug in driver where error messages were getting truncated by 1 character (thanks Shane Krueger)
- Made MySqlConnection serializable (thanks Mathias Hasselmann)

## D.7.17. Version 0.70

- Updated some of the character code pages to be more accurate
- Fixed problem where readers could be opened on connections that had readers open
- Release of 0.70
- Moved test to separate assembly MySqlConnectionTests
- Fixed stupid problem in driver with sequence out of order (Thanks Peter Belbin)
- Added some pipe tests
- Increased default max pool size to 50
- Compiles with Mono 0-24
- Fixed connection and data reader dispose problems
- Added String datatype handling to parameter serialization
- Fixed sequence problem in driver that occurred after thrown exception (thanks Burkhard Perken-Golomb)
- Added support for CommandBehavior.SingleRow to DataReader
- Fixed command sql processing so quotes are better handled (thanks Theo Spears)
- Fixed parsing of double, single, and decimal values to account for non-English separators. You still have to use the right syntax if you using hard coded sql, but if you use parameters the code will convert floating point types to use '.' appropriately internal both into the server and out. [ Thanks anonymous ]
- Added MySqlStream class to simplify timeOuts and driver coding.
- Fixed DataReader so that it is closed properly when the associated connection is closed. [thanks smishra]
- Made client more SqlConnection compliant so that DataReaders have to be closed before the connection can be used to run another command
- Improved DBNull.Value handling in the fields
- Added several unit tests
- Fixed MySqlConnection so that the base class is actually called :-o

- Improved driver coding
- Fixed bug where NextResult was returning false on the last resultset
- Added more tests for MySQL
- Improved casting problems by equating unsigned 32bit values to Int64 and unsigned 16bit values to Int32, and so forth.
- Added new ctor for MySqlParameter for (name, type, size, srccol)
- Fixed bug in MySqlDataReader where it didn't check for null fieldlist before returning field count
- Started adding MySqlClient unit tests (added MySqlClient/Tests folder and some test cases)
- Fixed some things in Connection String handling
- Moved INIT\_DB to MySqlPool. I may move it again, this is in preparation of the conference.
- Fixed bug inside CommandBuilder that prevented inserts from happening properly
- Reworked some of the internals so that all three execute methods of Command worked properly
- Fixed many small bugs found during benchmarking
- The first cut of ConnectionPooling is working. "min pool size" and "max pool size" are respected.
- Work to enable multiple resultsets to be returned
- Character sets are handled much more intelligently now. The driver queries MySQL at startup for the default character set. That character set is then used for conversions if that code page can be loaded. If not, then the default code page for the current OS is used.
- Added code to save the inferred type in the name,value ctor of Parameter
- Also, inferred type if value of null parameter is changed using Value property
- Converted all files to use proper Camel case. MySQL is now MySql in all files. PostgreSQL is now Pg-Sql
- Added attribute to PgSql code to prevent designer from trying to show
- Added MySQLDbType property to Parameter object and added proper conversion code to convert from DbType to MySQLDbType)
- Removed unused ObjectToString method from MySQLParameter.cs
- Fixed Add(..) method in ParameterCollection so that it doesn't use Add(name, value) instead.
- Fixed IndexOf and Contains in ParameterCollection to be aware that parameter names are now stored without @
- Fixed Command.ConvertSQLToBytes so it only allows characters that can be in MySQL variable names
- Fixed DataReader and Field so that blob fields read their data from Field.cs and GetBytes works right
- Added simple query builder editor to CommandText property of MySQLCommand

- Fixed CommandBuilder and Parameter serialization to account for Parameters not storing @ in their names
- Removed MySQLFieldType enum from Field.cs. Now using MySQLDbType enum
- Added Designer attribute to several classes to prevent designer view when using VS.Net
- Fixed Initial catalog typo in ConnectionString designer
- Removed 3 parameter ctor for MySQLParameter that conflicted with (name, type, value)
- changed MySQLParameter so paramName is now stored without leading @ (this fixed null inserts when using designer)
- Changed TypeConverter for MySQLParameter to use the ctor with all properties

### **D.7.18. Version 0.68**

- Fixed sequence issue in driver
- Added DbParametersEditor to make parameter editing more like SqlClient
- Fixed Command class so that parameters can be edited using the designer
- Update connection string designer to support Use Compression flag
- Fixed string encoding so that European characters like ä will work correctly
- Creating base classes to aid in building new data providers
- Added support for UID key in connection string
- Field, parameter, command now using DBNull.Value instead of null
- CommandBuilder using DBNull.Value
- CommandBuilder now builds insert command correctly when an auto\_insert field is not present
- Field now uses typeof keyword to return System.Types (performance)

### **D.7.19. Version 0.65**

- MySQLCommandBuilder now implemented
- Transaction support now implemented (not all table types support this)
- GetSchemaTable fixed to not use xsd (for Mono)
- Driver is now Mono-compatible!!
- TIME data type now supported
- More work to improve Timestamp data type handling
- Changed signatures of all classes to match corresponding SqlClient classes

## D.7.20. Version 0.60

- Protocol compression using SharpZipLib ([www.icsharpcode.net](http://www.icsharpcode.net))
- Named pipes on Windows now working properly
- Work done to improve Timestamp data type handling
- Implemented IEnumerable on DataReader so DataGrid would work

## D.7.21. Version 0.50

- Speed increased dramatically by removing bugging network sync code
- Driver no longer buffers rows of data (more ADO.Net compliant)
- Conversion bugs related to TIMESTAMP and DATETIME fields fixed

## D.8. MySQL Connector/J Change History

### D.8.1. Changes in MySQL Connector/J 5.0.3 (26 July 2006)

- Fixed `Statement.cancel()` causes `NullPointerException` if underlying connection has been closed due to server failure. ([Bug#20650](http://bugs.mysql.com/20650) [<http://bugs.mysql.com/20650>])
- Added configuration option `noAccessToProcedureBodies` which will cause the driver to create basic parameter metadata for `CallableStatements` when the user does not have access to procedure bodies via `SHOW CREATE PROCEDURE` or selecting from `mysql.proc` instead of throwing an exception. The default value for this option is `false`

### D.8.2. Changes in MySQL Connector/J 5.0.2-beta (11 July 2006)

- Fixed can't use `XAConnection` for local transactions when no global transaction is in progress. ([Bug#17401](http://bugs.mysql.com/17401) [<http://bugs.mysql.com/17401>])
- Fixed driver fails on non-ASCII platforms. The driver was assuming that the platform character set would be a superset of MySQL's `latin1` when doing the handshake for authentication, and when reading error messages. We now use Cp1252 for all strings sent to the server during the handshake phase, and a hard-coded mapping of the `language` system variable to the character set that is used for error messages. ([Bug#18086](http://bugs.mysql.com/18086) [<http://bugs.mysql.com/18086>])
- Fixed `ConnectionProperties` (and thus some subclasses) are not serializable, even though some J2EE containers expect them to be. ([Bug#19169](http://bugs.mysql.com/19169) [<http://bugs.mysql.com/19169>])
- Fixed `MysqlValidConnectionChecker` for JBoss doesn't work with `MySQLXADataSourcees`. ([Bug#20242](http://bugs.mysql.com/20242) [<http://bugs.mysql.com/20242>])
- Better caching of character set converters (per-connection) to remove a bottleneck for multibyte character sets.

- Added connection/datasource property `pinGlobalTxToPhysicalConnection` (defaults to `false`). When set to `true`, when using `XAConnections`, the driver ensures that operations on a given XID are always routed to the same physical connection. This allows the `XAConnection` to support `XA START ... JOIN` after `XA END` has been called, and is also a workaround for transaction managers that don't maintain thread affinity for a global transaction (most either always maintain thread affinity, or have it as a configuration option).
- `MysqlXaConnection.recover(int flags)` now allows combinations of `XARE-source.TMSTARTRSCAN` and `TMENDRSCAN`. To simulate the “scanning” nature of the interface, we return all prepared XIDs for `TMSTARTRSCAN`, and no new XIDs for calls with `TMNOFLAGS`, or `TMENDRSCAN` when not in combination with `TMSTARTRSCAN`. This change was made for API compliance, as well as integration with IBM WebSphere's transaction manager.

### D.8.3. Changes in MySQL Connector/J 5.0.1-beta (Not Released)

Not released due to a packaging error

### D.8.4. Changes in MySQL Connector/J 5.0.0-beta (22 December 2005)

- `XADataSource` implemented (ported from 3.2 branch which won't be released as a product). Use `com.mysql.jdbc.jdbc2.optional.MysqlXADataSource` as your datasource class name in your application server to utilize XA transactions in MySQL-5.0.10 and newer.
- `PreparedStatement.setString()` didn't work correctly when `sql_mode` on server contained `NO_BACKSLASH_ESCAPES` and no characters that needed escaping were present in the string.
- Attempt detection of the MySQL type `BINARY` (it's an alias, so this isn't always reliable), and use the `java.sql.Types.BINARY` type mapping for it.
- Moved `-bin-g.jar` file into separate `debug` subdirectory to avoid confusion.
- Don't allow `.setAutoCommit(true)`, or `.commit()` or `.rollback()` on an XA-managed connection as per the JDBC specification.
- If the connection `useTimezone` is set to `true`, then also respect time zone conversions in escape-processed string literals (for example, `"{ts ...}"` and `"{t ...}"`).
- Return original column name for `RSMD.getColumnNames()` if the column was aliased, alias name for `.getColumnLabel()` (if aliased), and original table name for `.getTableName()`. Note this only works for MySQL-4.1 and newer, as older servers don't make this information available to clients.
- Setting `useJDBCCompliantTimezoneShift=true` (it's not the default) causes the driver to use GMT for *all* `TIMESTAMP/DATETIME` time zones, and the current VM time zone for any other type that refers to time zones. This feature can not be used when `useTimezone=true` to convert between server and client time zones.
- Add one level of indirection of internal representation of `CallableStatement` parameter metadata to avoid class not found issues on JDK-1.3 for `ParameterMetadata` interface (which doesn't exist prior to JDBC-3.0).

- Added unit tests for `XADatasource`, as well as friendlier exceptions for XA failures compared to the "stock" `XAException` (which has no messages).
- Idle timeouts cause `XAConnections` to whine about rolling themselves back. ([Bug#14729](http://bugs.mysql.com/14729) [<http://bugs.mysql.com/14729>])
- Added support for Connector/MXJ integration via url subprotocol `jdbc:mysql:mxj://....`
- Moved all `SQLException` constructor usage to a factory in `SQLException` (ground-work for JDBC-4.0 `SQLState`-based exception classes).
- Removed Java5-specific calls to `BigDecimal` constructor (when result set value is `' '`, `(int)0` was being used as an argument indirectly via method return value. This signature doesn't exist prior to Java5.)
- Added service-provider entry to `META-INF/services/java.sql.Driver` for JDBC-4.0 support.
- Return "[VAR]BINARY" for `RSMD.getColumnTypeName()` when that is actually the type, and it can be distinguished (MySQL-4.1 and newer).
- When fix for [Bug#14562](http://bugs.mysql.com/14562) [<http://bugs.mysql.com/14562>] was merged from 3.1.12, added functionality for `CallableStatement`'s parameter metadata to return correct information for `.getParameterClassName()`.
- Fuller synchronization of `Connection` to avoid deadlocks when using multithreaded frameworks that multithread a single connection (usually not recommended, but the JDBC spec allows it anyways), part of fix to [Bug#14972](http://bugs.mysql.com/14972) [<http://bugs.mysql.com/14972>]).
- Implementation of `Statement.cancel()` and `Statement.setQueryTimeout()`. Both require MySQL-5.0.0 or newer server, require a separate connection to issue the `KILL QUERY` statement, and in the case of `setQueryTimeout()` creates an additional thread to handle the timeout functionality.

Note: Failures to cancel the statement for `setQueryTimeout()` may manifest themselves as `RuntimeExceptions` rather than failing silently, as there is currently no way to unblock the thread that is executing the query being cancelled due to timeout expiration and have it throw the exception instead.

## D.8.5. Changes in MySQL Connector/J 3.1.14 (not yet released)

- Fixed updatable result set throws `ClassCastException` when there is row data and `moveToInsertRow()` is called. (Fixes [Bug#20479](http://bugs.mysql.com/20479) [<http://bugs.mysql.com/20479>])
- Fixed Updatable result set that contains a BIT column fails when server-side prepared statements are used. (Fixes [Bug#20485](http://bugs.mysql.com/20485) [<http://bugs.mysql.com/20485>])
- Fixed memory leak with `profileSQL=true`. (Fixes [Bug#16987](http://bugs.mysql.com/16987) [<http://bugs.mysql.com/16987>])
- Connection fails to localhost when using timeout and IPv6 is configured. (Fixes [Bug#19726](http://bugs.mysql.com/19726) [<http://bugs.mysql.com/19726>])
- Fixed `NullPointerException` in `MysqlDataSourceFactory` due to Reference containing `RefAddr`s with null content. (Fixes [Bug#16791](http://bugs.mysql.com/16791) [<http://bugs.mysql.com/16791>])

- Fixed `ResultSet.getShort()` for `UNSIGNED TINYINT` returns incorrect values when using server-side prepared statements. (Fixes [Bug#20306](http://bugs.mysql.com/20306) [<http://bugs.mysql.com/20306>])
- Fixed can't pool server-side prepared statements, exception raised when re-using them. (Fixes [Bug#20687](http://bugs.mysql.com/20687) [<http://bugs.mysql.com/20687>])
- Fixed [Bug#21062](http://bugs.mysql.com/21062) [<http://bugs.mysql.com/21062>] - `ResultSet.getSomeInteger()` doesn't work for `BIT(>1)`.
- Fixed [Bug#18880](http://bugs.mysql.com/18880) [<http://bugs.mysql.com/18880>] - `ResultSet.getFloatFromString()` can't retrieve values near `Float.MIN/MAX_VALUE`.
- Fixed [Bug#20888](http://bugs.mysql.com/20888) [<http://bugs.mysql.com/20888>] - escape of quotes in client-side prepared statements parsing not respected. Patch covers more than bug report, including `NO_BACKSLASH_ESCAPES` being set, and stacked quote characters forms of escaping (i.e. " or """).
- Fixed [Bug#19993](http://bugs.mysql.com/19993) [<http://bugs.mysql.com/19993>] - `ReplicationDriver` does not always round-robin load balance depending on URL used for slaves list.

## D.8.6. Changes in MySQL Connector/J 3.1.13 (26 May 2006)

- `INOUT` parameter does not store `IN` value. ([Bug#15464](http://bugs.mysql.com/15464) [<http://bugs.mysql.com/15464>])
- Exception thrown for new decimal type when using updatable result sets. ([Bug#14609](http://bugs.mysql.com/14609) [<http://bugs.mysql.com/14609>])
- No "dos" character set in MySQL > 4.1.0. ([Bug#15544](http://bugs.mysql.com/15544) [<http://bugs.mysql.com/15544>])
- `PreparedStatement.setObject()` serializes `BigInteger` as object, rather than sending as numeric value (and is thus not complementary to `.getObject()` on an `UNSIGNED LONG` type). ([Bug#15383](http://bugs.mysql.com/15383) [<http://bugs.mysql.com/15383>])
- `ResultSet.getShort()` for `UNSIGNED TINYINT` returned wrong values. ([Bug#11874](http://bugs.mysql.com/11874) [<http://bugs.mysql.com/11874>])
- `lib-nodist` directory missing from package breaks out-of-box build. ([Bug#15676](http://bugs.mysql.com/15676) [<http://bugs.mysql.com/15676>])
- `DBMD.getColumns()` returns wrong type for `BIT`. ([Bug#15854](http://bugs.mysql.com/15854) [<http://bugs.mysql.com/15854>])
- Fixed issue where driver was unable to initialize character set mapping tables. Removed reliance on `.properties` files to hold this information, as it turns out to be too problematic to code around class loader hierarchies that change depending on how an application is deployed. Moved information back into the `CharsetMapping` class. ([Bug#14938](http://bugs.mysql.com/14938) [<http://bugs.mysql.com/14938>])
- Fixed updatable result set doesn't return `AUTO_INCREMENT` values for `insertRow()` when multiple column primary keys are used. (the driver was checking for the existence of single-column primary keys and an autoincrement value > 0 instead of a straightforward `isAutoIncrement()` check). ([Bug#16841](http://bugs.mysql.com/16841) [<http://bugs.mysql.com/16841>])
- Fixed `Statement.getGeneratedKeys()` throws `NullPointerException` when no query has been processed. ([Bug#17099](http://bugs.mysql.com/17099) [<http://bugs.mysql.com/17099>])
- Fixed driver trying to call methods that don't exist on older and newer versions of Log4j. The fix is not trying to auto-detect presence of log4j, too many different incompatible versions out there in the wild to do this reliably. ([Bug#13469](http://bugs.mysql.com/13469) [<http://bugs.mysql.com/13469>])

If you relied on autodetection before, you will need to add "logger=com.mysql.jdbc.log.Log4JLogger" to your JDBC URL to enable Log4J usage, or alternatively use the new "CommonsLogger" class to take care of this.

- Added support for Apache Commons logging, use "com.mysql.jdbc.log.CommonsLogger" as the value for the "logger" configuration property.
- LogFactory now prepends "com.mysql.jdbc.log" to log class name if it can't be found as-specified. This allows you to use "short names" for the built-in log factories, for example "logger=CommonsLogger" instead of "logger=com.mysql.jdbc.log.CommonsLogger".
- Fixed issue with `ReplicationConnection` incorrectly copying state, doesn't transfer connection context correctly when transitioning between the same read-only states. ([Bug#15570](http://bugs.mysql.com/15570) [<http://bugs.mysql.com/15570>])
- Fixed issue where server-side prepared statements don't cause truncation exceptions to be thrown when truncation happens. ([Bug#18041](http://bugs.mysql.com/18041) [<http://bugs.mysql.com/18041>])
- Added performance feature, re-writing of batched executes for `Statement.executeBatch()` (for all DML statements) and `PreparedStatement.executeBatch()` (for INSERTs with VALUE clauses only). Enable by using "rewriteBatchedStatements=true" in your JDBC URL.
- Fixed `CallableStatement.registerOutParameter()` not working when some parameters pre-populated. Still waiting for feedback from JDBC experts group to determine what correct parameter count from `getMetaData()` should be, however. ([Bug#17898](http://bugs.mysql.com/17898) [<http://bugs.mysql.com/17898>])
- Fixed calling `clearParameters()` on a closed prepared statement causes NPE. ([Bug#17587](http://bugs.mysql.com/17587) [<http://bugs.mysql.com/17587>])
- Map "latin1" on MySQL server to CP1252 for MySQL > 4.1.0.
- Added additional accessor and mutator methods on `ConnectionProperties` so that `DataSource` users can use same naming as regular URL properties.
- Fixed data truncation and `getWarnings()` only returns last warning in set. ([Bug#18740](http://bugs.mysql.com/18740) [<http://bugs.mysql.com/18740>])
- Improved performance of retrieving `BigDecimal`, `Time`, `Timestamp` and `Date` values from server-side prepared statements by creating fewer short-lived instances of `Strings` when the native type is not an exact match for the requested type. Fixes [Bug#18496](http://bugs.mysql.com/18496) [<http://bugs.mysql.com/18496>] for `BigDecimals`.
- Fixed aliased column names where length of name > 251 are corrupted. ([Bug#18554](http://bugs.mysql.com/18554) [<http://bugs.mysql.com/18554>])
- Fixed `ResultSet.wasNull()` not always reset correctly for booleans when done via conversion for server-side prepared statements. ([Bug#17450](http://bugs.mysql.com/17450) [<http://bugs.mysql.com/17450>])
- Fixed invalid classname returned for `ResultSetMetaData.getColumnClassName()` for `BIGINT` type. ([Bug#19282](http://bugs.mysql.com/19282) [<http://bugs.mysql.com/19282>])
- Fixed case where driver wasn't reading server status correctly when fetching server-side prepared statement rows, which in some cases could cause warning counts to be off, or multiple result sets to not be read off the wire.
- Driver now aware of fix for `BIT` type metadata that went into MySQL-5.0.21 for server not reporting length consistently ([Bug#13601](http://bugs.mysql.com/13601) [<http://bugs.mysql.com/13601>]).



- Fixed `PreparedStatement.setObject(int, Object, int)` doesn't respect scale of `BigDecimal`s. (Bug#19615 [<http://bugs.mysql.com/19615>])
- Fixed `ResultSet.wasNull()` returns incorrect value when extracting native string from server-side prepared statement generated result set. (Bug#19282 [<http://bugs.mysql.com/19282>])

## D.8.7. Changes in MySQL Connector/J 3.1.12 (30 November 2005)

- Fixed client-side prepared statement bug with embedded `?` characters inside quoted identifiers (it was recognized as a placeholder, when it was not).
- Don't allow `executeBatch()` for `CallableStatements` with registered `OUT/INOUT` parameters (JDBC compliance).
- Fall back to platform-encoding for `URLDecoder.decode()` when parsing driver URL properties if the platform doesn't have a two-argument version of this method.
- Java type conversion may be incorrect for `MEDIUMINT`. (Bug#14562 [<http://bugs.mysql.com/14562>])
- Added configuration property `useGmtMillisForDatetimes` which when set to `true` causes `ResultSet.getDate()`, `ResultSet.getTimestamp()` to return correct millis-since GMT when `ResultSet.getTime()` is called on the return value (currently default is `false` for legacy behavior).
- Fixed `DatabaseMetaData.stores*Identifiers()`:
  - If `lower_case_table_names=0` (on server):
    - `storesLowerCaseIdentifiers()` returns `false`
    - `storesLowerCaseQuotedIdentifiers()` returns `false`
    - `storesMixedCaseIdentifiers()` returns `true`
    - `storesMixedCaseQuotedIdentifiers()` returns `true`
    - `storesUpperCaseIdentifiers()` returns `false`
    - `storesUpperCaseQuotedIdentifiers()` returns `true`
  - If `lower_case_table_names=1` (on server):
    - `storesLowerCaseIdentifiers()` returns `true`
    - `storesLowerCaseQuotedIdentifiers()` returns `true`
    - `storesMixedCaseIdentifiers()` returns `false`
    - `storesMixedCaseQuotedIdentifiers()` returns `false`
    - `storesUpperCaseIdentifiers()` returns `false`
    - `storesUpperCaseQuotedIdentifiers()` returns `true`
- `DatabaseMetaData.getColumns()` doesn't return `TABLE_NAME` correctly. (Bug#14815)

[<http://bugs.mysql.com/14815>])

- Escape processor replaces quote character in quoted string with string delimiter. (Bug#14909 [<http://bugs.mysql.com/14909>])
- OpenOffice expects `DBMD.supportsIntegrityEnhancementFacility()` to return `true` if foreign keys are supported by the datasource, even though this method also covers support for check constraints, which MySQL *doesn't* have. Setting the configuration property `overrideSupportsIntegrityEnhancementFacility` to `true` causes the driver to return `true` for this method. (Bug#12975 [<http://bugs.mysql.com/12975>])
- Added `com.mysql.jdbc.testsuite.url.default` system property to set default JDBC url for testsuite (to speed up bug resolution when I'm working in Eclipse).
- Unable to initialize character set mapping tables (due to J2EE classloader differences). (Bug#14938 [<http://bugs.mysql.com/14938>])
- Deadlock while closing server-side prepared statements from multiple threads sharing one connection. (Bug#14972 [<http://bugs.mysql.com/14972>])
- `logSlowQueries` should give better info. (Bug#12230 [<http://bugs.mysql.com/12230>])
- Extraneous sleep on `autoReconnect`. (Bug#13775 [<http://bugs.mysql.com/13775>])
- Driver incorrectly closes streams passed as arguments to `PreparedStatement`. Reverts to legacy behavior by setting the JDBC configuration property `autoClosePstmtStreams` to `true` (also included in the 3-0-Compat configuration “bundle”). (Bug#15024 [<http://bugs.mysql.com/15024>])
- `maxQuerySizeToLog` is not respected. Added logging of bound values for `execute()` phase of server-side prepared statements when `profileSQL=true` as well. (Bug#13048 [<http://bugs.mysql.com/13048>])
- Usage advisor complains about unreferenced columns, even though they've been referenced. (Bug#15065 [<http://bugs.mysql.com/15065>])
- Don't increase timeout for failover/reconnect. (Bug#6577 [<http://bugs.mysql.com/6577>])
- Process escape tokens in `Connection.prepareStatement(...)`. (Bug#15141 [<http://bugs.mysql.com/15141>]) You can disable this behavior by setting the JDBC URL configuration property `processEscapeCodesForPrepStmts` to `false`.
- Reconnect during middle of `executeBatch()` should not occur if `autoReconnect` is enabled. (Bug#13255 [<http://bugs.mysql.com/13255>])

## D.8.8. Changes in MySQL Connector/J 3.1.11-stable (07 October 2005)

- Spurious `!` on console when character encoding is `utf8`. (Bug#11629 [<http://bugs.mysql.com/11629>])
- Fixed statements generated for testcases missing `;` for “plain” statements.
- Incorrect generation of testcase scripts for server-side prepared statements. (Bug#11663 [<http://bugs.mysql.com/11663>])

- Fixed regression caused by fix for [Bug#11552](http://bugs.mysql.com/11552) [http://bugs.mysql.com/11552] that caused driver to return incorrect values for unsigned integers when those integers were within the range of the positive signed type.
- Moved source code to Subversion repository.
- Escape tokenizer doesn't respect stacked single quotes for escapes. ([Bug#11797](http://bugs.mysql.com/11797) [http://bugs.mysql.com/11797])
- `GEOMETRY` type not recognized when using server-side prepared statements.
- `ReplicationConnection` won't switch to slave, throws “Catalog can't be null” exception. ([Bug#11879](http://bugs.mysql.com/11879) [http://bugs.mysql.com/11879])
- Properties shared between master and slave with replication connection. ([Bug#12218](http://bugs.mysql.com/12218) [http://bugs.mysql.com/12218])
- `Statement.getWarnings()` fails with NPE if statement has been closed. ([Bug#10630](http://bugs.mysql.com/10630) [http://bugs.mysql.com/10630])
- Only get `char[]` from SQL in `PreparedStatement.ParseInfo()` when needed.
- Geometry types not handled with server-side prepared statements. ([Bug#12104](http://bugs.mysql.com/12104) [http://bugs.mysql.com/12104])
- `StringUtils.getBytes()` doesn't work when using multi-byte character encodings and a length in `characters` is specified. ([Bug#11614](http://bugs.mysql.com/11614) [http://bugs.mysql.com/11614])
- `Pstmt.setObject(..., Types.BOOLEAN)` throws exception. ([Bug#11798](http://bugs.mysql.com/11798) [http://bugs.mysql.com/11798])
- `maxPerformance.properties` mis-spells “elideSetAutoCommits”. ([Bug#11976](http://bugs.mysql.com/11976) [http://bugs.mysql.com/11976])
- `DBMD.storesLower/Mixed/UpperIdentifiers()` reports incorrect values for servers deployed on Windows. ([Bug#11575](http://bugs.mysql.com/11575) [http://bugs.mysql.com/11575])
- `ResultSet.moveToCurrentRow()` fails to work when preceded by a call to `ResultSet.moveToInsertRow()`. ([Bug#11190](http://bugs.mysql.com/11190) [http://bugs.mysql.com/11190])
- `VARBINARY` data corrupted when using server-side prepared statements and `.setBytes()`. ([Bug#11115](http://bugs.mysql.com/11115) [http://bugs.mysql.com/11115])
- `explainSlowQueries` hangs with server-side prepared statements. ([Bug#12229](http://bugs.mysql.com/12229) [http://bugs.mysql.com/12229])
- Escape processor didn't honor strings demarcated with double quotes. ([Bug#11498](http://bugs.mysql.com/11498) [http://bugs.mysql.com/11498])
- Lifted restriction of changing streaming parameters with server-side prepared statements. As long as `all` streaming parameters were set before execution, `.clearParameters()` does not have to be called. (due to limitation of client/server protocol, prepared statements can not reset *individual* stream data on the server side).
- Reworked `Field` class, `*Buffer`, and `MysqlIO` to be aware of field lengths > `Integer.MAX_VALUE`.
- Updated `DBMD.supportsCorrelatedQueries()` to return `true` for versions > 4.1, `supportsGroupByUnrelated()` to return `true` and `getResultSetHoldability()` to re-

turn `HOLD_CURSORS_OVER_COMMIT`.

- Handling of catalog argument in `DatabaseMetaData.getIndexInfo()`, which also means changes to the following methods in `DatabaseMetaData`: ([Bug#12541](http://bugs.mysql.com/12541) [<http://bugs.mysql.com/12541>])
  - `getBestRowIdentifier()`
  - `getColumns()`
  - `getCrossReference()`
  - `getExportedKeys()`
  - `getImportedKeys()`
  - `getIndexInfo()`
  - `getPrimaryKeys()`
  - `getProcedures()` (and thus indirectly `getProcedureColumns()`)
  - `getTables()`

The `catalog` argument in all of these methods now behaves in the following way:

- Specifying `NULL` means that catalog will not be used to filter the results (thus all databases will be searched), unless you've set `nullCatalogMeansCurrent=true` in your JDBC URL properties.
- Specifying `" "` means “current” catalog, even though this isn't quite JDBC spec compliant, it's there for legacy users.
- Specifying a catalog works as stated in the API docs.
- Made `Connection.clientPrepare()` available from “wrapped” connections in the `jdbc2.optional` package (connections built by `ConnectionPoolDataSource` instances).
- Added `Connection.isMasterConnection()` for clients to be able to determine if a multi-host master/slave connection is connected to the first host in the list.
- Tokenizer for `=` in URL properties was causing `sessionVariables=...` to be parameterized incorrectly. ([Bug#12753](http://bugs.mysql.com/12753) [<http://bugs.mysql.com/12753>])
- Foreign key information that is quoted is parsed incorrectly when `DatabaseMetaData` methods use that information. ([Bug#11781](http://bugs.mysql.com/11781) [<http://bugs.mysql.com/11781>])
- The `sendBlobChunkSize` property is now clamped to `max_allowed_packet` with consideration of stream buffer size and packet headers to avoid `PacketTooBigExceptions` when `max_allowed_packet` is similar in size to the default `sendBlobChunkSize` which is 1M.
- `CallableStatement.clearParameters()` now clears resources associated with `INOUT/OUTPUT` parameters as well as `INPUT` parameters.
- `Connection.prepareCall()` is database name case-sensitive (on Windows systems). ([Bug#12417](http://bugs.mysql.com/12417) [<http://bugs.mysql.com/12417>])
- `cp1251` incorrectly mapped to `win1251` for servers newer than 4.0.x. ([Bug#12752](http://bugs.mysql.com/12752) [<http://bugs.mysql.com/12752>])

- `java.sql.Types.OTHER` returned for `BINARY` and `VARBINARY` columns when using `DatabaseMetaData.getColumns()`. (Bug#12970 [http://bugs.mysql.com/12970])
- `ServerPreparedStatement.getBinding()` now checks if the statement is closed before attempting to reference the list of parameter bindings, to avoid throwing a `NullPointerException`.
- `ResultSetMetaData` from `Statement.getGeneratedKeys()` caused a `NullPointerException` to be thrown whenever a method that required a connection reference was called. (Bug#13277 [http://bugs.mysql.com/13277])
- Backport of `Field` class, `ResultSetMetaData.getColumnClassName()`, and `ResultSet.getObject(int)` changes from 5.0 branch to fix behavior surrounding `VARCHAR BINARY/VARBINARY` and related types.
- Fixed `NullPointerException` when converting `catalog` parameter in many `DatabaseMetaDataMethods` to `byte[]`s (for the result set) when the parameter is `null`. (`null` isn't technically allowed by the JDBC specification, but we've historically allowed it).
- Backport of `VAR[BINARY|CHAR] [BINARY]` types detection from 5.0 branch.
- Read response in `MysqlIO.sendFileToServer()`, even if the local file can't be opened, otherwise next query issued will fail, because it's reading the response to the empty `LOAD DATA INFILE` packet sent to the server.
- Workaround for Bug#13374 [http://bugs.mysql.com/13374]: `ResultSet.getStatement()` on closed result set returns `NULL` (as per JDBC 4.0 spec, but not backward-compatible). Set the connection property `retainStatementAfterResultSetClose` to `true` to be able to retrieve a `ResultSet`'s statement after the `ResultSet` has been closed via `.getStatement()` (the default is `false`, to be JDBC-compliant and to reduce the chance that code using JDBC leaks `Statement` instances).
- URL configuration parameters don't allow `'&'` or `'='` in their values. The JDBC driver now parses configuration parameters as if they are encoded using the application/x-www-form-urlencoded format as specified by `java.net.URLDecoder` (<http://java.sun.com/j2se/1.5.0/docs/api/java/net/URLDecoder.html>). (Bug#13453 [http://bugs.mysql.com/13453])  
  
If the `'%'` character is present in a configuration property, it must now be represented as `%25`, which is the encoded form of `'%'` when using application/x-www-form-urlencoded encoding.
- The configuration property `sessionVariables` now allows you to specify variables that start with the `'@'` sign.
- When `gatherPerfMetrics` is enabled for servers older than 4.1.0, a `NullPointerException` is thrown from the constructor of `ResultSet` if the query doesn't use any tables. (Bug#13043 [http://bugs.mysql.com/13043])

## D.8.9. Changes in MySQL Connector/J 3.1.10-stable (23 June 2005)

- Fixed connecting without a database specified raised an exception in `MysqlIO.changeDatabaseTo()`.
- Initial implementation of `ParameterMetadata` for `PreparedStatement.getParameterMetadata()`. Only works fully for `CallableStatements`, as current

server-side prepared statements return every parameter as a `VARCHAR` type.

## D.8.10. Changes in MySQL Connector/J 3.1.9-stable (22 June 2005)

- Overhaul of character set configuration, everything now lives in a properties file.
- Driver now correctly uses CP932 if available on the server for Windows-31J, CP932 and MS932 java encoding names, otherwise it resorts to SJIS, which is only a close approximation. Currently only MySQL-5.0.3 and newer (and MySQL-4.1.12 or .13, depending on when the character set gets backported) can reliably support any variant of CP932.
- `com.mysql.jdbc.PreparedStatement.ParseInfo` does unnecessary call to `toCharArray()`. (Bug#9064 [<http://bugs.mysql.com/9064>])
- Memory leak in `ServerPreparedStatement` if `serverPrepare()` fails. (Bug#10144 [<http://bugs.mysql.com/10144>])
- Actually write manifest file to correct place so it ends up in the binary jar file.
- Added `createDatabaseIfNotExist` property (default is `false`), which will cause the driver to ask the server to create the database specified in the URL if it doesn't exist. You must have the appropriate privileges for database creation for this to work.
- Unsigned `SMALLINT` treated as signed for `ResultSet.getInt()`, fixed all cases for `UNSIGNED` integer values and server-side prepared statements, as well as `ResultSet.getObject()` for `UNSIGNED TINYINT`. (Bug#10156 [<http://bugs.mysql.com/10156>])
- Double quotes not recognized when parsing client-side prepared statements. (Bug#10155 [<http://bugs.mysql.com/10155>])
- Made `enableStreamingResults()` visible on `com.mysql.jdbc.jdbc2.optional.StatementWrapper`.
- Made `ServerPreparedStatement.asSql()` work correctly so auto-explain functionality would work with server-side prepared statements.
- Made JDBC2-compliant wrappers public in order to allow access to vendor extensions.
- Cleaned up logging of profiler events, moved code to dump a profiler event as a string to `com.mysql.jdbc.log.LogUtils` so that third parties can use it.
- `DatabaseMetaData.supportsMultipleOpenResults()` now returns `true`. The driver has supported this for some time, DBMD just missed that fact.
- Driver doesn't support `{?=CALL(...)}` for calling stored functions. This involved adding support for function retrieval to `DatabaseMetaData.getProcedures()` and `getProcedureColumns()` as well. (Bug#10310 [<http://bugs.mysql.com/10310>])
- `SQLException` thrown when retrieving `YEAR(2)` with `ResultSet.getString()`. The driver will now always treat `YEAR` types as `java.sql.Date` and return the correct values for `getString()`. Alternatively, the `yearIsDateType` connection property can be set to `false` and the values will be treated as `SHORTs`. (Bug#10485 [<http://bugs.mysql.com/10485>])
- The datatype returned for `TINYINT(1)` columns when `tinyIntIsBit=true` (the default) can be switched between `Types.BOOLEAN` and `Types.BIT` using the new configuration property

`transformedBitIsBoolean`, which defaults to `false`. If set to `false` (the default), `DatabaseMetaData.getColumns()` and `ResultSetMetaData.getColumnType()` will return `Types.BOOLEAN` for `TINYINT(1)` columns. If `true`, `Types.BOOLEAN` will be returned instead. Regardless of this configuration property, if `tinyIntIsBit` is enabled, columns with the type `TINYINT(1)` will be returned as `java.lang.Boolean` instances from `ResultSet.getObject(...)`, and `ResultSetMetaData.getColumnClassName()` will return `java.lang.Boolean`.

- `SQLException` is thrown when using property `characterSetResults` with `cp932` or `eu-cjpm`. ([Bug#10496](http://bugs.mysql.com/10496) [<http://bugs.mysql.com/10496>])
- Reorganized directory layout. Sources now are in `src` folder. Don't pollute parent directory when building, now output goes to `./build`, distribution goes to `./dist`.
- Added support/bug hunting feature that generates `.sql` test scripts to `STDERR` when `autoGenerateTestcaseScript` is set to `true`.
- 0-length streams not sent to server when using server-side prepared statements. ([Bug#10850](http://bugs.mysql.com/10850) [<http://bugs.mysql.com/10850>])
- Setting `cachePrepStmts=true` now causes the `Connection` to also cache the check the driver performs to determine if a prepared statement can be server-side or not, as well as caches server-side prepared statements for the lifetime of a connection. As before, the `prepStmtCacheSize` parameter controls the size of these caches.
- Try to handle `OutOfMemoryErrors` more gracefully. Although not much can be done, they will in most cases close the connection they happened on so that further operations don't run into a connection in some unknown state. When an OOM has happened, any further operations on the connection will fail with a "Connection closed" exception that will also list the OOM exception as the reason for the implicit connection close event.
- Don't send `COM_RESET_STMT` for each execution of a server-side prepared statement if it isn't required.
- Driver detects if you're running MySQL-5.0.7 or later, and does not scan for `LIMIT ?[,?]` in statements being prepared, as the server supports those types of queries now.
- `VARBINARY` data corrupted when using server-side prepared statements and `ResultSet.getBytes()`. ([Bug#11115](http://bugs.mysql.com/11115) [<http://bugs.mysql.com/11115>])
- `Connection.setCatalog()` is now aware of the `useLocalSessionState` configuration property, which when set to `true` will prevent the driver from sending `USE ...` to the server if the requested catalog is the same as the current catalog.
- Added the following configuration bundles, use one or many via the `useConfigs` configuration property:
  - `maxPerformance` — maximum performance without being reckless
  - `solarisMaxPerformance` — maximum performance for Solaris, avoids syscalls where it can
  - `3-0-Compat` — Compatibility with Connector/J 3.0.x functionality
- Added `maintainTimeStats` configuration property (defaults to `true`), which tells the driver whether or not to keep track of the last query time and the last successful packet sent to the server's time. If set to `false`, removes two syscalls per query.
- `autoReconnect` ping causes exception on connection startup. ([Bug#11259](http://bugs.mysql.com/11259))

[<http://bugs.mysql.com/11259>])

- Connector/J dumping query into `SQLException` twice. ([Bug#11360](#) [<http://bugs.mysql.com/11360>])
- Fixed `PreparedStatement.setClob()` not accepting `null` as a parameter.
- Production package doesn't include JBoss integration classes. ([Bug#11411](#) [<http://bugs.mysql.com/11411>])
- Removed nonsensical “costly type conversion” warnings when using usage advisor.

## D.8.11. Changes in MySQL Connector/J 3.1.8-stable (14 April 2005)

- Fixed `DatabaseMetaData.getTables()` returning views when they were not asked for as one of the requested table types.
- Added support for new precision-math `DECIMAL` type in MySQL 5.0.3 and up.
- Fixed `ResultSet.getTime()` on a `NULL` value for server-side prepared statements throws NPE.
- Made `Connection.ping()` a public method.
- `DATE_FORMAT()` queries returned as `BLOBs` from `getObject()`. ([Bug#8868](#) [<http://bugs.mysql.com/8868>])
- `ServerPreparedStatements` now correctly “stream” `BLOB/CLOB` data to the server. You can configure the threshold chunk size using the JDBC URL property `blobSendChunkSize` (the default is 1MB).
- `BlobFromLocator` now uses correct identifier quoting when generating prepared statements.
- Server-side session variables can be preset at connection time by passing them as a comma-delimited list for the connection property `sessionVariables`.
- Fixed regression in `ping()` for users using `autoReconnect=true`.
- `PreparedStatement.addBatch()` doesn't work with server-side prepared statements and streaming `BINARY` data. ([Bug#9040](#) [<http://bugs.mysql.com/9040>])
- `DBMD.supportsMixedCase*Identifiers()` returns wrong value on servers running on case-sensitive filesystems. ([Bug#8800](#) [<http://bugs.mysql.com/8800>])
- Cannot use `UTF-8` for `characterSetResults` configuration property. ([Bug#9206](#) [<http://bugs.mysql.com/9206>])
- A continuation of [Bug#8868](#) [<http://bugs.mysql.com/8868>], where functions used in queries that should return non-string types when resolved by temporary tables suddenly become opaque binary strings (work-around for server limitation). Also fixed fields with type of `CHAR(n) CHARACTER SET BINARY` to return correct/matching classes for `RSMD.getColumnClassName()` and `ResultSet.getObject()`. ([Bug#9236](#) [<http://bugs.mysql.com/9236>])
- `DBMD.supportsResultSetConcurrency()` not returning `true` for forward-only/read-only result sets (we obviously support this). ([Bug#8792](#) [<http://bugs.mysql.com/8792>])



- `DATA_TYPE` column from `DBMD.getBestRowIdentifier()` causes `ArrayIndexOutOfBoundsException` when accessed (and in fact, didn't return any value). (Bug#8803 [http://bugs.mysql.com/8803])
- Check for empty strings ( ' ') when converting `CHAR/VARCHAR` column data to numbers, throw exception if `emptyStringsConvertToZero` configuration property is set to `false` (for backward-compatibility with 3.0, it is now set to `true` by default, but will most likely default to `false` in 3.2).
- `PreparedStatement.getMetaData()` inserts blank row in database under certain conditions when not using server-side prepared statements. (Bug#9320 [http://bugs.mysql.com/9320])
- `Connection.canHandleAsPreparedStatement()` now makes “best effort” to distinguish `LIMIT` clauses with placeholders in them from ones without in order to have fewer false positives when generating work-arounds for statements the server cannot currently handle as server-side prepared statements.
- Fixed `build.xml` to not compile `log4j` logging if `log4j` not available.
- Added support for the c3p0 connection pool's (<http://c3p0.sf.net/>) validation/connection checker interface which uses the lightweight `COM_PING` call to the server if available. To use it, configure your c3p0 connection pool's `connectionTesterClassName` property to use `com.mysql.jdbc.integration.c3p0.MysqlConnectionTester`.
- Better detection of `LIMIT` inside/outside of quoted strings so that the driver can more correctly determine whether a prepared statement can be prepared on the server or not.
- Stored procedures with same name in different databases confuse the driver when it tries to determine parameter counts/types. (Bug#9319 [http://bugs.mysql.com/9319])
- Added finalizers to `ResultSet` and `Statement` implementations to be JDBC spec-compliant, which requires that if not explicitly closed, these resources should be closed upon garbage collection.
- Stored procedures with `DECIMAL` parameters with storage specifications that contained ' , ' in them would fail. (Bug#9682 [http://bugs.mysql.com/9682])
- `PreparedStatement.setObject(int, Object, int type, int scale)` now uses scale value for `BigDecimal` instances.
- `Statement.getMoreResults()` could throw NPE when existing result set was `.close()`d. (Bug#9704 [http://bugs.mysql.com/9704])
- The performance metrics feature now gathers information about number of tables referenced in a `SELECT`.
- The logging system is now automatically configured. If the value has been set by the user, via the URL property `logger` or the system property `com.mysql.jdbc.logger`, then use that, otherwise, autodetect it using the following steps:
  1. Log4j, if it's available,
  2. Then JDK1.4 logging,
  3. Then fallback to our `STDERR` logging.
- `DBMD.getTables()` shouldn't return tables if views are asked for, even if the database version doesn't support views. (Bug#9778 [http://bugs.mysql.com/9778])
- Fixed driver not returning `true` for `-1` when `ResultSet.getBoolean()` was called on result

sets returned from server-side prepared statements.

- Added a `Manifest.MF` file with implementation information to the `.jar` file.
- More tests in `Field.isOpaqueBinary()` to distinguish opaque binary (that is, fields with type `CHAR(n)` and `CHARACTER SET BINARY`) from output of various scalar and aggregate functions that return strings.
- Should accept `null` for catalog (meaning use current) in DBMD methods, even though it's not JDBC-compliant for legacy's sake. Disable by setting connection property `nullCatalogMeansCurrent` to `false` (which will be the default value in C/J 3.2.x). ([Bug#9917](http://bugs.mysql.com/9917) [<http://bugs.mysql.com/9917>])
- Should accept `null` for name patterns in DBMD (meaning '%'), even though it isn't JDBC compliant, for legacy's sake. Disable by setting connection property `nullNamePatternMatchesAll` to `false` (which will be the default value in C/J 3.2.x). ([Bug#9769](http://bugs.mysql.com/9769) [<http://bugs.mysql.com/9769>])

## D.8.12. Changes in MySQL Connector/J 3.1.7-stable (18 February 2005)

- Timestamp key column data needed `_binary` stripped for `UpdatableResultSet.refreshRow()`. ([Bug#7686](http://bugs.mysql.com/7686) [<http://bugs.mysql.com/7686>])
- Timestamps converted incorrectly to strings with server-side prepared statements and updatable result sets. ([Bug#7715](http://bugs.mysql.com/7715) [<http://bugs.mysql.com/7715>])
- Detect new `sql_mode` variable in string form (it used to be integer) and adjust quoting method for strings appropriately.
- Added `holdResultsOpenOverStatementClose` property (default is `false`), that keeps result sets open over `statement.close()` or new execution on same statement (suggested by Kevin Burton).
- Infinite recursion when “falling back” to master in failover configuration. ([Bug#7952](http://bugs.mysql.com/7952) [<http://bugs.mysql.com/7952>])
- Disable multi-statements (if enabled) for MySQL-4.1 versions prior to version 4.1.10 if the query cache is enabled, as the server returns wrong results in this configuration.
- Fixed duplicated code in `configureClientCharset()` that prevented `useOldUTF8Behavior=true` from working properly.
- Removed `dontUnpackBinaryResults` functionality, the driver now always stores results from server-side prepared statements as is from the server and unpacks them on demand.
- Emulated locators corrupt binary data when using server-side prepared statements. ([Bug#8096](http://bugs.mysql.com/8096) [<http://bugs.mysql.com/8096>])
- Fixed synchronization issue with `ServerPreparedStatement.serverPrepare()` that could cause deadlocks/crashes if connection was shared between threads.
- By default, the driver now scans SQL you are preparing via all variants of `Connection.prepareStatement()` to determine if it is a supported type of statement to prepare on the server side, and if it is not supported by the server, it instead prepares it as a client-side emulated prepared statement. You can disable this by passing `emulateUnsupportedPstmts=false` in your JDBC URL. ([Bug#4718](http://bugs.mysql.com/4718) [<http://bugs.mysql.com/4718>])

- Remove `_binary` introducer from parameters used as in/out parameters in `CallableStatement`.
- Always return `byte[]`s for output parameters registered as `*BINARY`.
- Send correct value for “boolean” `true` to server for `PreparedStatement.setObject(n, "true", Types.BIT)`.
- Fixed bug with `Connection` not caching statements from `prepareStatement()` when the statement wasn't a server-side prepared statement.
- Choose correct “direction” to apply time adjustments when both client and server are in GMT time zone when using `ResultSet.get(..., cal)` and `PreparedStatement.set(..., cal)`.
- Added `dontTrackOpenResources` option (default is `false`, to be JDBC compliant), which helps with memory use for non-well-behaved apps (that is, applications that don't close `Statement` objects when they should).
- `ResultSet.getString()` doesn't maintain format stored on server, bug fix only enabled when `noDatetimeStringSync` property is set to `true` (the default is `false`). (Bug#8428 [http://bugs.mysql.com/8428])
- Fixed NPE in `ResultSet.realClose()` when using usage advisor and result set was already closed.
- `PreparedStatements` not creating streaming result sets. (Bug#8487 [http://bugs.mysql.com/8487])
- Don't pass `NULL` to `String.valueOf()` in `ResultSet.getNativeConvertToString()`, as it stringifies it (that is, returns `null`), which is not correct for the method in question.
- `ResultSet.getBigDecimal()` throws exception when rounding would need to occur to set scale. The driver now chooses a rounding mode of “half up” if non-rounding `BigDecimal.setScale()` fails. (Bug#8424 [http://bugs.mysql.com/8424])
- Added `useLocalSessionState` configuration property, when set to `true` the JDBC driver trusts that the application is well-behaved and only sets autocommit and transaction isolation levels using the methods provided on `java.sql.Connection`, and therefore can manipulate these values in many cases without incurring round-trips to the database server.
- Added `enableStreamingResults()` to `Statement` for connection pool implementations that check `Statement.setFetchSize()` for specification-compliant values. Call `Statement.setFetchSize(>=0)` to disable the streaming results for that statement.
- Added support for `BIT` type in MySQL-5.0.3. The driver will treat `BIT(1-8)` as the JDBC standard `BIT` type (which maps to `java.lang.Boolean`), as the server does not currently send enough information to determine the size of a bitfield when `< 9` bits are declared. `BIT(>9)` will be treated as `VARBINARY`, and will return `byte[]` when `getObject()` is called.

### D.8.13. Changes in MySQL Connector/J 3.1.6-stable (23 December 2004)

- Fixed hang on `SocketInputStream.read()` with `Statement.setMaxRows()` and multiple result sets when driver has to truncate result set directly, rather than tacking a `LIMIT n` on the

end of it.

- `DBMD.getProcedures()` doesn't respect catalog parameter. ([Bug#7026](http://bugs.mysql.com/7026) [<http://bugs.mysql.com/7026>])

## D.8.14. Changes in MySQL Connector/J 3.1.5-gamma (02 December 2004)

- Fix comparisons made between string constants and dynamic strings that are converted with either `toUpperCase()` or `toLowerCase()` to use `Locale.ENGLISH`, as some locales “override” case rules for English. Also use `StringUtils.indexOfIgnoreCase()` instead of `.toUpperCase().indexOf()`, avoids creating a very short-lived transient `String` instance.
- Server-side prepared statements did not honor `zeroDateTimeBehavior` property, and would cause class-cast exceptions when using `ResultSet.getObject()`, as the all-zero string was always returned. ([Bug#5235](http://bugs.mysql.com/5235) [<http://bugs.mysql.com/5235>])
- Fixed batched updates with server prepared statements weren't looking if the types had changed for a given batched set of parameters compared to the previous set, causing the server to return the error “Wrong arguments to mysql\_stmt\_execute()”.
- Handle case when string representation of timestamp contains trailing ‘.’ with no numbers following it.
- Inefficient detection of pre-existing string instances in `ResultSet.getNativeString()`. ([Bug#5706](http://bugs.mysql.com/5706) [<http://bugs.mysql.com/5706>])
- Don't throw exceptions for `Connection.releaseSavepoint()`.
- Use a per-session `Calendar` instance by default when decoding dates from `ServerPreparedStatements` (set to old, less performant behavior by setting property `dynamicCalendars=true`).
- Added experimental configuration property `dontUnpackBinaryResults`, which delays unpacking binary result set values until they're asked for, and only creates object instances for non-numerical values (it is set to `false` by default). For some usecase/jvm combinations, this is friendlier on the garbage collector.
- `UNSIGNED BIGINT` unpacked incorrectly from server-side prepared statement result sets. ([Bug#5729](http://bugs.mysql.com/5729) [<http://bugs.mysql.com/5729>])
- `ServerSidePreparedStatement` allocating short-lived objects unnecessarily. ([Bug#6225](http://bugs.mysql.com/6225) [<http://bugs.mysql.com/6225>])
- Removed unwanted new `Throwable()` in `ResultSet` constructor due to bad merge (caused a new object instance that was never used for every result set created). Found while profiling for [Bug#6359](http://bugs.mysql.com/6359) [<http://bugs.mysql.com/6359>].
- Fixed too-early creation of `StringBuffer` in `EscapeProcessor.escapeSQL()`, also return `String` when escaping not needed (to avoid unnecessary object allocations). Found while profiling for [Bug#6359](http://bugs.mysql.com/6359) [<http://bugs.mysql.com/6359>].
- Use null-safe-equals for key comparisons in updatable result sets.
- `SUM()` on `DECIMAL` with server-side prepared statement ignores scale if zero-padding is needed (this ends up being due to conversion to `DOUBLE` by server, which when converted to a string to

- parse into `BigDecimal`, loses all “padding” zeros). ([Bug#6537](http://bugs.mysql.com/6537) [<http://bugs.mysql.com/6537>])
- Use `DatabaseMetaData.getIdentiferQuoteString()` when building DBMD queries.
- Use 1MB packet for sending file for `LOAD DATA LOCAL INFILE` if that is < `max_allowed_packet` on server.
- `ResultSetMetaData.getColumnDisplaySize()` returns incorrect values for multi-byte charsets. ([Bug#6399](http://bugs.mysql.com/6399) [<http://bugs.mysql.com/6399>])
- Make auto-deserialization of `java.lang.Objects` stored in `BLOB` columns configurable via `autoDeserialize` property (defaults to `false`).
- Re-work `Field.isOpaqueBinary()` to detect `CHAR(n) CHARACTER SET BINARY` to support fixed-length binary fields for `ResultSet.getObject()`.
- Use our own implementation of buffered input streams to get around blocking behavior of `java.io.BufferedReader`. Disable this with `useReadAheadInput=false`.
- Failing to connect to the server when one of the addresses for the given host name is IPV6 (which the server does not yet bind on). The driver now loops through *all* IP addresses for a given host, and stops on the first one that `accepts()` a `socket.connect()`. ([Bug#6348](http://bugs.mysql.com/6348) [<http://bugs.mysql.com/6348>])

## D.8.15. Changes in MySQL Connector/J 3.1.4-beta (04 September 2004)

- Connector/J 3.1.3 beta does not handle integers correctly (caused by changes to support unsigned reads in `Buffer.readInt()` -> `Buffer.readShort()`). ([Bug#4510](http://bugs.mysql.com/4510) [<http://bugs.mysql.com/4510>])
- Added support in `DatabaseMetaData.getTables()` and `getTableTypes()` for views, which are now available in MySQL server 5.0.x.
- `ServerPreparedStatement.execute*()` sometimes threw `ArrayIndexOutOfBoundsException` when unpacking field metadata. ([Bug#4642](http://bugs.mysql.com/4642) [<http://bugs.mysql.com/4642>])
- Optimized integer number parsing, enable “old” slower integer parsing using JDK classes via `useFastIntParsing=false` property.
- Added `useOnlyServerErrorMessages` property, which causes message text in exceptions generated by the server to only contain the text sent by the server (as opposed to the `SQLState`'s “standard” description, followed by the server's error message). This property is set to `true` by default.
- `ResultSet.isNull()` does not work for primitives if a previous `null` was returned. ([Bug#4689](http://bugs.mysql.com/4689) [<http://bugs.mysql.com/4689>])
- Track packet sequence numbers if `enablePacketDebug=true`, and throw an exception if packets received out-of-order.
- `ResultSet.getObject()` returns wrong type for strings when using prepared statements. ([Bug#4482](http://bugs.mysql.com/4482) [<http://bugs.mysql.com/4482>])
- Calling `MysqlPooledConnection.close()` twice (even though an application error), caused NPE. Fixed.

- `ServerPreparedStatement`s dealing with return of `DECIMAL` type don't work. (Bug#5012 [<http://bugs.mysql.com/5012>])
- `ResultSet.getObject()` doesn't return type `Boolean` for pseudo-bit types from prepared statements on 4.1.x (shortcut for avoiding extra type conversion when using binary-encoded result sets obscured test in `getObject()` for "pseudo" bit type). (Bug#5032 [<http://bugs.mysql.com/5032>])
- You can now use URLs in `LOAD DATA LOCAL INFILE` statements, and the driver will use Java's built-in handlers for retrieving the data and sending it to the server. This feature is not enabled by default, you must set the `allowUrlInLocalInfile` connection property to `true`.
- The driver is more strict about truncation of numerics on `ResultSet.get*()`, and will throw an `SQLException` when truncation is detected. You can disable this by setting `jdbcCompliant-Truncation` to `false` (it is enabled by default, as this functionality is required for JDBC compliance).
- Added three ways to deal with all-zero datetimes when reading them from a `ResultSet: exception` (the default), which throws an `SQLException` with an `SQLState` of `S1009`; `convertToNull`, which returns `NULL` instead of the date; and `round`, which rounds the date to the nearest closest value which is `'0001-01-01'`.
- Fixed `ServerPreparedStatement` to read prepared statement metadata off the wire, even though it's currently a placeholder instead of using `MysqlIO.clearInputStream()` which didn't work at various times because data wasn't available to read from the server yet. This fixes sporadic errors users were having with `ServerPreparedStatements` throwing `ArrayIndexOutOfBoundsException`s.
- Use `com.mysql.jdbc.Message`'s classloader when loading resource bundle, should fix sporadic issues when the caller's classloader can't locate the resource bundle.

## D.8.16. Changes in MySQL Connector/J 3.1.3-beta (07 July 2004)

- Mangle output parameter names for `CallableStatements` so they will not clash with user variable names.
- Added support for `INOUT` parameters in `CallableStatements`.
- Null bitmask sent for server-side prepared statements was incorrect. (Bug#4119 [<http://bugs.mysql.com/4119>])
- Use SQL Standard SQL states by default, unless `useSqlStateCodes` property is set to `false`.
- Added packet debugging code (see the `enablePacketDebug` property documentation).
- Added constants for MySQL error numbers (publicly accessible, see `com.mysql.jdbc.MysqlErrorNumbers`), and the ability to generate the mappings of vendor error codes to `SQLStates` that the driver uses (for documentation purposes).
- Externalized more messages (on-going effort).
- Error in retrieval of `mediumint` column with prepared statements and binary protocol. (Bug#4311 [<http://bugs.mysql.com/4311>])
- Support new time zone variables in MySQL-4.1.3 when `useTimezone=true`.

- Support for unsigned numerics as return types from prepared statements. This also causes a change in `ResultSet.getObject()` for the `bigint unsigned` type, which used to return `BigDecimal` instances, it now returns instances of `java.lang.BigInteger`.

## D.8.17. Changes in MySQL Connector/J 3.1.2-alpha (09 June 2004)

- Fixed stored procedure parameter parsing info when size was specified for a parameter (for example, `char()`, `varchar()`).
- Enabled callable statement caching via `cacheCallableStmts` property.
- Fixed case when no output parameters specified for a stored procedure caused a bogus query to be issued to retrieve out parameters, leading to a syntax error from the server.
- Fixed case when no parameters could cause a `NullPointerException` in `CallableStatement.setOutputParameters()`.
- Removed wrapping of exceptions in `MysqlIO.changeUser()`.
- Fixed sending of split packets for large queries, enabled nio ability to send large packets as well.
- Added `.toString()` functionality to `ServerPreparedStatement`, which should help if you're trying to debug a query that is a prepared statement (it shows SQL as the server would process).
- Added `gatherPerformanceMetrics` property, along with properties to control when/where this info gets logged (see docs for more info).
- `ServerPreparedStatements` weren't actually de-allocating server-side resources when `.close()` was called.
- Added `logSlowQueries` property, along with `slowQueriesThresholdMillis` property to control when a query should be considered "slow."
- Correctly map output parameters to position given in `prepareCall()` versus. order implied during `registerOutParameter()`. ([Bug#3146](http://bugs.mysql.com/3146) [<http://bugs.mysql.com/3146>])
- Correctly detect initial character set for servers  $\geq 4.1.0$ .
- Cleaned up detection of server properties.
- Support placeholder for parameter metadata for server  $\geq 4.1.2$ .
- `getProcedures()` does not return any procedures in result set. ([Bug#3539](http://bugs.mysql.com/3539) [<http://bugs.mysql.com/3539>])
- `getProcedureColumns()` doesn't work with wildcards for procedure name. ([Bug#3540](http://bugs.mysql.com/3540) [<http://bugs.mysql.com/3540>])
- `DBMD.getSQLStateType()` returns incorrect value. ([Bug#3520](http://bugs.mysql.com/3520) [<http://bugs.mysql.com/3520>])
- Added `connectionCollation` property to cause driver to issue `set collation_connection=...` query on connection init if default collation for given charset is not appropriate.

- Fixed `DatabaseMetaData.getProcedures()` when run on MySQL-5.0.0 (output of `SHOW PROCEDURE STATUS` changed between 5.0.0 and 5.0.1).
- `getWarnings()` returns `SQLWarning` instead of `DataTruncation`. (Bug#3804 [<http://bugs.mysql.com/3804>])
- Don't enable server-side prepared statements for server version 5.0.0 or 5.0.1, as they aren't compatible with the '4.1.2+' style that the driver uses (the driver expects information to come back that isn't there, so it hangs).

## D.8.18. Changes in MySQL Connector/J 3.1.1-alpha (14 February 2004)

- Fixed bug with `UpdatableResultSets` not using client-side prepared statements.
- Fixed character encoding issues when converting bytes to ASCII when MySQL doesn't provide the character set, and the JVM is set to a multi-byte encoding (usually affecting retrieval of numeric values).
- Unpack “unknown” data types from server prepared statements as `Strings`.
- Implemented long data (Blobs, Clobs, InputStreams, Readers) for server prepared statements.
- Implemented `Statement.getWarnings()` for MySQL-4.1 and newer (using `SHOW WARNINGS`).
- Default result set type changed to `TYPE_FORWARD_ONLY` (JDBC compliance).
- Centralized setting of result set type and concurrency.
- Refactored how connection properties are set and exposed as `DriverPropertyInfo` as well as `Connection` and `DataSource` properties.
- Support for NIO. Use `useNIO=true` on platforms that support NIO.
- Support for transaction savepoints (MySQL  $\geq$  4.0.14 or 4.1.1).
- Support for `mysql_change_user()`. See the `changeUser()` method in `com.mysql.jdbc.Connection`.
- Reduced number of methods called in average query to be more efficient.
- Prepared `Statements` will be re-prepared on auto-reconnect. Any errors encountered are postponed until first attempt to re-execute the re-prepared statement.
- Ensure that warnings are cleared before executing queries on prepared statements, as-per JDBC spec (now that we support warnings).
- Support “old” `profileSql` capitalization in `ConnectionProperties`. This property is deprecated, you should use `profileSQL` if possible.
- Optimized `Buffer.readLenByteArray()` to return shared empty byte array when length is 0.
- Allow contents of `PreparedStatement.setBlob()` to be retained between calls to `.execute*()`.



- Deal with 0-length tokens in `EscapeProcessor` (caused by callable statement escape syntax).
- Check for closed connection on delete/update/insert row operations in `UpdatableResultSet`.
- Fix support for table aliases when checking for all primary keys in `UpdatableResultSet`.
- Removed `useFastDates` connection property.
- Correctly initialize datasource properties from JNDI Refs, including explicitly specified URLs.
- `DatabaseMetaData` now reports `supportsStoredProcedures()` for MySQL versions  $\geq 5.0.0$
- Fixed stack overflow in `Connection.prepareCall()` (bad merge).
- Fixed `IllegalAccessError` to `Calendar.getTimeInMillis()` in `DateTimeValue` (for JDK  $< 1.4$ ).
- `DatabaseMetaData.getColumns()` is not returning correct column ordinal info for non-'%' column name patterns. (Bug#1673 [http://bugs.mysql.com/1673])
- Merged fix of datatype mapping from MySQL type `FLOAT` to `java.sql.Types.REAL` from 3.0 branch.
- Detect collation of column for `RSMD.isCaseSensitive()`.
- Fixed sending of queries larger than 16M.
- Added named and indexed input/output parameter support to `CallableStatement`. MySQL-5.0.x or newer.
- Fixed `NullPointerException` in `ServerPreparedStatement.setTimestamp()`, as well as year and month discrepancies in `ServerPreparedStatement.setTimestamp()`, `setDate()`.
- Added ability to have multiple database/JVM targets for compliance and regression/unit tests in `build.xml`.
- Fixed NPE and year/month bad conversions when accessing some datetime functionality in `ServerPreparedStatements` and their resultant result sets.
- Display where/why a connection was implicitly closed (to aid debugging).
- `CommunicationsException` implemented, that tries to determine why communications was lost with a server, and displays possible reasons when `.getMessage()` is called.
- `NULL` values for numeric types in binary encoded result sets causing `NullPointerExceptions`. (Bug#2359 [http://bugs.mysql.com/2359])
- Implemented `Connection.prepareCall()`, and `DatabaseMetaData.getProcedures()` and `getProcedureColumns()`.
- Reset `long binary` parameters in `ServerPreparedStatement` when `clearParameters()` is called, by sending `COM_RESET_STMT` to the server.
- Merged prepared statement caching, and `.getMetaData()` support from 3.0 branch.
- Fixed off-by-1900 error in some cases for years in `TimeUtil.fastDate/TimeCreate()` when unpacking results from server-side prepared statements.

- Fixed charset conversion issue in `getTables()`. ([Bug#2502](http://bugs.mysql.com/2502) [<http://bugs.mysql.com/2502>])
- Implemented multiple result sets returned from a statement or stored procedure.
- Server-side prepared statements were not returning datatype `YEAR` correctly. ([Bug#2606](http://bugs.mysql.com/2606) [<http://bugs.mysql.com/2606>])
- Enabled streaming of result sets from server-side prepared statements.
- Class-cast exception when using scrolling result sets and server-side prepared statements. ([Bug#2623](http://bugs.mysql.com/2623) [<http://bugs.mysql.com/2623>])
- Merged unbuffered input code from 3.0.
- Fixed `ConnectionProperties` that weren't properly exposed via accessors, cleaned up `ConnectionProperties` code.
- `NULL` fields were not being encoded correctly in all cases in server-side prepared statements. ([Bug#2671](http://bugs.mysql.com/2671) [<http://bugs.mysql.com/2671>])
- Fixed rare buffer underflow when writing numbers into buffers for sending prepared statement execution requests.
- Use DocBook version of docs for shipped versions of drivers.

### D.8.19. Changes in MySQL Connector/J 3.1.0-alpha (18 February 2003)

- Added `requireSSL` property.
- Added `useServerPrepStmts` property (default `false`). The driver will use server-side prepared statements when the server version supports them (4.1 and newer) when this property is set to `true`. It is currently set to `false` by default until all bind/fetch functionality has been implemented. Currently only DML prepared statements are implemented for 4.1 server-side prepared statements.
- Track open `Statements`, close all when `Connection.close()` is called (JDBC compliance).

### D.8.20. Changes in MySQL Connector/J 3.0.17-ga (23 June 2005)

- `Timestamp/Time` conversion goes in the wrong “direction” when `useTimeZone=true` and server time zone differs from client time zone. ([Bug#5874](http://bugs.mysql.com/5874) [<http://bugs.mysql.com/5874>])
- `DatabaseMetaData.getIndexInfo()` ignored `unique` parameter. ([Bug#7081](http://bugs.mysql.com/7081) [<http://bugs.mysql.com/7081>])
- Support new protocol type `MYSQL_TYPE_VARCHAR`.
- Added `useOldUTF8Behavior` configuration property, which causes JDBC driver to act like it did with MySQL-4.0.x and earlier when the character encoding is `utf-8` when connected to MySQL-4.1 or newer.

- Statements created from a pooled connection were returning physical connection instead of logical connection when `getConnection()` was called. (Bug#7316 [<http://bugs.mysql.com/7316>])
- `PreparedStatement` don't encode Big5 (and other multi-byte) character sets correctly in static SQL strings. (Bug#7033 [<http://bugs.mysql.com/7033>])
- Connections starting up failed-over (due to down master) never retry master. (Bug#6966 [<http://bugs.mysql.com/6966>])
- `PreparedStatement.fixDecimalExponent()` adding extra +, making number unparseable by MySQL server. (Bug#7061 [<http://bugs.mysql.com/7061>])
- Timestamp key column data needed `_binary` stripped for `UpdatableResultSet.refreshRow()`. (Bug#7686 [<http://bugs.mysql.com/7686>])
- Backported SQLState codes mapping from Connector/J 3.1, enable with `useSqlStateCodes=true` as a connection property, it defaults to `false` in this release, so that we don't break legacy applications (it defaults to `true` starting with Connector/J 3.1).
- `PreparedStatement.fixDecimalExponent()` adding extra +, making number unparseable by MySQL server. (Bug#7601 [<http://bugs.mysql.com/7601>])
- Escape sequence `{fn convert(..., type)}` now supports ODBC-style types that are prepended by `SQL_`.
- Fixed duplicated code in `configureClientCharset()` that prevented `useOldUTF8Behavior=true` from working properly.
- Handle streaming result sets with more than 2 billion rows properly by fixing wraparound of row number counter.
- `MS932`, `SHIFT_JIS`, and `Windows_31J` not recognized as aliases for `sjis`. (Bug#7607 [<http://bugs.mysql.com/7607>])
- Adding `CP943` to aliases for `sjis`. (Bug#6549 [<http://bugs.mysql.com/6549>], fixed while fixing Bug#7607 [<http://bugs.mysql.com/7607>])
- Which requires hex escaping of binary data when using multi-byte charsets with prepared statements. (Bug#8064 [<http://bugs.mysql.com/8064>])
- `NON_UNIQUE` column from `DBMD.getIndexInfo()` returned inverted value. (Bug#8812 [<http://bugs.mysql.com/8812>])
- Workaround for server Bug#9098 [<http://bugs.mysql.com/9098>]: Default values of `CURRENT_*` for `DATE`, `TIME`, `DATETIME`, and `TIMESTAMP` columns can't be distinguished from `string` values, so `UpdatableResultSet.moveToInsertRow()` generates bad SQL for inserting default values.
- `EUCKR` charset is sent as `SET NAMES euc_kr` which MySQL-4.1 and newer doesn't understand. (Bug#8629 [<http://bugs.mysql.com/8629>])
- `DatabaseMetaData.supportsSelectForUpdate()` returns correct value based on server version.
- Use hex escapes for `PreparedStatement.setBytes()` for double-byte charsets including "aliases" `Windows-31J`, `CP934`, `MS932`.
- Added support for the `EUC_JP_Solaris` character encoding, which maps to a MySQL encoding of `eucjpms` (backported from 3.1 branch). This only works on servers that support `eucjpms`,

namely 5.0.3 or later.

## D.8.21. Changes in MySQL Connector/J 3.0.16-ga (15 November 2004)

- Re-issue character set configuration commands when re-using pooled connections and/or `Connection.changeUser()` when connected to MySQL-4.1 or newer.
- Fixed `ResultSetMetaData.isReadOnly()` to detect non-writable columns when connected to MySQL-4.1 or newer, based on existence of “original” table and column names.
- `ResultSet.updateByte()` when on insert row throws `ArrayOutOfBoundsException`. (Bug#5664 [<http://bugs.mysql.com/5664>])
- Fixed `DatabaseMetaData.getTypes()` returning incorrect (this is, non-negative) scale for the `NUMERIC` type.
- Off-by-one bug in `Buffer.readString(string)`. (Bug#5664 [<http://bugs.mysql.com/5664>])
- Made `TINYINT(1)` -> `BIT/Boolean` conversion configurable via `tinyIntIsBit` property (default `true` to be JDBC compliant out of the box).
- Only set `character_set_results` during connection establishment if server version  $\geq$  4.1.1.
- Fixed regression where `useUnbufferedInput` was defaulting to `false`.
- `ResultSet.getTimestamp()` on a column with `TIME` in it fails. (Bug#5664 [<http://bugs.mysql.com/5664>])

## D.8.22. Changes in MySQL Connector/J 3.0.15-production (04 September 2004)

- `StringUtils.escapeEasternUnicodeByteStream` was still broken for GBK. (Bug#4010 [<http://bugs.mysql.com/4010>])
- Failover for `autoReconnect` not using port numbers for any hosts, and not retrying all hosts. (**Warning:** This required a change to the `SocketFactory.connect()` method signature, which is now `public Socket connect(String host, int portNumber, Properties props)`; therefore, any third-party socket factories will have to be changed to support this signature. (Bug#4334 [<http://bugs.mysql.com/4334>])
- Logical connections created by `MysqlConnectionPoolDataSource` will now issue a `rollback()` when they are closed and sent back to the pool. If your application server/connection pool already does this for you, you can set the `rollbackOnPooledClose` property to `false` to avoid the overhead of an extra `rollback()`.
- Removed redundant calls to `checkRowPos()` in `ResultSet`.
- `DOUBLE` mapped twice in `DBMD.getTypeInfo()`. (Bug#4742 [<http://bugs.mysql.com/4742>])
- Added FLOSS license exemption.
- Calling `.close()` twice on a `PooledConnection` causes NPE. (Bug#4808)

[<http://bugs.mysql.com/4808>])

- `DBMD.getColumns()` returns incorrect JDBC type for unsigned columns. This affects type mappings for all numeric types in the `RSMD.getColumnType()` and `RSMD.getColumnTypeNames()` methods as well, to ensure that “like” types from `DBMD.getColumns()` match up with what `RSMD.getColumnType()` and `getColumnTypeNames()` return. (Bug#4138 [<http://bugs.mysql.com/4138>], Bug#4860 [<http://bugs.mysql.com/4860>])
- “Production” is now “GA” (General Availability) in naming scheme of distributions.
- `RSMD.getPrecision()` returning 0 for non-numeric types (should return max length in chars for non-binary types, max length in bytes for binary types). This fix also fixes mapping of `RSMD.getColumnType()` and `RSMD.getColumnTypeName()` for the BLOB types based on the length sent from the server (the server doesn’t distinguish between TINYBLOB, BLOB, MEDIUMBLOB or LONGBLOB at the network protocol level). (Bug#4880 [<http://bugs.mysql.com/4880>])
- `ResultSet` should release `Field[]` instance in `.close()`. (Bug#5022 [<http://bugs.mysql.com/5022>])
- `ResultSet.getMetaData()` should not return incorrectly initialized metadata if the result set has been closed, but should instead throw an `SQLException`. Also fixed for `getRow()` and `getWarnings()` and traversal methods by calling `checkClosed()` before operating on instance-level fields that are nullified during `.close()`. (Bug#5069 [<http://bugs.mysql.com/5069>])
- Parse new time zone variables from 4.1.x servers.
- Use `_binary` introducer for `PreparedStatement.setBytes()` and `set*Stream()` when connected to MySQL-4.1.x or newer to avoid misinterpretation during character conversion.

### D.8.23. Changes in MySQL Connector/J 3.0.14-production (28 May 2004)

- Fixed URL parsing error.

### D.8.24. Changes in MySQL Connector/J 3.0.13-production (27 May 2004)

- Using a `MySQLDataSource` without server name fails. (Bug#3848 [<http://bugs.mysql.com/3848>])
- `No Database Selected` when using `MysqlConnectionPoolDataSource`. (Bug#3920 [<http://bugs.mysql.com/3920>])
- `PreparedStatement.getGeneratedKeys()` method returns only 1 result for batched insertions. (Bug#3873 [<http://bugs.mysql.com/3873>])

### D.8.25. Changes in MySQL Connector/J 3.0.12-production (18 May 2004)

- Add unsigned attribute to `DatabaseMetaData.getColumns()` output in the `TYPE_NAME`

column.

- Added `failOverReadOnly` property, to allow end-user to configure state of connection (read-only/writable) when failed over.
- Backported “change user” and “reset server state” functionality from 3.1 branch, to allow clients of `MysqlConnectionPoolDataSource` to reset server state on `getConnection()` on a pooled connection.
- Don't escape SJIS/GBK/BIG5 when using MySQL-4.1 or newer.
- Allow `url` parameter for `MysqlDataSource` and `MysqlConnectionPool DataSource` so that passing of other properties is possible from inside appservers.
- Map duplicate key and foreign key errors to `SQLState` of `23000`.
- Backport documentation tooling from 3.1 branch.
- Return creating statement for `ResultSets` created by `getGeneratedKeys()`. (Bug#2957 [http://bugs.mysql.com/2957])
- Allow `java.util.Date` to be sent in as parameter to `PreparedStatement.setObject()`, converting it to a `Timestamp` to maintain full precision. (Bug#103 [http://bugs.mysql.com/103]).
- Don't truncate `BLOB` or `CLOB` values when using `setBytes()` and/or `setBinary/CharacterStream()`. (Bug#2670 [http://bugs.mysql.com/2670]).
- Dynamically configure character set mappings for field-level character sets on MySQL-4.1.0 and newer using `SHOW COLLATION` when connecting.
- Map `binary` character set to `US-ASCII` to support `DATETIME` charset recognition for servers  $\geq$  4.1.2.
- Use `SET character_set_results` during initialization to allow any charset to be returned to the driver for result sets.
- Use `charsetnr` returned during connect to encode queries before issuing `SET NAMES` on MySQL  $\geq$  4.1.0.
- Add helper methods to `ResultSetMetaData` (`getColumnCharacterEncoding()` and `getColumnCharacterSet()`) to allow end-users to see what charset the driver thinks it should be using for the column.
- Only set `character_set_results` for MySQL  $\geq$  4.1.0.
- `StringUtils.escapeSJISByteStream()` not covering all eastern double-byte charsets correctly. (Bug#3511 [http://bugs.mysql.com/3511])
- Renamed `StringUtils.escapeSJISByteStream()` to more appropriate `escapeEasternUnicodeByteStream()`.
- Not specifying database in URL caused `MalformedURLException` exception. (Bug#3554 [http://bugs.mysql.com/3554])
- Auto-convert MySQL encoding names to Java encoding names if used for `characterEncoding` property.
- Added encoding names that are recognized on some JVMs to fix case where they were reverse-

mapped to MySQL encoding names incorrectly.

- Use `junit.textui.TestRunner` for all unit tests (to allow them to be run from the command line outside of Ant or Eclipse).
- `UpdatableResultSet` not picking up default values for `moveToInsertRow()`. (Bug#3557 [http://bugs.mysql.com/3557])
- Inconsistent reporting of data type. The server still doesn't return all types for \*BLOBs \*TEXT correctly, so the driver won't return those correctly. (Bug#3570 [http://bugs.mysql.com/3570])
- `DBMD.getSQLStateType()` returns incorrect value. (Bug#3520 [http://bugs.mysql.com/3520])
- Fixed regression in `PreparedStatement.setString()` and eastern character encodings.
- Made `StringRegressionTest` 4.1-unicode aware.

## D.8.26. Changes in MySQL Connector/J 3.0.11-stable (19 February 2004)

- Trigger a `SET NAMES utf8` when encoding is forced to `utf8` or `utf-8` via the `characterEncoding` property. Previously, only the Java-style encoding name of `utf-8` would trigger this.
- `AutoReconnect` time was growing faster than exponentially. (Bug#2447 [http://bugs.mysql.com/2447])
- Fixed failover always going to last host in list. (Bug#2578 [http://bugs.mysql.com/2578])
- Added `useUnbufferedInput` parameter, and now use it by default (due to JVM issue <http://developer.java.sun.com/developer/bugParade/bugs/4401235.html>)
- Detect `on/off` or `1, 2, 3` form of `lower_case_table_names` value on server.
- Return `java.lang.Integer` for `TINYINT` and `SMALLINT` types from `ResultSetMetaData.getColumnClassName()`. (Bug#2852 [http://bugs.mysql.com/2852])
- Return `java.lang.Double` for `FLOAT` type from `ResultSetMetaData.getColumnClassName()`. (Bug#2855 [http://bugs.mysql.com/2855])
- Return `[B` instead of `java.lang.Object` for `BINARY`, `VARBINARY` and `LONGVARBINARY` types from `ResultSetMetaData.getColumnClassName()` (JDBC compliance).
- Issue connection events on all instances created from a `ConnectionPoolDataSource`.

## D.8.27. Changes in MySQL Connector/J 3.0.10-stable (13 January 2004)

- Don't count quoted IDs when inside a 'string' in `PreparedStatement` parsing. (Bug#1511 [http://bugs.mysql.com/1511])
- “Friendlier” exception message for `PacketTooLargeException`. (Bug#1534 [http://bugs.mysql.com/1534])

- Backported fix for aliased tables and `UpdatableResultSets` in `checkUpdatability()` method from 3.1 branch.
- Fix for `ArrayIndexOutOfBoundsException` exception when using `Statement.setMaxRows()`. (Bug#1695 [http://bugs.mysql.com/1695])
- Barge blobs and split packets not being read correctly. (Bug#1576 [http://bugs.mysql.com/1576])
- Fixed regression of `Statement.getGeneratedKeys()` and `REPLACE` statements.
- Subsequent call to `ResultSet.updateFoo()` causes NPE if result set is not updatable. (Bug#1630 [http://bugs.mysql.com/1630])
- Fix for 4.1.1-style authentication with no password.
- Foreign Keys column sequence is not consistent in `DatabaseMetaData.getImported/Exported/CrossReference()`. (Bug#1731 [http://bugs.mysql.com/1731])
- `DatabaseMetaData.getSystemFunction()` returning bad function `VResultsSion`. (Bug#1775 [http://bugs.mysql.com/1775])
- Cross-database updatable result sets are not checked for updatability correctly. (Bug#1592 [http://bugs.mysql.com/1592])
- `DatabaseMetaData.getColumns()` should return `Types.LONGVARCHAR` for MySQL `LONGTEXT` type.
- `ResultSet.getObject()` on `TINYINT` and `SMALLINT` columns should return Java type `Integer`. (Bug#1913 [http://bugs.mysql.com/1913])
- Added `alwaysClearStream` connection property, which causes the driver to always empty any remaining data on the input stream before each query.
- Added more descriptive error message `Server Configuration Denies Access to DataSource`, as well as retrieval of message from server.
- Autoreconnect code didn't set catalog upon reconnect if it had been changed.
- Implement `ResultSet.updateClob()`.
- `ResultSetMetaData.isCaseSensitive()` returned wrong value for `CHAR/VARCHAR` columns.
- Connection property `maxRows` not honored. (Bug#1933 [http://bugs.mysql.com/1933])
- Statements being created too many times in `DBMD.extractForeignKeyFromCreateTable()`. (Bug#1925 [http://bugs.mysql.com/1925])
- Support escape sequence `{fn convert ... }`. (Bug#1914 [http://bugs.mysql.com/1914])
- `ArrayIndexOutOfBoundsException` when parameter number == number of parameters + 1. (Bug#1958 [http://bugs.mysql.com/1958])
- `ResultSet.findColumn()` should use first matching column name when there are duplicate column names in `SELECT` query (JDBC-compliance). (Bug#2006 [http://bugs.mysql.com/2006])
- Removed static synchronization bottleneck from `PreparedStatement.setTimestamp()`.



- Removed static synchronization bottleneck from instance factory method of `SingleByteCharsetConverter`.
- Enable caching of the parsing stage of prepared statements via the `cachePrepStmts`, `prepStmtCacheSize`, and `prepStmtCacheSqlLimit` properties (disabled by default).
- Speed up parsing of `PreparedStatement`s, try to use one-pass whenever possible.
- Fixed security exception when used in Applets (applets can't read the system property `file.encoding` which is needed for `LOAD DATA LOCAL INFILE`).
- Use constants for `SQLStates`.
- Map charset `ko18_ru` to `ko18r` when connected to MySQL-4.1.0 or newer.
- Ensure that `Buffer.writeString()` saves room for the `\0`.
- Fixed exception `Unknown character set 'danish'` on connect with JDK-1.4.0
- Fixed mappings in `SQLException` to report deadlocks with `SQLStates` of `41000`.
- `maxRows` property would affect internal statements, so check it for all statement creation internal to the driver, and set to 0 when it is not.

## D.8.28. Changes in MySQL Connector/J 3.0.9-stable (07 October 2003)

- Faster date handling code in `ResultSet` and `PreparedStatement` (no longer uses `Date` methods that synchronize on static calendars).
- Fixed test for end of buffer in `Buffer.readString()`.
- Fixed `ResultSet.previous()` behavior to move current position to before result set when on first row of result set. ([Bug#496](http://bugs.mysql.com/496) [<http://bugs.mysql.com/496>])
- Fixed `Statement` and `PreparedStatement` issuing bogus queries when `setMaxRows()` had been used and a `LIMIT` clause was present in the query.
- `refreshRow` didn't work when primary key values contained values that needed to be escaped (they ended up being doubly escaped). ([Bug#661](http://bugs.mysql.com/661) [<http://bugs.mysql.com/661>])
- Support `InnoDB` constraint names when extracting foreign key information in `DatabaseMetaData` (implementing ideas from Parwinder Sekhon). ([Bug#517](http://bugs.mysql.com/517) [<http://bugs.mysql.com/517>], [Bug#664](http://bugs.mysql.com/664) [<http://bugs.mysql.com/664>])
- Backported 4.1 protocol changes from 3.1 branch (server-side SQL states, new field information, larger client capability flags, connect-with-database, and so forth).
- Fix `UpdatableResultSet` to return values for `getXXX()` when on insert row. ([Bug#675](http://bugs.mysql.com/675) [<http://bugs.mysql.com/675>])
- The `insertRow` in an `UpdatableResultSet` is now loaded with the default column values when `moveToInsertRow()` is called. ([Bug#688](http://bugs.mysql.com/688) [<http://bugs.mysql.com/688>])
- `DatabaseMetaData.getColumns()` wasn't returning `NULL` for default values that are specified as `NULL`.

- Change default statement type/concurrency to `TYPE_FORWARD_ONLY` and `CONCUR_READ_ONLY` (spec compliance).
- Don't try and reset isolation level on reconnect if MySQL doesn't support them.
- Don't wrap `SQLExceptions` in `RowDataDynamic`.
- Don't change timestamp TZ twice if `useTimezone==true`. (Bug#774 [http://bugs.mysql.com/774])
- Fixed regression in large split-packet handling. (Bug#848 [http://bugs.mysql.com/848])
- Better diagnostic error messages in exceptions for “streaming” result sets.
- Issue exception on `ResultSet.getXXX()` on empty result set (wasn't caught in some cases).
- Don't hide messages from exceptions thrown in I/O layers.
- Don't fire connection closed events when closing pooled connections, or on `PooledConnection.getConnection()` with already open connections. (Bug#884 [http://bugs.mysql.com/884])
- Clip +/- INF (to smallest and largest representative values for the type in MySQL) and NaN (to 0) for `setDouble/setFloat()`, and issue a warning on the statement when the server does not support +/- INF or NaN.
- Double-escaping of `'\'` when charset is SJIS or GBK and `'\'` appears in non-escaped input. (Bug#879 [http://bugs.mysql.com/879])
- When emptying input stream of unused rows for “streaming” result sets, have the current thread `yield()` every 100 rows in order to not monopolize CPU time.
- `DatabaseMetaData.getColumns()` getting confused about the keyword “set” in character columns. (Bug#1099 [http://bugs.mysql.com/1099])
- Fixed deadlock issue with `Statement.setMaxRows()`.
- Fixed `CLOB.truncate()`. (Bug#1130 [http://bugs.mysql.com/1130])
- Optimized `CLOB.setCharacterStream()`. (Bug#1131 [http://bugs.mysql.com/1131])
- Made `databaseName`, `portNumber`, and `serverName` optional parameters for `MySQLDataSourceFactory`. (Bug#1246 [http://bugs.mysql.com/1246])
- `ResultSet.get/setString` mashing char 127. (Bug#1247 [http://bugs.mysql.com/1247])
- Backported authentication changes for 4.1.1 and newer from 3.1 branch.
- Added `com.mysql.jdbc.util.BaseBugReport` to help creation of testcases for bug reports.
- Added property to “clobber” streaming results, by setting the `clobberStreamingResults` property to `true` (the default is `false`). This will cause a “streaming” `ResultSet` to be automatically closed, and any outstanding data still streaming from the server to be discarded if another query is executed before all the data has been read from the server.

## D.8.29. Changes in MySQL Connector/J 3.0.8-stable (23 May 2003)

- Allow bogus URLs in `Driver.getPropertyInfo()`.
- Return list of generated keys when using multi-value `INSERTS` with `Statement.getGeneratedKeys()`.
- Use JVM charset with filenames and `LOAD DATA [LOCAL] INFILE`.
- Fix infinite loop with `Connection.cleanup()`.
- Changed Ant target `compile-core` to `compile-driver`, and made testsuite compilation a separate target.
- Fixed result set not getting set for `Statement.executeUpdate()`, which affected `getGeneratedKeys()` and `getUpdateCount()` in some cases.
- Unicode character `0xFFFF` in a string would cause the driver to throw an `ArrayOutOfBoundsException`. ([Bug#378](http://bugs.mysql.com/378) [<http://bugs.mysql.com/378>]).
- Return correct number of generated keys when using `REPLACE` statements.
- Fix problem detecting server character set in some cases.
- Fix row data decoding error when using *very* large packets.
- Optimized row data decoding.
- Issue exception when operating on an already closed prepared statement.
- Fixed SJIS encoding bug, thanks to Naoto Sato.
- Optimized usage of `EscapeProcessor`.
- Allow multiple calls to `Statement.close()`.

### D.8.30. Changes in MySQL Connector/J 3.0.7-stable (08 April 2003)

- Fixed `MysqlPooledConnection.close()` calling wrong event type.
- Fixed `StringIndexOutOfBoundsException` in `PreparedStatement.setClob()`.
- 4.1 Column Metadata fixes.
- Remove synchronization from `Driver.connect()` and `Driver.acceptsUrl()`.
- `IOExceptions` during a transaction now cause the `Connection` to be closed.
- Fixed missing conversion for `YEAR` type in `ResultSetMetaData.getColumnTypeName()`.
- Don't pick up indexes that start with `pri` as primary keys for `DBMD.getPrimaryKeys()`.
- Throw `SQLExceptions` when trying to do operations on a forcefully closed `Connection` (that is, when a communication link failure occurs).
- You can now toggle profiling on/off using `Connection.setProfileSql(boolean)`.
- Fixed charset issues with database metadata (charset was not getting set correctly).

- Updatable `ResultSets` can now be created for aliased tables/columns when connected to MySQL-4.1 or newer.
- Fixed `LOAD DATA LOCAL INFILE` bug when file > `max_allowed_packet`.
- Fixed escaping of 0x5c ('\\') character for GBK and Big5 charsets.
- Fixed `ResultSet.getTimestamp()` when underlying field is of type `DATE`.
- Ensure that packet size from `alignPacketSize()` does not exceed `max_allowed_packet` (JVM bug)
- Don't reset `Connection.isReadOnly()` when `autoReconnecting`.

### D.8.31. Changes in MySQL Connector/J 3.0.6-stable (18 February 2003)

- Fixed `ResultSetMetaData` to return "" when catalog not known. Fixes `NullPointerExceptions` with Sun's `CachedRowSet`.
- Fixed `DBMD.getTypeInfo()` and `DBMD.getColumns()` returning different value for precision in `TEXT` and `BLOB` types.
- Allow ignoring of warning for “non transactional tables” during rollback (compliance/usability) by setting `ignoreNonTxTables` property to `true`.
- Fixed `SQLExceptions` getting swallowed on initial connect.
- Fixed `Statement.setMaxRows()` to stop sending `LIMIT` type queries when not needed (performance).
- Clean up `Statement` query/method mismatch tests (that is, `INSERT` not allowed with `.executeQuery()`).
- More checks added in `ResultSet` traversal method to catch when in closed state.
- Fixed `ResultSetMetaData.isWritable()` to return correct value.
- Add “window” of different `NULL` sorting behavior to `DBMD.nullsAreSortedAtStart` (4.0.2 to 4.0.10, true; otherwise, no).
- Implemented `Blob.setBytes()`. You still need to pass the resultant `Blob` back into an updatable `ResultSet` or `PreparedStatement` to persist the changes, because MySQL does not support “locators”.
- Backported 4.1 charset field info changes from Connector/J 3.1.

### D.8.32. Changes in MySQL Connector/J 3.0.5-gamma (22 January 2003)

- Fixed `Buffer.fastSkipLenString()` causing `ArrayIndexOutOfBoundsException` exceptions with some queries when unpacking fields.

- Implemented an empty `TypeMap` for `Connection.getTypeMap()` so that some third-party apps work with MySQL (IBM WebSphere 5.0 Connection pool).
- Added missing `LONGTEXT` type to `DBMD.getColumns()`.
- Retrieve `TX_ISOLATION` from database for `Connection.getTransactionIsolation()` when the MySQL version supports it, instead of an instance variable.
- Quote table names in `DatabaseMetaData.getColumns()`, `getPrimaryKeys()`, `getIndexInfo()`, `getBestRowIdentifier()`.
- Greatly reduce memory required for `setBinaryStream()` in `PreparedStatement`.
- Fixed `ResultSet.isBeforeFirst()` for empty result sets.
- Added update options for foreign key metadata.

### D.8.33. Changes in MySQL Connector/J 3.0.4-gamma (06 January 2003)

- Added quoted identifiers to database names for `Connection.setCatalog`.
- Added support for quoted identifiers in `PreparedStatement` parser.
- Streamlined character conversion and `byte[]` handling in `PreparedStatement` for `setByte()`.
- Reduce memory footprint of `PreparedStatement` by sharing outbound packet with `MysqlIO`.
- Added `strictUpdates` property to allow control of amount of checking for “correctness” of updatable result sets. Set this to `false` if you want faster updatable result sets and you know that you create them from `SELECT` statements on tables with primary keys and that you have selected all primary keys in your query.
- Added support for 4.0.8-style large packets.
- Fixed `PreparedStatement.executeBatch()` parameter overwriting.

### D.8.34. Changes in MySQL Connector/J 3.0.3-dev (17 December 2002)

- Changed `charsToByte` in `SingleByteCharConverter` to be non-static.
- Changed `SingleByteCharConverter` to use lazy initialization of each converter.
- Fixed charset handling in `Fields.java`.
- Implemented `Connection.nativeSQL()`.
- More robust escape tokenizer: Recognize `--` comments, and allow nested escape sequences (see `testsuite.EscapeProcessingTest`).

- `DBMD.getImported/ExportedKeys()` now handles multiple foreign keys per table.
- Fixed `ResultSetMetaData.getPrecision()` returning incorrect values for some floating-point types.
- Fixed `ResultSetMetaData.getColumnTypeName()` returning `BLOB` for `TEXT` and `TEXT` for `BLOB` types.
- Fixed `Buffer.isLastDataPacket()` for 4.1 and newer servers.
- Added `CLIENT_LONG_FLAG` to be able to get more column flags (`isAutoIncrement()` being the most important).
- Because of above, implemented `ResultSetMetaData.isAutoIncrement()` to use `Field.isAutoIncrement()`.
- Honor `lower_case_table_names` when enabled in the server when doing table name comparisons in `DatabaseMetaData` methods.
- Some MySQL-4.1 protocol support (extended field info from selects).
- Use non-aliased table/column names and database names to fully qualify tables and columns in `UpdatableResultSet` (requires MySQL-4.1 or newer).
- Allow user to alter behavior of `Statement/PreparedStatement.executeBatch()` via `continueBatchOnError` property (defaults to `true`).
- Check for connection closed in more `Connection` methods (`createStatement`, `prepareStatement`, `setTransactionIsolation`, `setAutoCommit`).
- More robust implementation of updatable result sets. Checks that *all* primary keys of the table have been selected.
- `LOAD DATA LOCAL INFILE ...` now works, if your server is configured to allow it. Can be turned off with the `allowLoadLocalInfile` property (see the `README`).
- Substitute `' ? '` for unknown character conversions in single-byte character sets instead of `'\0'`.
- `NamedPipeSocketFactory` now works (only intended for Windows), see `README` for instructions.

## D.8.35. Changes in MySQL Connector/J 3.0.2-dev (08 November 2002)

- Fixed issue with updatable result sets and `PreparedStatements` not working.
- Fixed `ResultSet.setFetchDirection(FETCH_UNKNOWN)`.
- Fixed issue when calling `Statement.setFetchSize()` when using arbitrary values.
- Fixed incorrect conversion in `ResultSet.getLong()`.
- Implemented `ResultSet.updateBlob()`.
- Removed duplicate code from `UpdatableResultSet` (it can be inherited from `ResultSet`, the extra code for each method to handle updatability I thought might someday be necessary has not

been needed).

- Fixed `UnsupportedEncodingException` thrown when “forcing” a character encoding via properties.
- Fixed various non-ASCII character encoding issues.
- Added driver property `useHostsInPrivileges`. Defaults to true. Affects whether or not `@hostname` will be used in `DBMD.getColumn/TablePrivileges`.
- All `DBMD` result set columns describing schemas now return `NULL` to be more compliant with the behavior of other JDBC drivers for other database systems (MySQL does not support schemas).
- Added SSL support. See [README](#) for information on how to use it.
- Properly restore connection properties when `autoReconnecting` or failing-over, including `autoCommit` state, and isolation level.
- Use `SHOW CREATE TABLE` when possible for determining foreign key information for `DatabaseMetaData`. Also allows cascade options for `DELETE` information to be returned.
- Escape `0x5c` character in strings for the SJIS charset.
- Fixed start position off-by-1 error in `Clob.getSubString()`.
- Implemented `Clob.truncate()`.
- Implemented `Clob.setString()`.
- Implemented `Clob.setAsciiStream()`.
- Implemented `Clob.setCharacterStream()`.
- Added `com.mysql.jdbc.MinisAdmin` class, which allows you to send `shutdown` command to MySQL server. This is intended to be used when “embedding” Java and MySQL server together in an end-user application.
- Added `connectTimeout` parameter that allows users of JDK-1.4 and newer to specify a maximum time to wait to establish a connection.
- Failover and `autoReconnect` work only when the connection is in an `autoCommit(false)` state, in order to stay transaction-safe.
- Added `queriesBeforeRetryMaster` property that specifies how many queries to issue when failed over before attempting to reconnect to the master (defaults to 50).
- Fixed `DBMD.supportsResultSetConcurrency()` so that it returns true for `ResultSet.TYPE_SCROLL_INSENSITIVE` and `ResultSet.CONCUR_READ_ONLY` or `ResultSet.CONCUR_UPDATABLE`.
- Fixed `ResultSet.isLast()` for empty result sets (should return `false`).
- `PreparedStatement` now honors stream lengths in `setBinary/Ascii/Character Stream()` unless you set the connection property `useStreamLengthsInPrepStmts` to `false`.
- Removed some not-needed temporary object creation by smarter use of `Strings` in `EscapeProcessor`, `Connection` and `DatabaseMetaData` classes.

## D.8.36. Changes in MySQL Connector/J 3.0.1-dev (21 September 2002)

- Fixed `ResultSet.getRow()` off-by-one bug.
- Fixed `RowDataStatic.getAt()` off-by-one bug.
- Added limited `Clob` functionality (`ResultSet.getClob()`, `PreparedStatement.setClob()`, `PreparedStatement.setObject(Clob)`).
- Added `socketTimeout` parameter to URL.
- `Connection.isClosed()` no longer “pings” the server.
- `Connection.close()` issues `rollback()` when `getAutoCommit()` is `false`.
- Added `paranoid` parameter, which sanitizes error messages by removing “sensitive” information from them (such as hostnames, ports, or usernames), as well as clearing “sensitive” data structures when possible.
- Fixed `ResultSetMetaData.isSigned()` for `TINYINT` and `BIGINT`.
- Charsets now automatically detected. Optimized code for single-byte character set conversion.
- Implemented `ResultSet.getCharacterStream()`.
- Added `LOCAL TEMPORARY` to table types in `DatabaseMetaData.getTableTypes()`.
- Massive code clean-up to follow Java coding conventions (the time had come).

## D.8.37. Changes in MySQL Connector/J 3.0.0-dev (31 July 2002)

- **!!! LICENSE CHANGE !!!** The driver is now GPL. If you need non-GPL licenses, please contact me <mark@mysql.com>.
- JDBC-3.0 functionality including `Statement/PreparedStatement.getGeneratedKeys()` and `ResultSet.getURL()`.
- Performance enhancements: Driver is now 50–100% faster in most situations, and creates fewer temporary objects.
- Repackaging: New driver name is `com.mysql.jdbc.Driver`, old name still works, though (the driver is now provided by MySQL-AB).
- Better checking for closed connections in `Statement` and `PreparedStatement`.
- Support for streaming (row-by-row) result sets (see [README](#)) Thanks to Doron.
- Support for large packets (new addition to MySQL-4.0 protocol), see [README](#) for more information.
- JDBC Compliance: Passes all tests besides stored procedure tests.
- Fix and sort primary key names in `DBMetaData` (SF bugs 582086 and 582086).



- Float types now reported as `java.sql.Types.FLOAT` (SF bug 579573).
- `ResultSet.getTimestamp()` now works for `DATE` types (SF bug 559134).
- `ResultSet.getDate/Time/Timestamp` now recognizes all forms of invalid values that have been set to all zeros by MySQL (SF bug 586058).
- Testsuite now uses Junit (which you can get from <http://www.junit.org>).
- The driver now only works with JDK-1.2 or newer.
- Added multi-host failover support (see [README](#)).
- General source-code cleanup.
- Overall speed improvements via controlling transient object creation in `MysqlIO` class when reading packets.
- Performance improvements in string handling and field metadata creation (lazily instantiated) contributed by Alex Twisleton-Wykeham-Fiennes.

### D.8.38. Changes in MySQL Connector/J 2.0.14 (16 May 2002)

- More code cleanup.
- `PreparedStatement` now releases resources on `.close()`. (SF bug 553268)
- Quoted identifiers not used if server version does not support them. Also, if server started with `--ansi` or `--sql-mode=ANSI_QUOTES`, `"` will be used as an identifier quote character, otherwise `'` will be used.
- `ResultSet.getDouble()` now uses code built into JDK to be more precise (but slower).
- `LogicalHandle.isClosed()` calls through to physical connection.
- Added SQL profiling (to `STDERR`). Set `profileSql=true` in your JDBC URL. See [README](#) for more information.
- Fixed typo for `relaxAutoCommit` parameter.

### D.8.39. Changes in MySQL Connector/J 2.0.13 (24 April 2002)

- More code cleanup.
- Fixed unicode chars being read incorrectly. (SF bug 541088)
- Faster blob escaping for `PrepStmt`.
- Added `set/getPortNumber()` to `DataSource(s)`. (SF bug 548167)
- Added `setURL()` to `MySQLXADataSource`. (SF bug 546019)
- `PreparedStatement.toString()` fixed. (SF bug 534026)
- `ResultSetMetaData.getColumnClassName()` now implemented.

- Rudimentary version of `Statement.getGeneratedKeys()` from JDBC-3.0 now implemented (you need to be using JDK-1.4 for this to work, I believe).
- `DBMetaData.getIndexInfo()` - bad PAGES fixed. (SF BUG 542201)

## D.8.40. Changes in MySQL Connector/J 2.0.12 (07 April 2002)

- General code cleanup.
- Added `getIdleFor()` method to `Connection` and `MysqlLogicalHandle`.
- Relaxed synchronization in all classes, should fix 520615 and 520393.
- Added `getTable/columnPrivileges()` to `DBMD` (fixes 484502).
- Added new types to `getTypeInfo()`, fixed existing types thanks to Al Davis and Kid Kalanon.
- Added support for `BIT` types (51870) to `PreparedStatement`.
- Fixed `getRow()` bug (527165) in `ResultSet`.
- Fixes for `ResultSet` updatability in `PreparedStatement`.
- Fixed time zone off-by-1-hour bug in `PreparedStatement` (538286, 528785).
- `ResultSet`: Fixed updatability (values being set to `null` if not updated).
- `DataSources` - fixed `setUrl` bug (511614, 525565), wrong datasource class name (532816, 528767).
- Added identifier quoting to all `DatabaseMetaData` methods that need them (should fix 518108).
- Added support for `YEAR` type (533556).
- `ResultSet.insertRow()` should now detect `auto_increment` fields in most cases and use that value in the new row. This detection will not work in multi-valued keys, however, due to the fact that the MySQL protocol does not return this information.
- `ResultSet.refreshRow()` implemented.
- Fixed `testsuite.Traversal afterLast()` bug, thanks to Igor Lastric.

## D.8.41. Changes in MySQL Connector/J 2.0.11 (27 January 2002)

- Fixed missing `DELETE_RULE` value in `DBMD.getImported/ExportedKeys()` and `getCrossReference()`.
- Full synchronization of `Statement.java`.
- More changes to fix `Unexpected end of input stream` errors when reading `BLOB` values. This should be the last fix.

## D.8.42. Changes in MySQL Connector/J 2.0.10 (24 January 2002)

- Fixed spurious `Unexpected end of input stream` errors in `MysqlIO` (bug 507456).
- Fixed null-pointer-exceptions when using `MysqlConnectionPoolDataSource` with WebSphere 4 (bug 505839).

## D.8.43. Changes in MySQL Connector/J 2.0.9 (13 January 2002)

- `Ant` build was corrupting included `jar` files, fixed (bug 487669).
- Fixed extra memory allocation in `MysqlIO.readPacket()` (bug 488663).
- Implementation of `DatabaseMetaData.getExported/ImportedKeys()` and `getCrossReference()`.
- Full synchronization on methods modifying instance and class-shared references, driver should be entirely thread-safe now (please let me know if you have problems).
- `DataSource` implementations moved to `org.gjt.mm.mysql.jdbc2.optional` package, and (initial) implementations of `PooledConnectionDataSource` and `XADataSource` are in place (thanks to Todd Wolff for the implementation and testing of `PooledConnectionDataSource` with IBM WebSphere 4).
- Added detection of network connection being closed when reading packets (thanks to Todd Lizambri).
- Fixed quoting error with escape processor (bug 486265).
- Report batch update support through `DatabaseMetaData` (bug 495101).
- Fixed off-by-one-hour error in `PreparedStatement.setTimestamp()` (bug 491577).
- Removed concatenation support from driver (the `||` operator), as older versions of VisualAge seem to be the only thing that use it, and it conflicts with the logical `||` operator. You will need to start `mysqld` with the `--ansi` flag to use the `||` operator as concatenation (bug 491680).
- Fixed casting bug in `PreparedStatement` (bug 488663).

## D.8.44. Changes in MySQL Connector/J 2.0.8 (25 November 2001)

- Batch updates now supported (thanks to some inspiration from Daniel Rall).
- `XADataSource/ConnectionPoolDataSource` code (experimental)
- `PreparedStatement.setAnyNumericType()` now handles positive exponents correctly (adds `+` so MySQL can understand it).
- `DatabaseMetaData.getPrimaryKeys()` and `getBestRowIdentifier()` are now

more robust in identifying primary keys (matches regardless of case or abbreviation/full spelling of `Primary Key` in `Key_type` column).

## D.8.45. Changes in MySQL Connector/J 2.0.7 (24 October 2001)

- `PreparedStatement.setCharacterStream()` now implemented
- Fixed dangling socket problem when in high availability (`autoReconnect=true`) mode, and finalizer for `Connection` will close any dangling sockets on GC.
- Fixed `ResultSetMetaData.getPrecision()` returning one less than actual on newer versions of MySQL.
- `ResultSet.getBlob()` now returns `null` if column value was `null`.
- Character sets read from database if `useUnicode=true` and `characterEncoding` is not set. (thanks to Dmitry Vereshchagin)
- Initial transaction isolation level read from database (if available). (thanks to Dmitry Vereshchagin)
- Fixed `DatabaseMetaData.supportsTransactions()`, and `supportsTransactionIsolationLevel()` and `getTypeInfo()` `SQL_DATETIME_SUB` and `SQL_DATA_TYPE` fields not being readable.
- Fixed `PreparedStatement` generating SQL that would end up with syntax errors for some queries.
- Fixed `ResultSet.isAfterLast()` always returning `false`.
- Fixed time zone issue in `PreparedStatement.setTimestamp()`. (thanks to Erik Olofsson)
- Capitalize type names when `capitalizeTypeNames=true` is passed in URL or properties (for WebObjects). (thanks to Anjo Krank)
- Updatable result sets now correctly handle `NULL` values in fields.
- `PreparedStatement.setDouble()` now uses full-precision doubles (reverting a fix made earlier to truncate them).
- `PreparedStatement.setBoolean()` will use 1/0 for values if your MySQL version is 3.21.23 or higher.

## D.8.46. Changes in MySQL Connector/J 2.0.6 (16 June 2001)

- Fixed `PreparedStatement` parameter checking.
- Fixed case-sensitive column names in `ResultSet.java`.

## D.8.47. Changes in MySQL Connector/J 2.0.5 (13 June 2001)

- Fixed `ResultSet.getBlob()` `ArrayIndex` out-of-bounds.

- Fixed `ResultSetMetaData.getColumnTypeName()` for `TEXT/BLOB`.
- Fixed `ArrayIndexOutOfBoundsException` when sending large `BLOB` queries. (Max size packet was not being set)
- Added `ISOLATION` level support to `Connection.setIsolationLevel()`
- Fixed NPE on `PreparedStatement.executeUpdate()` when all columns have not been set.
- Fixed data parsing of `TIMESTAMP` values with 2-digit years.
- Added `Byte` to `PreparedStatement.setObject()`.
- `ResultSet.getBoolean()` now recognizes `-1` as `true`.
- `ResultSet` has `+/-Inf/inf` support.
- `ResultSet.insertRow()` works now, even if not all columns are set (they will be set to `NULL`).
- `DatabaseMetaData.getCrossReference()` no longer `ArrayIndexOOB`.
- `getObject()` on `ResultSet` correctly does `TINYINT->Byte` and `SMALLINT->Short`.

## D.8.48. Changes in MySQL Connector/J 2.0.3 (03 December 2000)

- Implemented `getBigDecimal()` without scale component for JDBC2.
- Fixed composite key problem with updatable result sets.
- Added detection of `-/+INF` for doubles.
- Faster ASCII string operations.
- Fixed incorrect detection of `MAX_ALLOWED_PACKET`, so sending large blobs should work now.
- Fixed off-by-one error in `java.sql.Blob` implementation code.
- Added `ultraDevHack` URL parameter, set to `true` to allow (broken) Macromedia UltraDev to use the driver.

## D.8.49. Changes in MySQL Connector/J 2.0.1 (06 April 2000)

- Fixed `RSMD.isWritable()` returning wrong value. Thanks to Moritz Maass.
- Cleaned up exception handling when driver connects.
- Columns that are of type `TEXT` now return as `Strings` when you use `getObject()`.
- `DatabaseMetaData.getPrimaryKeys()` now works correctly with respect to `key_seq`. Thanks to Brian Slesinsky.
- No escape processing is done on `PreparedStatements` anymore per JDBC spec.

- Fixed many JDBC-2.0 traversal, positioning bugs, especially with respect to empty result sets. Thanks to Ron Smits, Nick Brook, Cessar Garcia and Carlos Martinez.
- Fixed some issues with updatability support in [ResultSet](#) when using multiple primary keys.

### **D.8.50. Changes in MySQL Connector/J 2.0.0pre5 (21 February 2000)**

- Fixed Bad Handshake problem.

### **D.8.51. Changes in MySQL Connector/J 2.0.0pre4 (10 January 2000)**

- Fixes to [ResultSet](#) for [insertRow\(\)](#) - Thanks to Cesar Garcia
- Fix to Driver to recognize JDBC-2.0 by loading a JDBC-2.0 class, instead of relying on JDK version numbers. Thanks to John Baker.
- Fixed [ResultSet](#) to return correct row numbers
- [Statement.updateCount\(\)](#) now returns rows matched, instead of rows actually updated, which is more SQL-92 like.

10-29-99

- [Statement/PreparedStatement.getMoreResults\(\)](#) bug fixed. Thanks to Noel J. Bergman.
- Added [Short](#) as a type to [PreparedStatement.setObject\(\)](#). Thanks to Jeff Crowder
- Driver now automatically configures maximum/preferred packet sizes by querying server.
- Autoreconnect code uses fast ping command if server supports it.
- Fixed various bugs with respect to packet sizing when reading from the server and when alloc'ing to write to the server.

### **D.8.52. Changes in MySQL Connector/J 2.0.0pre (17 August 1999)**

- Now compiles under JDK-1.2. The driver supports both JDK-1.1 and JDK-1.2 at the same time through a core set of classes. The driver will load the appropriate interface classes at runtime by figuring out which JVM version you are using.
- Fixes for result sets with all nulls in the first row. (Pointed out by Tim Endres)
- Fixes to column numbers in [SQLExceptions](#) in [ResultSet](#) (Thanks to Blas Rodriguez Somoza)
- The database no longer needs to be specified to connect. (Thanks to Christian Motschke)

### D.8.53. Changes in MySQL Connector/J 1.2b (04 July 1999)

- Better Documentation (in progress), in doc/mm.doc/book1.html
- DBMD now allows null for a column name pattern (not in spec), which it changes to '%'
- DBMD now has correct types/lengths for getXXX().
- ResultSet.getDate(), getTime(), and getTimestamp() fixes. (contributed by Alan Wilken)
- EscapeProcessor now handles \{ \} and { or } inside quotes correctly. (thanks to Alik for some ideas on how to fix it)
- Fixes to properties handling in Connection. (contributed by Juho Tikkala)
- ResultSet.getObject() now returns null for NULL columns in the table, rather than bombing out. (thanks to Ben Grosman)
- ResultSet.getObject() now returns Strings for types from MySQL that it doesn't know about. (Suggested by Chris Perdue)
- Removed DataInput/Output streams, not needed, 1/2 number of method calls per IO operation.
- Use default character encoding if one is not specified. This is a work-around for broken JVMs, because according to spec, EVERY JVM must support "ISO8859\_1", but they don't.
- Fixed Connection to use the platform character encoding instead of "ISO8859\_1" if one isn't explicitly set. This fixes problems people were having loading the character- converter classes that didn't always exist (JVM bug). (thanks to Fritz Elfert for pointing out this problem)
- Changed MySQLIO to re-use packets where possible to reduce memory usage.
- Fixed escape-processor bugs pertaining to { } inside quotes.

### D.8.54. Changes in MySQL Connector/J 1.2a (14 April 1999)

- Fixed character-set support for non-Javasoft JVMs (thanks to many people for pointing it out)
- Fixed ResultSet.getBoolean() to recognize 'y' & 'n' as well as '1' & '0' as boolean flags. (thanks to Tim Pizey)
- Fixed ResultSet.getTimestamp() to give better performance. (thanks to Richard Swift)
- Fixed getByte() for numeric types. (thanks to Ray Bellis)
- Fixed DatabaseMetaData.getTypeInfo() for DATE type. (thanks to Paul Johnston)
- Fixed EscapeProcessor for "fn" calls. (thanks to Piyush Shah at locomotive.org)
- Fixed EscapeProcessor to not do extraneous work if there are no escape codes. (thanks to Ryan Gustafson)
- Fixed Driver to parse URLs of the form "jdbc:mysql://host:port" (thanks to Richard Lobb)

### D.8.55. Changes in MySQL Connector/J 1.1i (24 March 1999)

- Fixed Timestamps for PreparedStatements
- Fixed null pointer exceptions in RSMD and RS
- Re-compiled with jikes for valid class files (thanks ms!)

### **D.8.56. Changes in MySQL Connector/J 1.1h (08 March 1999)**

- Fixed escape processor to deal with unmatched { and } (thanks to Craig Coles)
- Fixed escape processor to create more portable (between DATETIME and TIMESTAMP types) representations so that it will work with BETWEEN clauses. (thanks to Craig Longman)
- MysqlIO.quit() now closes the socket connection. Before, after many failed connections some OS's would run out of file descriptors. (thanks to Michael Brinkman)
- Fixed NullPointerException in Driver.getPropertyInfo. (thanks to Dave Potts)
- Fixes to MysqlDefs to allow all \*text fields to be retrieved as Strings. (thanks to Chris at Leverage)
- Fixed setDouble in PreparedStatement for large numbers to avoid sending scientific notation to the database. (thanks to J.S. Ferguson)
- Fixed getScale() and getPrecision() in RSMD. (contrib'd by James Klicman)
- Fixed getObject() when field was DECIMAL or NUMERIC (thanks to Bert Hobbs)
- DBMD.getTables() bombed when passed a null table-name pattern. Fixed. (thanks to Richard Lobb)
- Added check for "client not authorized" errors during connect. (thanks to Hannes Wallnoefer)

### **D.8.57. Changes in MySQL Connector/J 1.1g (19 February 1999)**

- Result set rows are now byte arrays. Blobs and Unicode work bidirectionally now. The useUnicode and encoding options are implemented now.
- Fixes to PreparedStatement to send binary set by setXXXStream to be sent untouched to the MySQL server.
- Fixes to getDriverPropertyInfo().

### **D.8.58. Changes in MySQL Connector/J 1.1f (31 December 1998)**

- Changed all ResultSet fields to Strings, this should allow Unicode to work, but your JVM must be able to convert between the character sets. This should also make reading data from the server be a bit quicker, because there is now no conversion from StringBuffer to String.
- Changed PreparedStatement.streamToString() to be more efficient (code from Uwe Schaefer).



- URL parsing is more robust (throws SQL exceptions on errors rather than NullPointerExceptions)
- PreparedStatement now can convert Strings to Time/Date values via setObject() (code from Robert Currey).
- IO no longer hangs in Buffer.readInt(), that bug was introduced in 1.1d when changing to all byte-arrays for result sets. (Pointed out by Samo Login)

### **D.8.59. Changes in MySQL Connector/J 1.1b (03 November 1998)**

- Fixes to DatabaseMetaData to allow both IBM VA and J-Builder to work. Let me know how it goes. (thanks to Jac Kersing)
- Fix to ResultSet.getBoolean() for NULL strings (thanks to Barry Lagerweij)
- Beginning of code cleanup, and formatting. Getting ready to branch this off to a parallel JDBC-2.0 source tree.
- Added "final" modifier to critical sections in MySQLIO and Buffer to allow compiler to inline methods for speed.

9-29-98

- If object references passed to setXXX() in PreparedStatement are null, setNull() is automatically called for you. (Thanks for the suggestion goes to Erik Ostrom)
- setObject() in PreparedStatement will now attempt to write a serialized representation of the object to the database for objects of Types.OTHER and objects of unknown type.
- Util now has a static method readObject() which given a ResultSet and a column index will re-instantiate an object serialized in the above manner.

### **D.8.60. Changes in MySQL Connector/J 1.1 (02 September 1998)**

- Got rid of "ugly hack" in MySQLIO.nextRow(). Rather than catch an exception, Buffer.isLastDataPacket() was fixed.
- Connection.getCatalog() and Connection.setCatalog() should work now.
- Statement.setMaxRows() works, as well as setting by property maxRows. Statement.setMaxRows() overrides maxRows set via properties or url parameters.
- Automatic re-connection is available. Because it has to "ping" the database before each query, it is turned off by default. To use it, pass in "autoReconnect=true" in the connection URL. You may also change the number of reconnect tries, and the initial timeout value via "maxReconnects=n" (default 3) and "initialTimeout=n" (seconds, default 2) parameters. The timeout is an exponential backoff type of timeout; for example, if you have initial timeout of 2 seconds, and maxReconnects of 3, then the driver will timeout 2 seconds, 4 seconds, then 16 seconds between each re-connection attempt.

## D.8.61. Changes in MySQL Connector/J 1.0 (24 August 1998)

- Fixed handling of blob data in Buffer.java
- Fixed bug with authentication packet being sized too small.
- The JDBC Driver is now under the LPGL

8-14-98

- Fixed Buffer.readLenString() to correctly read data for BLOBS.
- Fixed PreparedStatement.toString() to correctly read data for BLOBS.
- Fixed PreparedStatement.setDate() to not add a day. (above fixes thanks to Vincent Partington)
- Added URL parameter parsing (?user=... and so forth).

## D.8.62. Changes in MySQL Connector/J 0.9d (04 August 1998)

- Big news! New package name. Tim Endres from ICE Engineering is starting a new source tree for GNU GPL'd Java software. He's graciously given me the org.gjt.mm package directory to use, so now the driver is in the org.gjt.mm.mysql package scheme. I'm "legal" now. Look for more information on Tim's project soon.
- Now using dynamically sized packets to reduce memory usage when sending commands to the DB.
- Small fixes to getTypeInfo() for parameters, and so forth.
- DatabaseMetaData is now fully implemented. Let me know if these drivers work with the various IDEs out there. I've heard that they're working with JBuilder right now.
- Added JavaDoc documentation to the package.
- Package now available in .zip or .tar.gz.

## D.8.63. Changes in MySQL Connector/J 0.9 (28 July 1998)

- Implemented getTypeInfo(). Connection.rollback() now throws an SQLException per the JDBC spec.
- Added PreparedStatement that supports all JDBC API methods for PreparedStatement including InputStreams. Please check this out and let me know if anything is broken.
- Fixed a bug in ResultSet that would break some queries that only returned 1 row.
- Fixed bugs in DatabaseMetaData.getTables(), DatabaseMetaData.getColumns() and DatabaseMetaData.getCatalogs().
- Added functionality to Statement that allows executeUpdate() to store values for IDs that are auto-

matically generated for AUTO\_INCREMENT fields. Basically, after an executeUpdate(), look at the SQLWarnings for warnings like "LAST\_INSERTED\_ID = 'some number', COMMAND = 'your SQL query'". If you are using AUTO\_INCREMENT fields in your tables and are executing a lot of executeUpdate(s) on one Statement, be sure to clearWarnings() every so often to save memory.

### D.8.64. Changes in MySQL Connector/J 0.8 (06 July 1998)

- Split MysqlIO and Buffer to separate classes. Some ClassLoaders gave an IllegalAccessException error for some fields in those two classes. Now mm.mysql works in applets and all classloaders. Thanks to Joe Ennis <jce@mail.boone.com> for pointing out the problem and working on a fix with me.

### D.8.65. Changes in MySQL Connector/J 0.7 (01 July 1998)

- Fixed DatabaseMetadata problems in getColumnns() and bug in switch statement in the Field constructor. Thanks to Costin Manolache <costin@tdiinc.com> for pointing these out.

### D.8.66. Changes in MySQL Connector/J 0.6 (21 May 1998)

- Incorporated efficiency changes from Richard Swift <Richard.Swift@kanatek.ca> in [MysqlIO.java](#) and [ResultSet.java](#):
- We're now 15% faster than gwe's driver.
- Started working on [DatabaseMetaData](#).
- The following methods are implemented:
  - [getTables\(\)](#)
  - [getTableTypes\(\)](#)
  - [getColumnns](#)
  - [getCatalogs\(\)](#)

---

## Appendix E. Porting to Other Systems

This appendix helps you port MySQL to other operating systems. Do check the list of currently supported operating systems first. See [Section 2.1.1, “Operating Systems Supported by MySQL”](#). If you have created a new port of MySQL, please let us know so that we can list it here and on our Web site (<http://www.mysql.com/>), recommending it to other users.

Note: If you create a new port of MySQL, you are free to copy and distribute it under the GPL license, but it does not make you a copyright holder of MySQL.

A working POSIX thread library is needed for the server. On Solaris 2.5 we use Sun PThreads (the native thread support in 2.4 and earlier versions is not good enough), on Linux we use LinuxThreads by Xavier Leroy, <[Xavier.Leroy@inria.fr](mailto:Xavier.Leroy@inria.fr)>.

The hard part of porting to a new Unix variant without good native thread support is probably to port MIT-pthreads. See [mit-pthreads/README](#) and Programming POSIX Threads (<http://www.humanfactor.com/pthreads/>).

Up to MySQL 4.0.2, the MySQL distribution included a patched version of Chris Provenzano's Pthreads from MIT (see the MIT Pthreads Web page at <http://www.mit.edu/afs/sipb/project/pthreads/> and a programming introduction at [http://www.mit.edu:8001/people/proven/IAP\\_2000/](http://www.mit.edu:8001/people/proven/IAP_2000/)). These can be used for some operating systems that do not have POSIX threads. See [Section 2.9.5, “MIT-pthreads Notes”](#).

It is also possible to use another user level thread package named FSU Pthreads (see <http://moss.csc.ncsu.edu/~mueller/pthreads/>). This implementation is being used for the SCO port.

See the `thr_lock.c` and `thr_alarm.c` programs in the `mysys` directory for some tests/examples of these problems.

Both the server and the client need a working C++ compiler. We use `gcc` on many platforms. Other compilers that are known to work are SPARCworks, Sun Forte, Irix `cc`, HP-UX `aCC`, IBM AIX `xlC_r`, Intel `ecc/icc` and Compaq `cxx`).

To compile only the client use `./configure --without-server`.

There is currently no support for only compiling the server, nor is it likely to be added unless someone has a good reason for it.

If you want/need to change any `Makefile` or the configure script you also need GNU Automake and Autoconf. See [Section 2.9.3, “Installing from the Development Source Tree”](#).

All steps needed to remake everything from the most basic files.

```
/bin/rm */.deps/*.P
/bin/rm -f config.cache
aclocal
autoheader
aclocal
automake
autoconf
./configure --with-debug=full --prefix='your installation directory'

# The makefiles generated above need GNU make 3.75 or newer.
# (called gmake below)
gmake clean all install init-db
```

If you run into problems with a new port, you may have to do some debugging of MySQL! See [Section E.1, “Debugging a MySQL Server”](#).

**Note:** Before you start debugging `mysqld`, first get the test programs `mysys/thr_alarm` and

`mysys/thr_lock` to work. This ensures that your thread installation has even a remote chance to work!

## E.1. Debugging a MySQL Server

If you are using some functionality that is very new in MySQL, you can try to run `mysqld` with the `--skip-new` (which disables all new, potentially unsafe functionality) or with `--safe-mode` which disables a lot of optimization that may cause problems. See [Section A.4.2, “What to Do If MySQL Keeps Crashing”](#).

If `mysqld` doesn't want to start, you should verify that you don't have any `my.cnf` files that interfere with your setup! You can check your `my.cnf` arguments with `mysqld --print-defaults` and avoid using them by starting with `mysqld --no-defaults ...`.

If `mysqld` starts to eat up CPU or memory or if it “hangs,” you can use `mysqladmin processlist status` to find out if someone is executing a query that takes a long time. It may be a good idea to run `mysqladmin -i10 processlist status` in some window if you are experiencing performance problems or problems when new clients can't connect.

The command `mysqladmin debug` dumps some information about locks in use, used memory and query usage to the MySQL log file. This may help solve some problems. This command also provides some useful information even if you haven't compiled MySQL for debugging!

If the problem is that some tables are getting slower and slower you should try to optimize the table with `OPTIMIZE TABLE` or `myisamchk`. See [Chapter 5, Database Administration](#). You should also check the slow queries with `EXPLAIN`.

You should also read the OS-specific section in this manual for problems that may be unique to your environment. See [Section 2.13, “Operating System-Specific Notes”](#).

### E.1.1. Compiling MySQL for Debugging

If you have some very specific problem, you can always try to debug MySQL. To do this you must configure MySQL with the `--with-debug` or the `--with-debug=full` option. You can check whether MySQL was compiled with debugging by doing: `mysqld --help`. If the `--debug` flag is listed with the options then you have debugging enabled. `mysqladmin ver` also lists the `mysqld` version as `mysql ... --debug` in this case.

If you are using `gcc` or `egcs`, the recommended `configure` line is:

```
CC=gcc CFLAGS="-O2" CXX=gcc CXXFLAGS="-O2 -felide-constructors \  
-fno-exceptions -fno-rtti" ./configure --prefix=/usr/local/mysql \  
--with-debug --with-extra-charsets=complex
```

This avoids problems with the `libstdc++` library and with C++ exceptions (many compilers have problems with C++ exceptions in threaded code) and compile a MySQL version with support for all character sets.

If you suspect a memory overrun error, you can configure MySQL with `--with-debug=full`, which installs a memory allocation (`SAFEMALLOC`) checker. However, running with `SAFEMALLOC` is quite slow, so if you get performance problems you should start `mysqld` with the `--skip-safemalloc` option. This disables the memory overrun checks for each call to `malloc()` and `free()`.

If `mysqld` stops crashing when you compile it with `--with-debug`, you probably have found a compiler bug or a timing bug within MySQL. In this case, you can try to add `-g` to the `CFLAGS` and `CXXFLAGS` variables above and not use `--with-debug`. If `mysqld` dies, you can at least attach to it with `gdb` or use `gdb` on the core file to find out what happened.

When you configure MySQL for debugging you automatically enable a lot of extra safety check functions that monitor the health of `mysqld`. If they find something “unexpected,” an entry is written to `stderr`, which `mysqld_safe` directs to the error log! This also means that if you are having some unexpected problems with MySQL and are using a source distribution, the first thing you should do is to configure MySQL for debugging! (The second thing is to send mail to a MySQL mailing list and ask for help. See [Section 1.7.1, “MySQL Mailing Lists”](#). If you believe that you have found a bug, please use the instructions at [Section 1.8, “How to Report Bugs or Problems”](#)).

In the Windows MySQL distribution, `mysqld.exe` is by default compiled with support for trace files.

## E.1.2. Creating Trace Files

If the `mysqld` server doesn't start or if you can cause it to crash quickly, you can try to create a trace file to find the problem.

To do this, you must have a `mysqld` that has been compiled with debugging support. You can check this by executing `mysqld -V`. If the version number ends with `-debug`, it's compiled with support for trace files. (On Windows, the debugging server is named `mysqld-debug` rather than `mysqld` as of MySQL 4.1.)

Start the `mysqld` server with a trace log in `/tmp/mysqld.trace` on Unix or `C:\mysqld.trace` on Windows:

```
shell> mysqld --debug
```

On Windows, you should also use the `--standalone` flag to not start `mysqld` as a service. In a console window, use this command:

```
C:\> mysqld-debug --debug --standalone
```

After this, you can use the `mysql.exe` command-line tool in a second console window to reproduce the problem. You can stop the `mysqld` server with `mysqladmin shutdown`.

Note that the trace file become **very big**! If you want to generate a smaller trace file, you can use debugging options something like this:

```
mysqld --debug=d,info,error,query,general,where:0,/tmp/mysqld.trace
```

This only prints information with the most interesting tags to the trace file.

If you make a bug report about this, please only send the lines from the trace file to the appropriate mailing list where something seems to go wrong! If you can't locate the wrong place, you can ftp the trace file, together with a full bug report, to <ftp://ftp.mysql.com/pub/mysql/upload/> so that a MySQL developer can take a look at it.

The trace file is made with the **DEBUG** package by Fred Fish. See [Section E.3, “The DEBUG Package”](#).

## E.1.3. Debugging `mysqld` under `gdb`

On most systems you can also start `mysqld` from `gdb` to get more information if `mysqld` crashes.

With some older `gdb` versions on Linux you must use `run --one-thread` if you want to be able to debug `mysqld` threads. In this case, you can only have one thread active at a time. We recommend you to upgrade to `gdb 5.1` ASAP as thread debugging works much better with this version!

NTPL threads (the new thread library on Linux) may cause problems while running `mysqld` under `gdb`. Some symptoms are:

- `mysqld` hangs during startup (before it writes `ready for connections`).
- `mysqld` crashes during a `pthread_mutex_lock()` or `pthread_mutex_unlock()` call.

In this case, you should set the following environment variable in the shell before starting `gdb`:

```
LD_ASSUME_KERNEL=2.4.1
export LD_ASSUME_KERNEL
```

When running `mysqld` under `gdb`, you should disable the stack trace with `--skip-stack-trace` to be able to catch segfaults within `gdb`.

In MySQL 4.0.14 and above you should use the `--gdb` option to `mysqld`. This installs an interrupt handler for `SIGINT` (needed to stop `mysqld` with `^C` to set breakpoints) and disable stack tracing and core file handling.

It's very hard to debug MySQL under `gdb` if you do a lot of new connections the whole time as `gdb` doesn't free the memory for old threads. You can avoid this problem by starting `mysqld` with `-thread_cache_size='max_connections+1'`. In most cases just using `-thread_cache_size=5` helps a lot!

If you want to get a core dump on Linux if `mysqld` dies with a `SIGSEGV` signal, you can start `mysqld` with the `--core-file` option. This core file can be used to make a backtrace that may help you find out why `mysqld` died:

```
shell> gdb mysqld core
gdb> backtrace full
gdb> quit
```

See [Section A.4.2, “What to Do If MySQL Keeps Crashing”](#).

If you are using `gdb` 4.17.x or above on Linux, you should install a `.gdb` file, with the following information, in your current directory:

```
set print sevenbit off
handle SIGUSR1 nostop noprint
handle SIGUSR2 nostop noprint
handle SIGWAITING nostop noprint
handle SIGLWP nostop noprint
handle SIGPIPE nostop
handle SIGALRM nostop
handle SIGHUP nostop
handle SIGTERM nostop noprint
```

If you have problems debugging threads with `gdb`, you should download `gdb` 5.x and try this instead. The new `gdb` version has very improved thread handling!

Here is an example how to debug `mysqld`:

```
shell> gdb /usr/local/libexec/mysqld
gdb> run
...
backtrace full # Do this when mysqld crashes
```

Include the above output in a bug report, which you can file using the instructions in [Section 1.8, “How to Report Bugs or Problems”](#).

If `mysqld` hangs you can try to use some system tools like `strace` or `/usr/proc/bin/pstack` to examine where `mysqld` has hung.

```
strace /tmp/log libexec/mysqld
```

If you are using the Perl [DBI](#) interface, you can turn on debugging information by using the `trace` method or by setting the `DBI_TRACE` environment variable.

## E.1.4. Using a Stack Trace

On some operating systems, the error log contains a stack trace if `mysqld` dies unexpectedly. You can use this to find out where (and maybe why) `mysqld` died. See [Section 5.11.1, “The Error Log”](#). To get a stack trace, you must not compile `mysqld` with the `-fomit-frame-pointer` option to gcc. See [Section E.1.1, “Compiling MySQL for Debugging”](#).

If the error file contains something like the following:

```
mysqld got signal 11;
The manual section 'Debugging a MySQL server' tells you how to use a
stack trace and/or the core file to produce a readable backtrace that may
help in finding out why mysqld died
Attempting backtrace. You can use the following information to find out
where mysqld died.  If you see no messages after this, something went
terribly wrong...
stack range sanity check, ok, backtrace follows
0x40077552
0x81281a0
0x8128f47
0x8127be0
0x8127995
0x8104947
0x80ff28f
0x810131b
0x80ee4bc
0x80c3c91
0x80c6b43
0x80c1fd9
0x80c1686
```

you can find where `mysqld` died by doing the following:

1. Copy the preceding numbers to a file, for example `mysqld.stack`.
2. Make a symbol file for the `mysqld` server:

```
nm -n libexec/mysqld > /tmp/mysqld.sym
```

Note that most MySQL binary distributions (except for the "debug" packages, where this information is included inside of the binaries themselves) ship with the above file, named `mysqld.sym.gz`. In this case, you can simply unpack it by doing:

```
gunzip < bin/mysqld.sym.gz > /tmp/mysqld.sym
```

3. Execute `resolve_stack_dump -s /tmp/mysqld.sym -n mysqld.stack`.

This prints out where `mysqld` died. If this doesn't help you find out why `mysqld` died, you should make a bug report and include the output from the above command with the bug report.

Note however that in most cases it does not help us to just have a stack trace to find the reason for the problem. To be able to locate the bug or provide a workaround, we would in most cases need to know the query that killed `mysqld` and preferably a test case so that we can repeat the problem! See [Section 1.8, “How to Report Bugs or Problems”](#).



## E.1.5. Using Server Logs to Find Causes of Errors in `mysqld`

Note that before starting `mysqld` with `--log` you should check all your tables with `myisamchk`. See [Chapter 5, Database Administration](#).

If `mysqld` dies or hangs, you should start `mysqld` with `--log`. When `mysqld` dies again, you can examine the end of the log file for the query that killed `mysqld`.

If you are using `--log` without a file name, the log is stored in the database directory as `host_name.log`. In most cases it is the last query in the log file that killed `mysqld`, but if possible you should verify this by restarting `mysqld` and executing the found query from the `mysql` command-line tools. If this works, you should also test all complicated queries that didn't complete.

You can also try the command `EXPLAIN` on all `SELECT` statements that takes a long time to ensure that `mysqld` is using indexes properly. See [Section 7.2.1, "Optimizing Queries with EXPLAIN"](#).

You can find the queries that take a long time to execute by starting `mysqld` with `--log-slow-queries`. See [Section 5.11.5, "The Slow Query Log"](#).

If you find the text `mysqld restarted` in the error log file (normally named `hostname.err`) you probably have found a query that causes `mysqld` to fail. If this happens, you should check all your tables with `myisamchk` (see [Chapter 5, Database Administration](#)), and test the queries in the MySQL log files to see whether one fails. If you find such a query, try first upgrading to the newest MySQL version. If this doesn't help and you can't find anything in the `mysql` mail archive, you should report the bug to a MySQL mailing list. The mailing lists are described at <http://lists.mysql.com/>, which also has links to online list archives.

If you have started `mysqld` with `myisam-recover`, MySQL automatically checks and tries to repair `MyISAM` tables if they are marked as 'not closed properly' or 'crashed'. If this happens, MySQL writes an entry in the `hostname.err` file `'Warning: Checking table ...'` which is followed by `Warning: Repairing table` if the table needs to be repaired. If you get a lot of these errors, without `mysqld` having died unexpectedly just before, then something is wrong and needs to be investigated further. See [Section 5.2.1, "Command Options"](#).

It is not a good sign if `mysqld` did die unexpectedly, but in this case, you should not investigate the `Checking table...` messages, but instead try to find out why `mysqld` died.

## E.1.6. Making a Test Case If You Experience Table Corruption

If you get corrupted tables or if `mysqld` always fails after some update commands, you can test whether this bug is reproducible by doing the following:

- Take down the MySQL daemon (with `mysqladmin shutdown`).
- Make a backup of the tables (to guard against the very unlikely case that the repair does something bad).
- Check all tables with `myisamchk -s database/*.MYI`. Repair any wrong tables with `myisamchk -r database/table.MYI`.
- Make a second backup of the tables.
- Remove (or move away) any old log files from the MySQL data directory if you need more space.
- Start `mysqld` with `--log-bin`. See [Section 5.11.4, "The Binary Log"](#). If you want to find a query that crashes `mysqld`, you should use `--log --log-bin`.
- When you have gotten a crashed table, stop the `mysqld server`.

- Restore the backup.
- Restart the `mysqld` server **without** `--log-bin`
- Re-execute the commands with `mysqlbinlog update-log-file | mysql`. The update log is saved in the MySQL database directory with the name `hostname-bin.#`.
- If the tables are corrupted again or you can get `mysqld` to die with the above command, you have found a reproducible bug that should be easy to fix! FTP the tables and the binary log to <ftp://ftp.mysql.com/pub/mysql/upload/> and report it in our bugs database using the instructions given in [Section 1.8, “How to Report Bugs or Problems”](#). (Please note that the `/pub/mysql/upload/` FTP directory is not listable, so you'll not see what you've uploaded in your FTP client.) If you are a support customer, you can use the MySQL Customer Support Center <https://support.mysql.com/> to alert the MySQL team about the problem and have it fixed as soon as possible.

You can also use the script `mysql_find_rows` to just execute some of the update statements if you want to narrow down the problem.

## E.2. Debugging a MySQL Client

To be able to debug a MySQL client with the integrated debug package, you should configure MySQL with `--with-debug` or `--with-debug=full`. See [Section 2.9.2, “Typical configure Options”](#).

Before running a client, you should set the `MYSQL_DEBUG` environment variable:

```
shell> MYSQL_DEBUG=d:t:0,/tmp/client.trace
shell> export MYSQL_DEBUG
```

This causes clients to generate a trace file in `/tmp/client.trace`.

If you have problems with your own client code, you should attempt to connect to the server and run your query using a client that is known to work. Do this by running `mysql` in debugging mode (assuming that you have compiled MySQL with debugging on):

```
shell> mysql --debug=d:t:0,/tmp/client.trace
```

This provides useful information in case you mail a bug report. See [Section 1.8, “How to Report Bugs or Problems”](#).

If your client crashes at some 'legal' looking code, you should check that your `mysql.h` include file matches your MySQL library file. A very common mistake is to use an old `mysql.h` file from an old MySQL installation with new MySQL library.

## E.3. The DBUG Package

The MySQL server and most MySQL clients are compiled with the DBUG package originally created by Fred Fish. When you have configured MySQL for debugging, this package makes it possible to get a trace file of what the program is debugging. See [Section E.1.2, “Creating Trace Files”](#).

This section summarizes the argument values that you can specify in debug options on the command line for MySQL programs that have been built with debugging support. For more information about programming with the DBUG package, see the DBUG manual in the `dbug` directory of MySQL source distributions. It's best to use a recent distribution for MySQL 5.0 to get the most updated DBUG manual.

You use the debug package by invoking a program with the `--debug="..."` or the `-#...` option.

Most MySQL programs have a default debug string that is used if you don't specify an option to `--debug`. The default trace file is usually `/tmp/program_name.trace` on Unix and `\program_name.trace` on Windows.

The debug control string is a sequence of colon-separated fields as follows:

```
<field_1>:<field_2>:...:<field_N>
```

Each field consists of a mandatory flag character followed by an optional `' , '` and comma-separated list of modifiers:

```
flag[,modifier,modifier,...,modifier]
```

The currently recognized flag characters are:

Flag	Description
<code>d</code>	Enable output from <code>DEBUG_&lt;N&gt;</code> macros for the current state. May be followed by a list of keywords which selects output only for the <code>DEBUG</code> macros with that keyword. An empty list of keywords implies output for all macros.
<code>D</code>	Delay after each debugger output line. The argument is the number of tenths of seconds to delay, subject to machine capabilities. For example, <code>--D,20</code> specifies a delay of two seconds.
<code>f</code>	Limit debugging, tracing, and profiling to the list of named functions. Note that a null list disables all functions. The appropriate <code>d</code> or <code>t</code> flags must still be given; this flag only limits their actions if they are enabled.
<code>F</code>	Identify the source file name for each line of debug or trace output.
<code>i</code>	Identify the process with the PID or thread ID for each line of debug or trace output.
<code>g</code>	Enable profiling. Create a file called <code>debugmon.out</code> containing information that can be used to profile the program. May be followed by a list of keywords that select profiling only for the functions in that list. A null list implies that all functions are considered.
<code>L</code>	Identify the source file line number for each line of debug or trace output.
<code>n</code>	Print the current function nesting depth for each line of debug or trace output.
<code>N</code>	Number each line of debug output.
<code>o</code>	Redirect the debugger output stream to the specified file. The default output is <code>stderr</code> .
<code>O</code>	Like <code>o</code> , but the file is really flushed between each write. When needed, the file is closed and reopened between each write.
<code>p</code>	Limit debugger actions to specified processes. A process must be identified with the <code>DEBUG_PROCESS</code> macro and match one in the list for debugger actions to occur.
<code>P</code>	Print the current process name for each line of debug or trace output.
<code>r</code>	When pushing a new state, do not inherit the previous state's function nesting level. Useful when the output is to start at the left margin.
<code>S</code>	Do function <code>_sanity(_file_,_line_)</code> at each debugged function until <code>_sanity()</code> returns something that differs from 0. (Mostly used with <code>safemalloc</code> to find memory leaks)
<code>t</code>	Enable function call/exit trace lines. May be followed by a list (containing only one modifier) giving a numeric maximum trace level, beyond which no output occurs for either debugging or tracing macros. The default is a compile time option.

Some examples of debug control strings that might appear on a shell command line (the `--#` is typically used to introduce a control string to an application program) are:

```

-#d:t
-#d:f,main,subr1:F:L:t,20
-#d,input,output,files:n
-#d:t:i:O,\\mysqld.trace

```

In MySQL, common tags to print (with the `d` option) are `enter`, `exit`, `error`, `warning`, `info`, and `loop`.

## E.4. Comments about RTS Threads

I have tried to use the RTS thread packages with MySQL but stumbled on the following problems:

They use old versions of many POSIX calls and it is very tedious to make wrappers for all functions. I am inclined to think that it would be easier to change the thread libraries to the newest POSIX specification.

Some wrappers are currently written. See `mysys/my_pthread.c` for more info.

At least the following should be changed:

`pthread_get_specific` should use one argument. `sigwait` should take two arguments. A lot of functions (at least `pthread_cond_wait`, `pthread_cond_timedwait()`) should return the error code on error. Now they return -1 and set `errno`.

Another problem is that user-level threads use the `ALRM` signal and this aborts a lot of functions (`read`, `write`, `open`...). MySQL should do a retry on interrupt on all of these but it is not that easy to verify it.

The biggest unsolved problem is the following:

To get thread-level alarms I changed `mysys/thr_alarm.c` to wait between alarms with `pthread_cond_timedwait()`, but this aborts with error `EINTR`. I tried to debug the thread library as to why this happens, but couldn't find any easy solution.

If someone wants to try MySQL with RTS threads I suggest the following:

- Change functions MySQL uses from the thread library to POSIX. This shouldn't take that long.
- Compile all libraries with the `-DHAVE_rts_threads`.
- Compile `thr_alarm`.
- If there are some small differences in the implementation, they may be fixed by changing `my_pthread.h` and `my_pthread.c`.
- Run `thr_alarm`. If it runs without any “warning,” “error,” or aborted messages, you are on the right track. Here is a successful run on Solaris:

```

Main thread: 1
Thread 0 (5) started
Thread: 5 Waiting
process_alarm
Thread 1 (6) started
Thread: 6 Waiting
process_alarm
process_alarm
thread_alarm
Thread: 6 Slept for 1 (1) sec
Thread: 6 Waiting
process_alarm
process_alarm
thread_alarm
Thread: 6 Slept for 2 (2) sec

```

```
Thread: 6 Simulation of no alarm needed
Thread: 6 Slept for 0 (3) sec
Thread: 6 Waiting
process_alarm
process_alarm
thread_alarm
Thread: 6 Slept for 4 (4) sec
Thread: 6 Waiting
process_alarm
thread_alarm
Thread: 5 Slept for 10 (10) sec
Thread: 5 Waiting
process_alarm
process_alarm
thread_alarm
Thread: 6 Slept for 5 (5) sec
Thread: 6 Waiting
process_alarm
process_alarm
...
thread_alarm
Thread: 5 Slept for 0 (1) sec
end
```

## E.5. Differences Between Thread Packages

MySQL is very dependent on the thread package used. So when choosing a good platform for MySQL, the thread package is very important.

There are at least three types of thread packages:

- User threads in a single process. Thread switching is managed with alarms and the threads library manages all non-thread-safe functions with locks. Read, write and select operations are usually managed with a thread-specific select that switches to another thread if the running threads have to wait for data. If the user thread packages are integrated in the standard libs (FreeBSD and BSDI threads) the thread package requires less overhead than thread packages that have to map all unsafe calls (MIT-pthreads, FSU Pthreads and RTS threads). In some environments (for example, SCO), all system calls are thread-safe so the mapping can be done very easily (FSU Pthreads on SCO). Downside: All mapped calls take a little time and it's quite tricky to be able to handle all situations. There are usually also some system calls that are not handled by the thread package (like MIT-pthreads and sockets). Thread scheduling isn't always optimal.
- User threads in separate processes. Thread switching is done by the kernel and all data are shared between threads. The thread package manages the standard thread calls to allow sharing data between threads. LinuxThreads is using this method. Downside: Lots of processes. Thread creating is slow. If one thread dies the rest are usually left hanging and you must kill them all before restarting. Thread switching is somewhat expensive.
- Kernel threads. Thread switching is handled by the thread library or the kernel and is very fast. Everything is done in one process, but on some systems, `ps` may show the different threads. If one thread aborts, the whole process aborts. Most system calls are thread-safe and should require very little overhead. Solaris, HP-UX, AIX and OSF/1 have kernel threads.

In some systems kernel threads are managed by integrating user level threads in the system libraries. In such cases, the thread switching can only be done by the thread library and the kernel isn't really “thread aware.”

---

## Appendix F. Environment Variables

This appendix lists all the environment variables that are used directly or indirectly by MySQL. Most of these can also be found in other places in this manual.

Note that any options on the command line take precedence over values specified in option files and environment variables, and values in option files take precedence over values in environment variables.

In many cases, it is preferable to use an option file instead of environment variables to modify the behavior of MySQL. See [Section 4.3.2, “Using Option Files”](#).

Variable	Description
CXX	The name of your C++ compiler (for running <code>configure</code> ).
CC	The name of your C compiler (for running <code>configure</code> ).
CFLAGS	Flags for your C compiler (for running <code>configure</code> ).
CXXFLAGS	Flags for your C++ compiler (for running <code>configure</code> ).
DBI_USER	The default username for Perl DBI.
DBI_TRACE	Trace options for Perl DBI.
HOME	The default path for the <code>mysql</code> history file is <code>\$HOME/.mysql_history</code> .
LD_RUN_PATH	Used to specify the location of <code>libmysqlclient.so</code> .
MYSQL_DEBUG	Debug trace options when debugging.
MYSQL_GROUP_SUFFIX	Option group suffix value (like specifying <code>-defaults-group-suffix</code> ).
MYSQL_HISTFILE	The path to the <code>mysql</code> history file. If this variable is set, its value overrides the default for <code>\$HOME/.mysql_history</code> .
MYSQL_HOME	The path to the directory in which the server-specific <code>my.cnf</code> file resides (as of MySQL 5.0.3).
MYSQL_HOST	The default hostname used by the <code>mysql</code> command-line client.
MYSQL_PS1	The command prompt to use in the <code>mysql</code> command-line client.
MYSQL_PWD	The default password when connecting to <code>mysqld</code> . Note that using this is insecure. See <a href="#">Section 5.8.6, “Keeping Your Password Secure”</a> .
MYSQL_TCP_PORT	The default TCP/IP port number.
MYSQL_UNIX_PORT	The default Unix socket filename; used for connections to <code>localhost</code> .
PATH	Used by the shell to find MySQL programs.
TMPDIR	The directory where temporary files are created.
TZ	This should be set to your local time zone. See <a href="#">Section A.4.6, “Time Zone Problems”</a> .
UMASK_DIR	The user-directory creation mask when creating directories. Note that this is <b>ANDed</b> with <code>UMASK</code> .
UMASK	The user-file creation mask when creating files.
USER	The default username on Windows and NetWare used when connecting to <code>mysqld</code> .

---

# Appendix G. Regular Expressions

A regular expression is a powerful way of specifying a pattern for a complex search.

MySQL uses Henry Spencer's implementation of regular expressions, which is aimed at conformance with POSIX 1003.2. See [Appendix C, Credits](#). MySQL uses the extended version to support pattern-matching operations performed with the `REGEXP` operator in SQL statements. See [Section 3.3.4.7, "Pattern Matching"](#), and [Section 12.3.1, "String Comparison Functions"](#).

This appendix is a summary, with examples, of the special characters and constructs that can be used in MySQL for `REGEXP` operations. It does not contain all the details that can be found in Henry Spencer's `regex(7)` manual page. That manual page is included in MySQL source distributions, in the `regex.7` file under the `regex` directory.

A regular expression describes a set of strings. The simplest regular expression is one that has no special characters in it. For example, the regular expression `hello` matches `hello` and nothing else.

Non-trivial regular expressions use certain special constructs so that they can match more than one string. For example, the regular expression `hello|word` matches either the string `hello` or the string `word`.

As a more complex example, the regular expression `B[an]*s` matches any of the strings `Bananas`, `Baaaaas`, `Bs`, and any other string starting with a `B`, ending with an `s`, and containing any number of `a` or `n` characters in between.

A regular expression for the `REGEXP` operator may use any of the following special characters and constructs:

- `^`

Match the beginning of a string.

```
mysql> SELECT 'fo\nfo' REGEXP '^fo$';      -> 0
mysql> SELECT 'fofo' REGEXP '^fo';        -> 1
```

- `$`

Match the end of a string.

```
mysql> SELECT 'fo\no' REGEXP '^fo\no$';    -> 1
mysql> SELECT 'fo\no' REGEXP '^fo$';       -> 0
```

- `.`

Match any character (including carriage return and newline).

```
mysql> SELECT 'fofo' REGEXP '^f.*$';       -> 1
mysql> SELECT 'fo\r\nfo' REGEXP '^f.*$';   -> 1
```

- `a*`

Match any sequence of zero or more `a` characters.

```
mysql> SELECT 'Ban' REGEXP '^Ba*n';        -> 1
mysql> SELECT 'Baaan' REGEXP '^Ba*n';     -> 1
mysql> SELECT 'Bn' REGEXP '^Ba*n';        -> 1
```

- `a+`

Match any sequence of one or more `a` characters.

```
mysql> SELECT 'Ban' REGEXP '^Ba+n';      -> 1
mysql> SELECT 'Bn' REGEXP '^Ba+n';      -> 0
```

- `a?`

Match either zero or one `a` character.

```
mysql> SELECT 'Bn' REGEXP '^Ba?n';      -> 1
mysql> SELECT 'Ban' REGEXP '^Ba?n';     -> 1
mysql> SELECT 'Baan' REGEXP '^Ba?n';    -> 0
```

- `de|abc`

Match either of the sequences `de` or `abc`.

```
mysql> SELECT 'pi' REGEXP 'pi|apa';     -> 1
mysql> SELECT 'axe' REGEXP 'pi|apa';    -> 0
mysql> SELECT 'apa' REGEXP 'pi|apa';    -> 1
mysql> SELECT 'apa' REGEXP '^(pi|apa)$'; -> 1
mysql> SELECT 'pi' REGEXP '^(pi|apa)$'; -> 1
mysql> SELECT 'pix' REGEXP '^(pi|apa)$'; -> 0
```

- `(abc)*`

Match zero or more instances of the sequence `abc`.

```
mysql> SELECT 'pi' REGEXP '^(pi)*$';    -> 1
mysql> SELECT 'pip' REGEXP '^(pi)*$';   -> 0
mysql> SELECT 'pipi' REGEXP '^(pi)*$';  -> 1
```

- `{1}, {2, 3}`

`{n}` or `{m, n}` notation provides a more general way of writing regular expressions that match many occurrences of the previous atom (or “piece”) of the pattern. `m` and `n` are integers.

- `a*`

Can be written as `a{0, }`.

- `a+`

Can be written as `a{1, }`.

- `a?`

Can be written as `a{0, 1}`.

To be more precise, `a{n}` matches exactly `n` instances of `a`. `a{n, }` matches `n` or more instances of `a`. `a{m, n}` matches `m` through `n` instances of `a`, inclusive.

`m` and `n` must be in the range from 0 to `RE_DUP_MAX` (default 255), inclusive. If both `m` and `n` are given, `m` must be less than or equal to `n`.

```
mysql> SELECT 'abcde' REGEXP 'a[bcd]{2}e'; -> 0
mysql> SELECT 'abcde' REGEXP 'a[bcd]{3}e'; -> 1
mysql> SELECT 'abcde' REGEXP 'a[bcd]{1,10}e'; -> 1
```



- `[a-dX], [^a-dX]`

Matches any character that is (or is not, if `^` is used) either `a`, `b`, `c`, `d` or `X`. A `-` character between two other characters forms a range that matches all characters from the first character to the second. For example, `[0-9]` matches any decimal digit. To include a literal `]` character, it must immediately follow the opening bracket `[`. To include a literal `-` character, it must be written first or last. Any character that does not have a defined special meaning inside a `[ ]` pair matches only itself.

```
mysql> SELECT 'aXbc' REGEXP '[a-dXYZ]';           -> 1
mysql> SELECT 'aXbc' REGEXP '^[a-dXYZ]$';       -> 0
mysql> SELECT 'aXbc' REGEXP '^a-dXYZ+$';       -> 1
mysql> SELECT 'aXbc' REGEXP '^[^a-dXYZ]+$';    -> 0
mysql> SELECT 'gheis' REGEXP '^[^a-dXYZ]+$';   -> 1
mysql> SELECT 'gheisa' REGEXP '^a-dXYZ+$';     -> 0
```

- `[.characters.]`

Within a bracket expression (written using `[` and `]`), matches the sequence of characters of that collating element. `characters` is either a single character or a character name like `newline`. You can find the full list of character names in the `regexp/cname.h` file.

```
mysql> SELECT '~' REGEXP '[[.~.]]';             -> 1
mysql> SELECT '~' REGEXP '[[.tilde.]]';        -> 1
```

- `[=character_class=]`

Within a bracket expression (written using `[` and `]`), `[=character_class=]` represents an equivalence class. It matches all characters with the same collation value, including itself. For example, if `o` and `(+)` are the members of an equivalence class, then `[=o=]`, `[=(+)=]`, and `[o(+)]` are all synonymous. An equivalence class may not be used as an endpoint of a range.

- `[:character_class:]`

Within a bracket expression (written using `[` and `]`), `[:character_class:]` represents a character class that matches all characters belonging to that class. The following table lists the standard class names. These names stand for the character classes defined in the `ctype(3)` manual page. A particular locale may provide other class names. A character class may not be used as an endpoint of a range.

<code>alnum</code>	Alphanumeric characters
<code>alpha</code>	Alphabetic characters
<code>blank</code>	Whitespace characters
<code>cntrl</code>	Control characters
<code>digit</code>	Digit characters
<code>graph</code>	Graphic characters
<code>lower</code>	Lowercase alphabetic characters
<code>print</code>	Graphic or space characters
<code>punct</code>	Punctuation characters
<code>space</code>	Space, tab, newline, and carriage return
<code>upper</code>	Uppercase alphabetic characters
<code>xdigit</code>	Hexadecimal digit characters

```
mysql> SELECT 'justalnums' REGEXP '[:alnum:]+'; -> 1
mysql> SELECT '!!!' REGEXP '[:alnum:]+';      -> 0
```

- `[[:<:]], [[:>:]]`

These markers stand for word boundaries. They match the beginning and end of words, respectively. A word is a sequence of word characters that is not preceded by or followed by word characters. A word character is an alphanumeric character in the `alnum` class or an underscore (`_`).

```
mysql> SELECT 'a word a' REGEXP '[:<:]word[:>:]'; -> 1
mysql> SELECT 'a xword a' REGEXP '[:<:]word[:>:]'; -> 0
```

To use a literal instance of a special character in a regular expression, precede it by two backslash (`\`) characters. The MySQL parser interprets one of the backslashes, and the regular expression library interprets the other. For example, to match the string `1+2` that contains the special `+` character, only the last of the following regular expressions is the correct one:

```
mysql> SELECT '1+2' REGEXP '1+2'; -> 0
mysql> SELECT '1+2' REGEXP '1\+2'; -> 0
mysql> SELECT '1+2' REGEXP '1\\+2'; -> 1
```

---

# Appendix H. Limits in MySQL

This Appendix lists current limits in MySQL 4.1.

## H.1. Limits of Joins

In MySQL 4.1, the maximum number of tables that can be referenced in a single join is 61. This also applies to the number of tables that can be referenced in the definition of a view.

---

# Appendix I. Feature Restrictions

The discussion here describes restrictions that apply to the use of MySQL features such as subqueries.

## I.1. Restrictions on Subqueries

- Known bug to be fixed later: If you compare a `NULL` value to a subquery using `ALL`, `ANY`, or `SOME`, and the subquery returns an empty result, the comparison might evaluate to the non-standard result of `NULL` rather than to `TRUE` or `FALSE`. This is to be fixed in MySQL 5.1.
- A subquery's outer statement can be any one of: `SELECT`, `INSERT`, `UPDATE`, `DELETE`, `SET`, or `DO`.
- Subquery optimization for `IN` is not as effective as for the `=` operator or for `IN(value_list)` constructs.

A typical case for poor `IN` subquery performance is when the subquery returns a small number of rows but the outer query returns a large number of rows to be compared to the subquery result.

The problem is that, for a statement that uses an `IN` subquery, the optimizer rewrites it as a correlated subquery. Consider the following statement that uses an uncorrelated subquery:

```
SELECT ... FROM t1 WHERE t1.a IN (SELECT b FROM t2);
```

The optimizer rewrites the statement to a correlated subquery:

```
SELECT ... FROM t1 WHERE EXISTS (SELECT 1 FROM t2 WHERE t2.b = t1.a);
```

If the inner and outer queries return  $M$  and  $N$  rows, respectively, the execution time becomes on the order of  $O(M \times N)$ , rather than  $O(M+N)$  as it would be for an uncorrelated subquery.

An implication is that an `IN` subquery can be much slower than a query written using an `IN(value_list)` construct that lists the same values that the subquery would return.

- In general, you cannot modify a table and select from the same table in a subquery. For example, this limitation applies to statements of the following forms:

```
DELETE FROM t WHERE ... (SELECT ... FROM t ...);  
UPDATE t ... WHERE col = (SELECT ... FROM t ...);  
{INSERT|REPLACE} INTO t (SELECT ... FROM t ...);
```

Exception: The preceding prohibition does not apply if you are using a subquery for the modified table in the `FROM` clause. Example:

```
UPDATE t ... WHERE col = (SELECT (SELECT ... FROM t...) AS _t ...);
```

Here the prohibition does not apply because a subquery in the `FROM` clause is materialized as a temporary table, so the relevant rows in `t` have already been selected by the time the update to `t` takes place.

- Row comparison operations are only partially supported:
  - For `expr IN (subquery)`, `expr` can be an  $n$ -tuple (specified via row constructor syntax) and the subquery can return rows of  $n$ -tuples.

- For `expr op {ALL|ANY|SOME} (subquery)`, `expr` must be a scalar value and the subquery must be a column subquery; it cannot return multiple-column rows.

In other words, for a subquery that returns rows of  $n$ -tuples, this is supported:

```
(val_1, ..., val_n) IN (subquery)
```

But this is not supported:

```
(val_1, ..., val_n) op {ALL|ANY|SOME} (subquery)
```

The reason for supporting row comparisons for `IN` but not for the others is that `IN` is implemented by rewriting it as a sequence of `=` comparisons and `AND` operations. This approach cannot be used for `ALL`, `ANY`, or `SOME`.

- Row constructors are not well optimized. The following two expressions are equivalent, but only the second can be optimized:

```
(col1, col2, ...) = (val1, val2, ...)
col1 = val1 AND col2 = val2 AND ...
```

- Subqueries in the `FROM` clause cannot be correlated subqueries. They are materialized (executed to produce a result set) before evaluating the outer query, so they cannot be evaluated per row of the outer query.
- The optimizer is more mature for joins than for subqueries, so in many cases a statement that uses a subquery can be executed more efficiently if you rewrite it as a join.

An exception occurs for the case where an `IN` subquery can be rewritten as a `SELECT DISTINCT` join. Example:

```
SELECT col FROM t1 WHERE id_col IN (SELECT id_col2 FROM t2 WHERE condition);
```

That statement can be rewritten as follows:

```
SELECT DISTINCT col FROM t1, t2 WHERE t1.id_col = t2.id_col AND condition;
```

But in this case, the join requires an extra `DISTINCT` operation and is not more efficient than the subquery.

- Possible future optimization: MySQL does not rewrite the join order for subquery evaluation. In some cases, a subquery could be executed more efficiently if MySQL rewrote it as a join. This would give the optimizer a chance to choose between more execution plans. For example, it could decide whether to read one table or the other first.

Example:

```
SELECT a FROM outer_table AS ot
WHERE a IN (SELECT a FROM inner_table AS it WHERE ot.b = it.b);
```

For that query, MySQL always scans `outer_table` first and then executes the subquery on `inner_table` for each row. If `outer_table` has a lot of rows and `inner_table` has few rows, the query probably will not be as fast as it could be.

The preceding query could be rewritten like this:

```
SELECT a FROM outer_table AS ot, inner_table AS it
```

```
WHERE ot.a = it.a AND ot.b = it.b;
```

In this case, we can scan the small table (`inner_table`) and look up rows in `outer_table`, which will be fast if there is an index on `(ot.a,ot.b)`.

- Possible future optimization: A correlated subquery is evaluated for each row of the outer query. A better approach is that if the outer row values do not change from the previous row, do not evaluate the subquery again. Instead, use its previous result.
- Possible future optimization: A subquery in the `FROM` clause is evaluated by materializing the result into a temporary table, and this table does not use indexes. This does not allow the use of indexes in comparison with other tables in the query, although that might be useful.
- Possible future optimization: If a subquery in the `FROM` clause resembles a view to which the merge algorithm can be applied, rewrite the query and apply the merge algorithm so that indexes can be used. The following statement contains such a subquery:

```
SELECT * FROM (SELECT * FROM t1 WHERE t1.t1_col) AS _t1, t2 WHERE t2.t2_col;
```

The statement can be rewritten as a join like this:

```
SELECT * FROM t1, t2 WHERE t1.t1_col AND t2.t2_col;
```

This type of rewriting would provide two benefits:

- It avoids the use of a temporary table for which no indexes can be used. In the rewritten query, the optimizer can use indexes on `t1`.
- It gives the optimizer more freedom to choose between different execution plans. For example, rewriting the query as a join allows the optimizer to use `t1` or `t2` first.
- Possible future optimization: For `IN`, `= ANY`, `<> ANY`, `= ALL`, and `<> ALL` with non-correlated subqueries, use an in-memory hash for a result result or a temporary table with an index for larger results. Example:

```
SELECT a FROM big_table AS bt
WHERE non_key_field IN (SELECT non_key_field FROM table WHERE condition)
```

In this case, we could create a temporary table:

```
CREATE TABLE t (key (non_key_field))
(SELECT non_key_field FROM table WHERE condition)
```

Then, for each row in `big_table`, do a key lookup in `t` based on `bt.non_key_field`.

---

# Appendix J. GNU General Public License

Version 2, June 1991

Copyright © 1989, 1991 Free Software Foundation, Inc.  
59 Temple Place - Suite 330, Boston, MA 02111-1307, USA  
Everyone is permitted to copy and distribute verbatim copies  
of this license document, but changing it is not allowed.

## Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change free software---to make sure the software is free for all its users. This General Public License applies to most of the Free Software Foundation's software and to any other program whose authors commit to using it. (Some other Free Software Foundation software is covered by the GNU Library General Public License instead.) You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

We protect your rights with two steps: (1) copyright the software, and (2) offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each author's protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors' reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that redistributors of a free program will individually obtain patent licenses, in effect making the program proprietary. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all.

The precise terms and conditions for copying, distribution and modification follow.

## GNU GENERAL PUBLIC LICENSE TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

1. This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this General Public License. The ``Program'', below, refers to any such program or work, and a ``work based on the Program'' means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with modifications and/or translated into another language. (Hereinafter, translation is included without limitation in the term ``modification".) Each licensee is

addressed as ``you''.

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running the Program is not restricted, and the output from the Program is covered only if its contents constitute a work based on the Program (independent of having been made by running the Program). Whether that is true depends on what the Program does.

2. You may copy and distribute verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and give any other recipients of the Program a copy of this License along with the Program.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

3. You may modify your copy or copies of the Program or any portion of it, thus forming a work based on the Program, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:
  - a. You must cause the modified files to carry prominent notices stating that you changed the files and the date of any change.
  - b. You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.
  - c. If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the most ordinary way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this License. (Exception: if the Program itself is interactive but does not normally print such an announcement, your work based on the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Program, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Program, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Program.

In addition, mere aggregation of another work not based on the Program with the Program (or with a work based on the Program) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

4. You may copy and distribute the Program (or a work based on it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the following:
  - a. Accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,



- b. Accompany it with a written offer, valid for at least three years, to give any third-party, for a charge no more than your cost of physically performing source distribution, a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
- c. Accompany it with the information you received as to the offer to distribute corresponding source code. (This alternative is allowed only for noncommercial distribution and only if you received the program in object code or executable form with such an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for making modifications to it. For an executable work, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the executable. However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

If distribution of executable or object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place counts as distribution of the source code, even though third parties are not compelled to copy the source along with the object code.

- 5. You may not copy, modify, sublicense, or distribute the Program except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense or distribute the Program is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.
- 6. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Program or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Program (or any work based on the Program), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Program or works based on it.
- 7. Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the original licensor to copy, distribute or modify the Program subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to this License.
- 8. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Program at all. For example, if a patent license would not permit royalty-free redistribution of the Program by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the

integrity of the free software distribution system, which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

9. If the distribution and/or use of the Program is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Program under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.
10. The Free Software Foundation may publish revised and/or new versions of the General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies a version number of this License which applies to it and ``any later version'', you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of this License, you may choose any version ever published by the Free Software Foundation.

11. If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

#### NO WARRANTY

12. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM ``AS IS'' WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.
13. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS

**How to Apply These Terms to Your New Programs**

If you develop a new program, and you want it to be of the greatest possible use to the public, the best way to achieve this is to make it free software which everyone can redistribute and change under these terms.

To do so, attach the following notices to the program. It is safest to attach them to the start of each source file to most effectively convey the exclusion of warranty; and each file should have at least the ``copyright" line and a pointer to where the full notice is found.

```
one line to give the program's name and a brief idea of what it does.
Copyright (C) yyyy name of author

This program is free software; you can redistribute it and/or modify
it under the terms of the GNU General Public License as published by
the Free Software Foundation; either version 2 of the License, or
(at your option) any later version.

This program is distributed in the hope that it will be useful,
but WITHOUT ANY WARRANTY; without even the implied warranty of
MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the
GNU General Public License for more details.

You should have received a copy of the GNU General Public License
along with this program; if not, write to the Free Software
Foundation, Inc., 59 Temple Place - Suite 330, Boston, MA 02111-1307, USA.
```

Also add information on how to contact you by electronic and paper mail.

If the program is interactive, make it output a short notice like this when it starts in an interactive mode:

```
Gnomovision version 69, Copyright (C) 19yy name of author
Gnomovision comes with ABSOLUTELY NO WARRANTY; for details type `show w'.
This is free software, and you are welcome to redistribute it
under certain conditions; type `show c' for details.
```

The hypothetical commands 'show w' and 'show c' should show the appropriate parts of the General Public License. Of course, the commands you use may be called something other than 'show w' and 'show c'; they could even be mouse-clicks or menu items---whatever suits your program.

You should also get your employer (if you work as a programmer) or your school, if any, to sign a ``copyright disclaimer" for the program, if necessary. Here is a sample; alter the names:

```
Yoyodyne, Inc., hereby disclaims all copyright interest in the program
`Gnomovision' (which makes passes at compilers) written by James Hacker.

signature of Ty Coon, 1 April 1989
Ty Coon, President of Vice
```

This General Public License does not permit incorporating your program into proprietary programs. If your program is a subroutine library, you may consider it more useful to permit linking proprietary applications with the library. If this is what you want to do, use the GNU Library General Public License instead of this License.

---

# Appendix K. MySQL FLOSS License Exception

The MySQL AB Exception for Free/Libre and Open Source Software-only Applications Using MySQL Client Libraries (the "FLOSS Exception").

*Version 0.5, 30 August 2006*

## Exception Intent

We want specified Free/Libre and Open Source Software ("FLOSS") applications to be able to use specified GPL-licensed MySQL client libraries (the "Program") despite the fact that not all FLOSS licenses are compatible with version 2 of the GNU General Public License (the "GPL").

## Legal Terms and Conditions

As a special exception to the terms and conditions of version 2.0 of the GPL:

1. You are free to distribute a Derivative Work that is formed entirely from the Program and one or more works (each, a "FLOSS Work") licensed under one or more of the licenses listed below in section 1, as long as:
  - a. You obey the GPL in all respects for the Program and the Derivative Work, except for identifiable sections of the Derivative Work which are not derived from the Program, and which can reasonably be considered independent and separate works in themselves,
  - b. all identifiable sections of the Derivative Work which are not derived from the Program, and which can reasonably be considered independent and separate works in themselves,
    - i. are distributed subject to one of the FLOSS licenses listed below, and
    - ii. the object code or executable form of those sections are accompanied by the complete corresponding machine-readable source code for those sections on the same medium and under the same FLOSS license as the corresponding object code or executable forms of those sections, and
  - c. any works which are aggregated with the Program or with a Derivative Work on a volume of a storage or distribution medium in accordance with the GPL, can reasonably be considered independent and separate works in themselves which are not derivatives of either the Program, a Derivative Work or a FLOSS Work.

If the above conditions are not met, then the Program may only be copied, modified, distributed or used under the terms and conditions of the GPL or another valid licensing option from MySQL AB.

## 2. FLOSS License List

License name	Version(s)/Copyright Date
Academic Free License	2.0
Apache Software License	1.0/1.1/2.0
Apple Public Source License	2.0
Artistic license	From Perl 5.8.0
BSD license	"July 22 1999"
Common Public License	1.0
GNU Library or "Lesser" General Public License (LGPL)	2.0/2.1
Jabber Open Source License	1.0

MIT license	—
Mozilla Public License (MPL)	1.0/1.1
Open Software License	2.0
OpenSSL license (with original SSLeay license)	"2003" ("1998")
PHP License	3.0
Python license (CNRI Python License)	—
Python Software Foundation License	2.1.1
Sleepycat License	"1999"
University of Illinois/NCSA Open Source License	—
W3C License	"2001"
X11 License	"2001"
Zlib/libpng License	—
Zope Public License	2.0

Due to the many variants of some of the above licenses, we require that any version follow the 2003 version of the Free Software Foundation's Free Software Definition (<http://www.gnu.org/philosophy/free-sw.html>) or version 1.9 of the Open Source Definition by the Open Source Initiative (<http://www.opensource.org/docs/definition.php>).

3. Definitions

- a. Terms used, but not defined, herein shall have the meaning provided in the GPL.
- b. Derivative Work means a derivative work under copyright law.

4. **Applicability:** This FLOSS Exception applies to all Programs that contain a notice placed by MySQL AB saying that the Program may be distributed under the terms of this FLOSS Exception. If you create or distribute a work which is a Derivative Work of both the Program and any other work licensed under the GPL, then this FLOSS Exception is not available for that work; thus, you must remove the FLOSS Exception notice from that work and comply with the GPL in all respects, including by retaining all GPL notices. You may choose to redistribute a copy of the Program exclusively under the terms of the GPL by removing the FLOSS Exception notice from that copy of the Program, provided that the copy has never been modified by you or any third party.

**Appendix A. Qualified Libraries and Packages**

The following is a non-exhaustive list of libraries and packages which are covered by the FLOSS License Exception. Please note that this appendix is provided merely as an additional service to specific FLOSS projects wishing to simplify licensing information for their users. Compliance with one of the licenses noted under the "FLOSS license list" section remains a prerequisite.

Package Name	Qualifying License and Version
Apache Portable Runtime (APR)	Apache Software License

	2.0
--	-----

---

# Index

## Symbols

! (logical NOT), 624  
!= (not equal), 620  
", 542  
% (modulo), 643  
% (wildcard character), 539  
& (bitwise AND), 676  
&& (logical AND), 624  
( ) (parentheses), 617  
(Control-Z) \Z, 539  
\* (multiplication), 639  
+ (addition), 639  
- (subtraction), 639  
- (unary minus), 639  
--password option, 328  
--with-raid link errors, 102  
-p option, 328  
.my.cnf file, 214, 214, 304, 312, 329, 375  
.mysql\_history file, 497  
.pid (process ID) file, 353  
/ (division), 639  
/etc/passwd, 293, 743  
< (less than), 620  
<<, 203  
<< (left shift), 677  
<= (less than or equal), 620  
<=> (equal to), 620  
<> (not equal), 620  
= (equal), 619  
> (greater than), 620  
>= (greater than or equal), 620  
>> (right shift), 677  
[API] (MySQL Cluster), 968  
[MGM] (MySQL Cluster), 966  
[NDBD] (MySQL Cluster), 959  
[NDBD\_DEFAULT] (MySQL Cluster), 959  
[NDB\_MGMD] (MySQL Cluster), 966  
[SQL] (MySQL Cluster), 968  
\" (double quote), 539  
' (single quote), 539  
\0 (ASCII 0), 539  
\b (backspace), 539  
\n (linefeed), 539  
\n (newline), 539  
\r (carriage return), 539  
\t (tab), 539  
\Z (Control-Z) ASCII 26, 539  
\\ (escape), 539  
^ (bitwise XOR), 676  
\_ (wildcard character), 539  
\_rowid, 710  
`, 542  
| (bitwise OR), 676  
|| (logical OR), 624

~, 677

## A

aborted clients, 1508  
aborted connection, 1508  
ABS(), 640  
access control, 305  
access denied errors, 1499  
access privileges, 296  
account privileges  
  adding, 322  
accounts  
  anonymous user, 121  
  root, 121  
ACID, 33, 838  
ACLs, 296  
ACOS(), 640  
Active Server Pages (ASP), 1228  
ActiveState Perl, 172  
add-drop-database option  
  mysqldump, 520  
add-drop-table option  
  mysqldump, 520  
add-locks option  
  mysqldump, 520  
ADDDATE(), 646  
adding  
  character sets, 356  
  native functions, 1495  
  new account privileges, 322  
  new functions, 1484  
  new user privileges, 322  
  new users, 91, 95  
  procedures, 1496  
  user-defined functions, 1486  
addition (+), 639  
ADDTIME(), 647  
addtodest option  
  mysqlhotcopy, 528  
administration of MySQL Cluster, 977  
AES\_DECRYPT(), 677  
AES\_ENCRYPT(), 677  
age  
  calculating, 187  
alias, 1523  
alias names  
  case sensitivity, 544  
aliases  
  for expressions, 696  
  for tables, 741  
  in GROUP BY clauses, 696  
  names, 541  
  on expressions, 740  
ALL, 744  
all-databases option  
  mysqlcheck, 517  
  mysqldump, 520  
all-in-1 option

- mysqlcheck, 517
  - allow-keywords option
    - mysqldump, 520
  - allow-suspicious-udfs option
    - mysqld, 223, 293
  - allowold option
    - mysqlhotcopy, 528
  - ALTER COLUMN, 700
  - ALTER DATABASE, 698
  - ALTER TABLE, 698, 701, 1528
  - altering
    - database, 698
  - analyze option
    - myisamchk, 483
    - mysqlcheck, 517
  - ANALYZE TABLE, 779
  - AND
    - bitwise, 676
    - logical, 624
  - anonymous user, 121, 122, 306, 309
  - ANSI mode
    - running, 29
  - ansi option
    - mysqld, 223
  - ANSI SQL mode, 273, 275
  - ANSI\_QUOTES SQL mode, 273
  - answering questions
    - etiquette, 23
  - Apache, 208
  - API node (MySQL Cluster)
    - defined, 912
  - API nodes, 974
  - API's
    - list of, 1573
  - APIs, 1060
    - Perl, 1158
  - ArbitrationDelay, 935, 953
  - ArbitrationRank, 934, 953
  - ArbitrationTimeout (MySQL Cluster configuration parameter), 949
  - arbitrator, 1025
  - ARCHIVE storage engine, 828, 905
  - Area(), 1051, 1052
  - argument processing, 1490
  - arithmetic expressions, 639
  - arithmetic functions, 676
  - AS, 740, 745
  - AsBinary(), 1047
  - ASCII(), 627
  - ASIN(), 640
  - AsText(), 1047
  - ATAN(), 640
  - ATAN2(), 640
  - attackers
    - security against, 291
  - AUTO-INCREMENT
    - ODBC, 1224
  - auto-rehash option
    - mysql, 492
  - auto-repair option
    - mysqlcheck, 517
  - autoclose option
    - mysqld\_safe, 281
  - AUTO\_INCREMENT, 203
    - and NULL values, 1523
  - AVG(), 691
- ## B
- backslash
    - escape character, 538
  - backspace (\b), 539
  - backup option
    - myisamchk, 481
    - myisampack, 486
  - BACKUP TABLE, 780
  - BackupDataBufferSize, 995
  - BackupDataBufferSize (MySQL Cluster configuration parameter), 952
  - BackupDataDir, 936
  - BackupLogBufferSize, 952, 996
  - BackupMaxWriteSize, 952, 996
  - BackupMemory, 952, 996
  - backups, 336
    - database, 780
    - in MySQL Cluster, 991, 992, 992, 993, 995
  - backups, troubleshooting
    - in MySQL Cluster, 996
  - BackupWriteSize, 952, 996
  - basedir option
    - mysqld, 223
    - mysqld\_safe, 281
  - batch mode, 196
  - batch option
    - mysql, 492
  - BatchByteSize, 953
  - BatchSize, 954
  - BatchSizePerLocalScan, 942
  - BDB storage engine, 828, 899
  - BDB tables, 33
  - bdb-home option
    - mysqld, 901
  - bdb-lock-detect option
    - mysqld, 901
  - bdb-logdir option
    - mysqld, 901
  - bdb-no-recover option
    - mysqld, 901
  - bdb-no-sync option
    - mysqld, 901
  - bdb-shared-data option
    - mysqld, 901
  - bdb-tmpdir option
    - mysqld, 901
  - BdMPolyFromText(), 1043
  - BdMPolyFromWKB(), 1043
  - BdPolyFromText(), 1043



- BdPolyFromWKB(), 1044
  - BEGIN, 765
  - benchmark suite, 418
  - BENCHMARK(), 681
  - benchmarks, 419
  - BerkeleyDB storage engine, 828, 899
  - BETWEEN ... AND, 621
  - big-tables option
    - mysqld, 223
  - Big5 Chinese character encoding, 1520
  - BIGINT data type, 586
  - BIN(), 627
  - BINARY, 673
  - BINARY data type, 590, 606
  - binary distributions, 48
    - installing, 89
    - on Linux, 139
  - binary log, 363
  - bind-address option
    - mysqld, 223
  - binlog-do-db option
    - mysqld, 364
  - binlog-ignore-db option
    - mysqld, 364
  - BIT data type, 584
  - BitKeeper tree, 98
  - BIT\_AND(), 691
  - BIT\_COUNT, 203
  - BIT\_COUNT(), 677
  - bit\_functions
    - example, 203
  - BIT\_LENGTH(), 627
  - BIT\_OR, 203
  - BIT\_OR(), 691
  - BIT\_XOR(), 691
  - BLACKHOLE storage engine, 828, 907
  - BLOB
    - inserting binary data, 540
    - size, 614
  - BLOB columns
    - default values, 608
    - indexing, 451, 710
  - BLOB data type, 591, 607
  - block-search option
    - myisamchk, 483
  - BOOL data type, 584
  - BOOLEAN data type, 584
  - bootstrap option
    - mysqld, 224
  - Borland Builder 4, 1229
  - Borland C++ compiler, 1159
  - Boundary(), 1049
  - brackets
    - square, 584
  - brief option
    - mysqlaccess, 505
  - buffer sizes
    - client, 1060
    - mysqld server, 464
  - Buffer(), 1054
  - bug reports
    - criteria for, 24
  - bugs
    - known, 1529
    - MySQL Cluster
      - reporting, 1000
    - reporting, 24
  - bugs database, 24
  - bugs.mysql.com, 24
  - building
    - client programs, 1155
- C**
- C API
    - data types, 1065
    - functions, 1069
    - linking problems, 1155
  - C Prepared statement API
    - functions, 1123
  - C++ APIs, 1159
  - C++ Builder, 1229
  - C++ compiler
    - gcc, 96
  - C++ compiler cannot create executables, 102
  - C:\my.cnf file, 375
  - CACHE INDEX, 809
  - caches
    - clearing, 810
  - calculating
    - dates, 187
  - calendar, 661
  - calling sequences for aggregate functions
    - UDF, 1489
  - calling sequences for simple functions
    - UDF, 1488
  - can't create/write to file, 1509
  - carriage return (\r), 539
  - CASE, 625
  - case sensitivity
    - in access checking, 299
    - in identifiers, 544
    - in names, 544
    - in searches, 1520
    - in string comparisons, 636
  - case-sensitivity
    - of database names, 30
    - of table names, 30
  - CAST, 673
  - cast functions, 673
  - cast operators, 673
  - casts, 618, 619, 673
  - CC environment variable, 96, 96, 103, 1886
  - cc1plus problems, 102
  - CEILING(), 641
  - Centroid(), 1052
  - CFLAGS environment variable, 96, 103, 1886

- CHANGE MASTER TO, 816
- ChangeLog, 1575
- changes
  - Cluster, 1796
  - InnoDB, 1771
  - log, 1575
  - MySQL 3.23, 1723
  - MySQL 4.0, 1664
  - MySQL 4.1, 1575
- changes to privileges, 310
- changing
  - column, 700
  - column order, 1528
  - field, 700
  - table, 698, 701, 1528
- changing socket location, 96, 118, 1520
- CHAR data type, 590, 605
- CHAR VARYING data type, 590
- CHAR(), 627
- CHARACTER data type, 590
- character sets, 97, 354
  - adding, 356
- Character sets, 551
- CHARACTER VARYING data type, 590
- character-set-client-handshake option
  - mysqld, 224
- character-set-server option
  - mysqld, 224
- character-sets-dir option
  - myisamchk, 481
  - myisampack, 487
  - mysql, 492
  - mysqladmin, 509
  - mysqlbinlog, 512
  - mysqlcheck, 517
  - mysqld, 224
  - mysqldump, 520
  - mysqlexport, 530
  - mysqlshow, 533
- characters
  - multi-byte, 358
- CHARACTER\_LENGTH(), 628
- CHARSET(), 682
- CHAR\_LENGTH(), 628
- check option
  - myisamchk, 480
  - mysqlcheck, 517
- check options
  - myisamchk, 480
- CHECK TABLE, 780
- check-only-changed option
  - myisamchk, 480
  - mysqlcheck, 517
- checking
  - tables for errors, 344
- CHECKPOINT Events (MySQL Cluster), 987
- checkpoint option
  - mysqlhotcopy, 528
- Checksum, 955
- Checksum (MySQL Cluster), 957, 958
- checksum errors, 147
- CHECKSUM TABLE, 782
- Chinese, 1520
- choosing
  - a MySQL version, 44
- choosing types, 615
- chroot option
  - mysqld, 224
  - mysqlhotcopy, 528
- clearing
  - caches, 810
- client programs, 474
  - building, 1155
- client tools, 1060
- clients
  - debugging, 1882
  - threaded, 1155
- closing
  - tables, 462
- cluster logs, 984, 985
- Clustering, 910
- CLUSTERLOG commands (MySQL Cluster), 985
- CLUSTERLOG STATISTICS command (MySQL Cluster), 990
- COALESCE(), 621
- COERCIBILITY(), 682
- ColdFusion, 1229
- collating
  - strings, 358
- COLLATION(), 683
- collation-server option
  - mysqld, 224
- column
  - changing, 700
  - types, 584
- column comments, 709
- column names
  - case sensitivity, 544
- column-names option
  - mysql, 493
- columns
  - changing, 1528
  - indexes, 451
  - names, 541
  - other types, 615
  - selecting, 185
  - storage requirements, 613
- columns option
  - mysqlexport, 530
- comma-separate values data, reading, 734, 744
- command options
  - mysql, 492
  - MySQL Cluster, 977, 978, 979, 980, 980
  - mysqladmin, 509
  - mysqld, 222
- command syntax, 4

- command-line history
  - mysql, 497
- commands
  - for binary distribution, 89
- commands out of sync, 1510
- Comment syntax, 546
- comments
  - adding, 546
  - starting, 37
- comments option
  - mysqldump, 520
- COMMIT, 33, 765
- commit option
  - mysqlaccess, 505
- compact option
  - mysqldump, 521
- comparison operators, 619
- compatibility
  - between MySQL versions, 125, 132
  - with mSQL, 637
  - with ODBC, 543, 587, 618, 621, 708, 746
  - with Oracle, 30, 693, 762
  - with PostgreSQL, 32
  - with standard SQL, 28
  - with Sybase, 765
- compatible option
  - mysqldump, 521
- compiler
  - C++ gcc, 96
- compiling
  - on Windows, 109
  - optimizing, 463
  - problems, 101
  - speed, 466
  - statically, 96
  - user-defined functions, 1492
- complete-insert option
  - mysqldump, 521
- compliance
  - Y2K, 11
- compress option
  - mysql, 493
  - mysqladmin, 509
  - mysqlcheck, 517
  - mysqldump, 521
  - mysqlimport, 530
  - mysqlshow, 533
- COMPRESS(), 678
- compressed tables, 835
- CONCAT(), 628
- CONCAT\_WS(), 628
- concurrent inserts, 447, 450
- config-file option
  - mysqld\_multi, 285
- config.cache, 101
- config.cache file, 101
- config.ini (MySQL Cluster), 920, 929, 930, 977
- configuration
  - MySQL Cluster, 958
  - configuration files, 312
  - configuration options, 95
  - configure
    - running after prior invocation, 101
  - configure option
    - with-charset, 97
    - with-collation, 97
    - with-extra-charsets, 97
    - with-low-memory, 102
  - configure script, 95
  - configuring backups
    - in MySQL Cluster, 995
  - configuring MySQL Cluster, 915, 926, 974, 977
  - Configuring MySQL Cluster (concepts), 912
  - connecting
    - remotely with SSH, 335
    - to the server, 175, 303
    - verification, 305
  - connection
    - aborted, 1508
  - CONNECTION Events (MySQL Cluster), 986
  - CONNECTION\_ID(), 683
  - Connector/JDBC, 1162
  - Connector/MXJ, 1162
  - Connector/NET, 1162, 1235
    - reporting problems, 1423
  - Connector/ODBC, 1162, 1162
    - reporting problems, 1234
  - Connectors
    - MySQL, 1162
  - connectstring, 932
  - connect\_timeout variable, 496, 511
  - console option
    - mysqld, 224
  - constant table, 422, 430
  - constraints, 38
  - Contains(), 1055
  - contributing companies
    - list of, 1574
  - contributors
    - list of, 1566
  - control access, 305
  - control flow functions, 625
  - CONV(), 628
  - conventions
    - typographical, 2
  - CONVERT, 673
  - CONVERT TO, 701
  - CONVERT\_TZ(), 647
  - ConvexHull(), 1054
  - copy option
    - mysqlaccess, 505
  - copying databases, 137
  - copying tables, 714
  - core-file option
    - mysqld, 224
  - core-file-size option

- 
- mysqld\_safe, 281
  - correct-checksum option
    - myisamchk, 481
  - COS(), 641
  - COT(), 641
  - count option
    - myisam\_ftdump, 476
    - mysqldadmin, 510
  - COUNT(), 691
  - COUNT(DISTINCT), 692
  - counting
    - table rows, 192
  - crash, 1877
    - recovery, 343
    - repeated, 1516
  - crash-me, 419
  - crash-me program, 417, 418
  - CRC32(), 641
  - CREATE DATABASE, 704
  - CREATE FUNCTION, 1485
  - CREATE INDEX, 704
  - CREATE TABLE, 706
  - create-options option
    - mysqldump, 521
  - creating
    - bug reports, 24
    - database, 704
    - databases, 179
    - default startup options, 214
    - function, 1485
    - tables, 181
  - CROSS JOIN, 745
  - Crosses(), 1055
  - CR\_SERVER\_GONE\_ERROR, 1505
  - CR\_SERVER\_LOST\_ERROR, 1505
  - CSV data, reading, 734, 744
  - CSV storage engine, 828, 906
  - CURDATE(), 647
  - CURRENT\_DATE, 647
  - CURRENT\_TIME, 648
  - CURRENT\_TIMESTAMP, 648
  - CURRENT\_USER(), 683
  - CURTIME(), 647
  - customers
    - of MySQL, 418
  - CXX environment variable, 96, 96, 102, 102, 103, 1886
  - CXXFLAGS environment variable, 96, 103, 1886
- D**
- data
    - character sets, 354
    - loading into tables, 182
    - retrieving, 183
    - size, 450
  - data node (MySQL Cluster)
    - defined, 912
  - data nodes (MySQL Cluster), 975
  - data type
    - BIGINT, 586
    - BINARY, 590, 606
    - BIT, 584
    - BLOB, 591, 607
    - BOOL, 584, 615
    - BOOLEAN, 584, 615
    - CHAR, 590, 605
    - CHAR VARYING, 590
    - CHARACTER, 590
    - CHARACTER VARYING, 590
    - DATE, 588, 596
    - DATETIME, 588, 596
    - DEC, 587
    - DECIMAL, 587
    - DOUBLE, 586
    - DOUBLE PRECISION, 587
    - ENUM, 591, 609
    - FIXED, 587
    - FLOAT, 586, 586, 587
    - GEOMETRY, 1041
    - GEOMETRYCOLLECTION, 1041
    - INT, 585
    - INTEGER, 586
    - LINestring, 1041
    - LONG, 607
    - LOB, 591
    - LONGTEXT, 591
    - MEDIUMBLOB, 591
    - MEDIUMINT, 585
    - MEDIUMTEXT, 591
    - MULTILINESTRING, 1041
    - MULTIPOINT, 1041
    - MULTIPOLYGON, 1041
    - NATIONAL CHAR, 590
    - NCHAR, 590
    - NUMERIC, 587
    - POINT, 1041
    - POLYGON, 1041
    - REAL, 587
    - SET, 591, 611
    - SMALLINT, 585
    - TEXT, 591, 607
    - TIME, 588, 603
    - TIMESTAMP, 588, 596
    - TINYBLOB, 591
    - TINYINT, 584
    - TINYTEXT, 591
    - VARBINARY, 591, 606
    - VARCHAR, 590, 605
    - VARCHARACTER, 590
    - YEAR, 588, 604
  - data types, 584
    - C API, 1065
    - overview, 584
  - data-file-length option
    - myisamchk, 481
  - database
    - altering, 698
-

- creating, 704
- deleting, 717
- database design, 450
- Database information
  - obtaining, 790
- database names
  - case sensitivity, 544
  - case-sensitivity, 30
- database option
  - mysql, 493
  - mysqlbinlog, 512
- DATABASE(), 683
- databases
  - backups, 336
  - copying, 137
  - creating, 179
  - defined, 5
  - information about, 196
  - names, 541
  - replicating, 381
  - selecting, 180
  - symbolic links, 471
  - using, 179
- databases option
  - mysqlcheck, 517
  - mysqldump, 521
- DataDir, 935, 936
- datadir option
  - mysqld, 224
  - mysqld\_safe, 281
- DataJunction, 1230
- DataMemory, 937, 968
- DATE, 1521
- date and time functions, 646
- Date and Time types, 595
- date calculations, 187
- DATE columns
  - problems, 1521
- DATE data type, 588, 596
- date functions
  - Y2K compliance, 11
- date option
  - mysql\_explain\_log, 504
- date types, 613
  - Y2K issues, 604
- date values
  - problems, 597
- DATE(), 648
- DATEDIFF(), 648
- DATETIME data type, 588, 596
- DATE\_ADD(), 648
- DATE\_FORMAT(), 650
- DATE\_SUB(), 648
- DAY(), 651
- DAYNAME(), 651
- DAYOFMONTH(), 651
- DAYOFWEEK(), 652
- DAYOFYEAR(), 652
- db option
  - mysqlaccess, 505
- db table
  - sorting, 309
- DB2 SQL mode, 275
- DBI interface, 1158
- DBI->quote, 540
- DBI->trace, 1880
- DBI/DBD interface, 1158
- DBI\_TRACE environment variable, 1880, 1886
- DBI\_USER environment variable, 1886
- DEBUG package, 1882
- DEALLOCATE PREPARE, 825, 826
- debug option
  - myisamchk, 478
  - myisampack, 487
  - mysql, 493
  - mysqlaccess, 505
  - mysqladmin, 510
  - mysqlbinlog, 512
  - mysqlcheck, 517
  - mysqld, 224
  - mysqldump, 521
  - mysqlhotcopy, 528
  - mysqlimport, 530
  - mysqlshow, 533
- debug-info option
  - mysql, 493
- debugging
  - client, 1882
  - server, 1877
- debugging support, 95
- DEC data type, 587
- DECIMAL data type, 587
- decimal point, 584
- DECODE(), 678
- decode\_bits myisamchk variable, 479
- DEFAULT
  - constraint, 39
- default
  - privileges, 121
- default hostname, 303
- default installation location, 57
- default options, 214
- DEFAULT value clause, 592, 709
- default values, 416, 592, 709, 725
  - BLOB and TEXT columns, 608
  - explicit, 592
  - implicit, 592
  - suppression, 39
- DEFAULT(), 687
- default-character-set option
  - mysql, 493
  - mysqladmin, 510
  - mysqlcheck, 517
  - mysqld, 225
  - mysqldump, 521
  - mysqlimport, 530

- mysqlshow, 533
- default-collation option
  - mysql, 225
- default-storage-engine option
  - mysql, 225
- default-table-type option
  - mysql, 225
- default-time-zone option
  - mysql, 225
- defaults
  - embedded, 1061
- defaults-extra-file option
  - mysql\_safe, 281
- defaults-file option
  - mysql\_safe, 282
- DEGREES(), 641
- delay-key-write option
  - mysql, 225, 832
- delay-key-write-for-all-tables option
  - mysql, 225
- DELAYED, 728
- delayed-insert option
  - mysqldump, 521
- delayed\_insert\_limit, 729
- DELETE, 719
- delete option
  - mysqlimport, 530
- delete-master-logs option
  - mysqldump, 522
- deleting
  - database, 717
  - foreign key, 701, 857
  - function, 1486
  - index, 700, 718
  - primary key, 700
  - rows, 1524
  - table, 718
  - user, 325, 771
  - users, 325, 771
- deletion
  - mysql.sock, 1519
- delimiter option
  - mysql, 493
- Delphi, 1229
- derived tables, 754
- des-key-file option
  - mysql, 225
- DESC, 761
- DESCRIBE, 196, 761
- description option
  - myisamchk, 483
- design
  - choices, 450
  - issues, 1529
  - limitations, 416
- DES\_DECRYPT(), 679
- DES\_ENCRYPT(), 679
- developers
  - list of, 1561
  - development source tree, 98
- Difference(), 1054
- digits, 584
- Dimension(), 1048
- directory structure
  - default, 57
- disable-keys option
  - mysqldump, 522
- disable-log-bin option
  - mysqlbinlog, 512
- DISCARD TABLESPACE, 702, 844
- disconnecting
  - from the server, 175
- Disjoint(), 1056
- disk full, 1518
- disk issues, 469
- Diskless, 945
- disks
  - splitting data across, 473
- display size, 584
- display width, 584
- displaying
  - information
    - Cardinality, 796
    - Collation, 796
    - SHOW, 790, 792, 795, 797, 806
    - table status, 804
- Distance(), 1056
- DISTINCT, 185, 435, 744
  - COUNT(), 692
- DISTINCTROW, 744
- DIV, 639
- division (/), 639
- DNS, 469
- DO, 722
- DocBook XML
  - documentation source format, 2
- Documenters
  - list of, 1570
- DOUBLE data type, 586
- DOUBLE PRECISION data type, 587
- double quote ("), 539
- downgrades
  - MySQL Cluster, 970, 970, 971
- downgrading, 138
- downloading, 54
- DROP DATABASE, 717
- DROP FOREIGN KEY, 701, 857
- DROP FUNCTION, 1486
- DROP INDEX, 700, 718
- DROP PREPARE, 826
- DROP PRIMARY KEY, 700
- DROP TABLE, 718
- DROP USER, 771
- dropping
  - user, 325, 771
- dryrun option

- mysqlhotcopy, 528
  - DUAL, 740
  - dump option
    - mysam\_ftdump, 477
  - DUMPFIL, 744
  - dynamic table characteristics, 834
- ## E
- Eiffel Wrapper, 1159
  - ELT(), 629
  - email lists, 20
  - embedded MySQL server library, 1060
  - enable-named-pipe option
    - mysqld, 225
  - ENCODE(), 678
  - ENCRYPT(), 680
  - encryption, 330
  - encryption functions, 677
  - EndPoint(), 1049
  - ENTER SINGLE USER MODE command (MySQL Cluster), 984
  - entering
    - queries, 176
  - ENUM
    - size, 615
  - ENUM data type, 591, 609
  - Envelope(), 1048
  - environment variable
    - CC, 96, 96, 103
    - CFLAGS, 96, 103
    - CXX, 96, 96, 103
    - CXXFLAGS, 96, 103
    - HOME, 497
    - LD\_RUN\_PATH, 142, 149
    - MYSQL\_DEBUG, 475
    - MYSQL\_HISTFILE, 497
    - MYSQL\_HOST, 304
    - MYSQL\_PWD, 304, 475
    - MYSQL\_TCP\_PORT, 374, 374, 475
    - MYSQL\_UNIX\_PORT, 374, 374, 475
    - PATH, 90, 210
    - USER, 304
  - Environment variable
    - CC, 1886
    - CFLAGS, 1886
    - CXX, 102, 1886
    - CXXFLAGS, 1886
    - DBI\_TRACE, 1880, 1886
    - DBI\_USER, 1886
    - HOME, 1886
    - LD\_LIBRARY\_PATH, 173
    - LD\_RUN\_PATH, 173, 1886
    - MYSQL\_DEBUG, 1882, 1886
    - MYSQL\_GROUP\_SUFFIX, 1886
    - MYSQL\_HISTFILE, 1886
    - MYSQL\_HOME, 1886
    - MYSQL\_HOST, 1886
    - MYSQL\_PS1, 1886
    - MYSQL\_PWD, 1886
    - MYSQL\_TCP\_PORT, 1886
    - MYSQL\_UNIX\_PORT, 115, 1886
    - PATH, 1886
    - TMPDIR, 115, 1886
    - TZ, 1520, 1886
    - UMASK, 1513, 1886
    - UMASK\_DIR, 1513, 1886
    - USER, 1886
  - Environment variables
    - CXX, 102
  - environment variables, 218, 312, 475
    - list of, 1886
  - equal (=), 619
  - Equals(), 1056
  - ERROR Events (MySQL Cluster), 989
  - error logs (MySQL Cluster), 975
  - error messages
    - can't find file, 1513
    - languages, 356, 356
  - errors
    - access denied, 1499
    - checking tables for, 344
    - common, 1498
    - directory checksum, 147
    - handling for UDFs, 1492
    - known, 1529
    - linking, 1512
    - list of, 1499
    - reporting, 1, 24, 24
  - escape (\), 539
  - escape characters, 538
  - estimating
    - query performance, 428
  - event log format (MySQL Cluster), 986
  - event logs (MySQL Cluster), 984, 985, 985
  - event severity levels (MySQL Cluster), 986
  - event types (MySQL Cluster), 985, 986
  - example option
    - mysqld\_multi, 285
  - EXAMPLE storage engine, 828, 905
  - examples
    - compressed tables, 487
    - mysamchk output, 348
    - queries, 198
  - EXECUTE, 825
  - execute option
    - mysql, 493
  - EXECUTE statement, 826
  - ExecuteOnComputer, 933, 935, 953
  - EXIT command (MySQL Cluster), 984
  - EXIT SINGLE USER MODE command (MySQL Cluster), 984
  - exit-info option
    - mysqld, 225
  - EXP(), 641
  - EXPLAIN, 420
  - explicit default values, 592

- EXPORT\_SET(), 629
- expression aliases, 696, 740
- expressions
  - extended, 190
- extend-check option
  - myisamchk, 480, 481
- extended option
  - mysqlcheck, 517
- extended-insert option
  - mysqldump, 522
- extensions
  - to standard SQL, 28
- ExteriorRing(), 1051
- external locking, 226, 230, 253, 343, 802
- external-locking option
  - mysqld, 226
- EXTRACT(), 652
- extracting
  - dates, 187
- F**
- FALSE, 540, 541
- fast option
  - myisamchk, 481
  - mysqlcheck, 517
- fatal signal 11, 102
- features of MySQL, 6
- field
  - changing, 700
- FIELD(), 629
- FILE, 631
- files
  - binary log, 363
  - config.cache, 101
  - error messages, 356
  - general query log, 362
  - log, 95, 367
  - my.cnf, 392
  - not found message, 1513
  - permissions, 1513
  - repairing, 481
  - script, 196
  - size limits, 10
  - slow query log, 367
  - tmp, 115
  - update log, 362
- filesort optimization, 437
- FileSystemPath, 936
- FIND\_IN\_SET(), 629
- FIXED data type, 587
- FLOAT data type, 586, 586, 587
- floating-point number, 587
- floats, 540
- FLOOR(), 641
- FLUSH, 810
- flush option
  - mysqld, 226
- flush tables, 509
  - flush-logs option
    - mysqldump, 522
  - flushlog option
    - mysqlhotcopy, 528
  - FOR UPDATE, 744
  - FORCE INDEX, 740, 746, 1527
  - FORCE KEY, 740, 747
  - force option
    - myisamchk, 481, 482
    - myisampack, 487
    - mysql, 493
    - mysqladmin, 510
    - mysqlcheck, 518
    - mysqldump, 522
    - mysqlimport, 531
  - force-read option
    - mysqlbinlog, 512
  - foreign key
    - constraint, 38
    - deleting, 701, 857
  - foreign keys, 35, 201, 701
  - FORMAT(), 629
  - Forums, 23
  - FOUND\_ROWS(), 684
  - FreeBSD troubleshooting, 103
  - frequently-asked questions about MySQL Cluster, 1020
  - FROM, 740
  - FROM\_DAYS(), 652
  - FROM\_UNIXTIME(), 652
  - ft\_max\_word\_len myisamchk variable, 479
  - ft\_min\_word\_len myisamchk variable, 479
  - ft\_stopword\_file myisamchk variable, 479
  - full disk, 1518
  - full-text search, 662
  - fulltext
    - stopword list, 671
  - function
    - creating, 1485
    - deleting, 1486
  - functions, 617
    - arithmetic, 676
    - bit, 676
    - C API, 1069
    - C Prepared statement API, 1123
    - cast, 673
    - control flow, 625
    - date and time, 646
    - encryption, 677
    - GROUP BY, 690
    - grouping, 617
    - information, 681
    - mathematical, 640
    - miscellaneous, 687
    - native
      - adding, 1495
      - new, 1484
      - string, 627
      - string comparison, 636



user-defined, 1484  
 adding, 1486  
 Functions  
 user-defined, 1485, 1486  
 functions for SELECT and WHERE clauses, 617

**G**

gcc, 96  
 gdb  
 using, 1878  
 general information, 1  
 General Public License, 5  
 general query log, 362  
 geographic feature, 1032  
 GeomCollFromText(), 1042  
 GeomCollFromWKB(), 1043  
 geometry, 1033  
 GEOMETRY data type, 1041  
 GEOMETRYCOLLECTION data type, 1041  
 GeometryCollection(), 1044  
 GeometryCollectionFromText(), 1042  
 GeometryCollectionFromWKB(), 1043  
 GeometryFromText(), 1042  
 GeometryFromWKB(), 1043  
 GeometryN(), 1053  
 GeometryType(), 1048  
 GeomFromText(), 1042, 1047  
 GeomFromWKB(), 1043, 1047  
 geospatial feature, 1032  
 getting MySQL, 54  
 GET\_FORMAT(), 653  
 GET\_LOCK(), 687  
 GIS, 1032, 1032  
 GLength(), 1050, 1051  
 global privileges, 771, 778  
 goals of MySQL, 6  
 GPL  
 General Public License, 1895  
 GNU General Public License, 1895  
 MySQL FLOSS License Exception, 1900  
 GRANT, 771  
 GRANT statement, 322  
 grant tables, 310  
 re-creating, 115  
 sorting, 307, 309  
 granting  
 privileges, 771  
 GRANTS, 795  
 greater than (>), 620  
 greater than or equal (>=), 620  
 GREATEST(), 621  
 GROUP BY, 439  
 aliases in, 696  
 extensions to standard SQL, 696, 741  
 GROUP BY functions, 690  
 grouping  
 expressions, 617  
 GROUP\_CONCAT(), 692

**H**

HANDLER, 722  
 handling  
 errors, 1492  
 HAVING, 742  
 HEAP storage engine, 828, 897  
 HeartbeatIntervalDbApi (MySQL Cluster configuration parameter), 946  
 HeartbeatIntervalDbDb (MySQL Cluster configuration parameter), 946  
 HELP command (MySQL Cluster), 983  
 help option  
 myisamchk, 478  
 myisampack, 486  
 myisam\_ftdump, 476  
 mysql, 492  
 mysqlaccess, 505  
 mysqladmin, 509  
 mysqlbinlog, 512  
 mysqlcheck, 516  
 mysqld, 223  
 mysqldump, 520  
 mysqld\_multi, 285  
 mysqlhotcopy, 528  
 mysqlimport, 530  
 mysqlshow, 533  
 perror, 536  
 HELP statement, 762  
 HEX(), 630  
 hex-blob option  
 mysqldump, 522  
 hexadecimal values, 540  
 HIGH\_PRIORITY, 744  
 hints, 29, 744, 746, 746  
 index, 741, 746  
 history of MySQL, 6  
 HOME environment variable, 497, 1886  
 host option  
 mysql, 493  
 mysqlaccess, 505  
 mysqladmin, 510  
 mysqlbinlog, 513  
 mysqlcheck, 518  
 mysqldump, 522  
 mysqlhotcopy, 528  
 mysqlimport, 531  
 mysqlshow, 533  
 mysql\_explain\_log, 504  
 host table, 310  
 sorting, 309  
 Host\*SciId\* parameters, 957  
 host.frm  
 problems finding, 112  
 hostname  
 default, 303  
 HostName, 934, 936, 953  
 hostname caching, 469

- hour(), 653
- howto option
  - mysqlaccess, 505
- html option
  - mysql, 493
- I**
- i-am-a-dummy option
  - mysql, 495
- id, 933, 935, 953
- ID
  - unique, 1154
- identifiers, 541
  - case sensitivity, 544
  - quoting, 542
- IF(), 625
- IFNULL(), 626
- IGNORE INDEX, 740, 746
- IGNORE KEY, 740, 747
- ignore option
  - mysqlimport, 531
- ignore-lines option
  - mysqlimport, 531
- ignore-spaces option
  - mysql, 493
- ignore-table option
  - mysqldump, 522
- IGNORE\_SPACE SQL mode, 273
- implicit default values, 592
- IMPORT TABLESPACE, 702, 844
- IN, 622
- increasing
  - performance, 410
- increasing with replication
  - speed, 381
- index
  - deleting, 700, 718
- index hints, 741, 746
- indexes, 704
  - and BLOB columns, 451, 710
  - and IS NULL, 454
  - and LIKE, 454
  - and NULL values, 710
  - and TEXT columns, 451, 710
  - assigning to key cache, 809
  - block size, 242
  - columns, 451
  - leftmost prefix of, 454
  - multi-column, 452
  - multiple-part, 704
  - names, 541
  - use of, 453
- IndexMemory, 938, 968
- INET\_ATON(), 688
- INET\_NTOA(), 688
- INFO Events (MySQL Cluster), 989
- information functions, 681
- information option
  - myisamchk, 481
- init-file option
  - mysqld, 226
- INNER JOIN, 745
- InnoDB, 838
  - Solaris 10 x86\_64 issues, 147
- innodb option
  - mysqld, 846
- InnoDB storage engine, 828, 838
- InnoDB tables, 33
- innodb-safe-binlog option
  - mysqld, 226
- innodb\_status\_file option
  - mysqld, 846
- INSERT, 441, 724
- INSERT ... SELECT, 727
- INSERT DELAYED, 728, 728
- INSERT statement
  - grant privileges, 323
- INSERT(), 630
- insert-ignore option
  - mysqldump, 523
- inserting
  - speed of, 441
- inserts
  - concurrent, 447, 450
- installation layouts, 57
- installation overview, 92
- installing
  - binary distribution, 89
  - Linux RPM packages, 81
  - Mac OS X PKG packages, 84
  - overview, 41
  - Perl, 170
  - Perl on Windows, 172
  - Solaris PKG packages, 86
  - source distribution, 92
  - user-defined functions, 1492
- installing MySQL Cluster, 915, 918, 927
- INSTR(), 630
- INT data type, 585
- INTEGER data type, 586
- integers, 540
- InteriorRingN(), 1052
- internal compiler errors, 102
- internal locking, 446
- internals, 1483
- Internet Relay Chat, 23
- Intersection(), 1054
- Intersects(), 1056
- INTERVAL(), 623
- introducer
  - string literal, 538, 556
- invalid data
  - constraint, 39
- IRC, 23
- IS NOT NULL, 620
- IS NULL, 434, 620

- and indexes, 454
  - ISAM storage engine, 828, 908
  - isamchk, 474
  - isamlog, 474
  - IsClosed(), 1051
  - IsEmpty(), 1049
  - ISNULL(), 622
  - ISOLATION LEVEL, 770
  - IsRing(), 1050
  - IsSimple(), 1049
  - IS\_FREE\_LOCK(), 689
  - IS\_USED\_LOCK(), 689
- J**
- JOIN, 745
  - join option
    - myisampack, 487
- K**
- keepold option
    - mysqlhotcopy, 528
  - Key cache
    - MyISAM, 455
  - key cache
    - assigning indexes to, 809
  - key space
    - MyISAM, 833
  - keys, 451
    - foreign, 35, 201
    - multi-column, 452
    - searching on two, 202
  - keys option
    - mysqlshow, 534
  - keys-used option
    - myisamchk, 482
  - keywords, 547
  - key\_buffer\_size myisamchk variable, 479
  - KILL, 811
  - known errors, 1529
- L**
- language option
    - mysqld, 226
  - language support
    - error messages, 356
  - last row
    - unique ID, 1154
  - LAST\_DAY(), 654
  - LAST\_INSERT\_ID(), 35, 727
  - LAST\_INSERT\_ID([expr]), 685
  - layout of installation, 57
  - LCASE(), 630
  - LD\_LIBRARY\_PATH environment variable, 173
  - LD\_RUN\_PATH environment variable, 142, 149, 173, 1886
  - LEAST(), 623
  - ledir option
    - mysqld\_safe, 282
  - LEFT JOIN, 435, 745
  - LEFT OUTER JOIN, 745
  - LEFT(), 630
  - leftmost prefix of indexes, 454
  - legal names, 541
  - length option
    - myisam\_ftdump, 477
  - LENGTH(), 631
  - less than (<), 620
  - less than or equal (<=), 620
  - libmysqld, 1060
    - options, 1061
  - libraries
    - list of, 1572
  - library
    - mysqlclient, 1060
    - mysqld, 1060
  - License, 1900
  - LIKE, 636
    - and indexes, 454
    - and wildcards, 454
  - LIMIT, 440, 684, 742
  - limitations
    - design, 416
    - MySQL Limitations, 1891
    - replication, 392
    - limitations of MySQL Cluster, 1014
  - limits
    - file-size, 10
    - MySQL Limits, limits in MySQL, 1891
  - line-numbers option
    - mysql, 493
  - linefeed (\n), 539
  - LineFromText(), 1042
  - LineFromWKB(), 1043
  - LINESTRING data type, 1041
  - LineString(), 1044
  - LineStringFromText(), 1042
  - LineStringFromWKB(), 1043
  - linking, 1155
    - errors, 1512
    - problems, 1155
    - speed, 466
  - links
    - symbolic, 470
  - Linux
    - binary distribution, 139
    - source distribution, 140
  - literals, 538
  - LN(), 642
  - LOAD DATA FROM MASTER, 818
  - LOAD DATA INFILE, 730, 1523
  - LOAD TABLE FROM MASTER, 819
  - loading
    - tables, 182
  - LOAD\_FILE(), 631
  - local checkpoints (MySQL Cluster), 968

- 
- local option
    - mysqlexport, 531
  - local-infile option
    - mysql, 493
    - mysqld, 293
  - local-load option
    - mysqbinlog, 513
  - LOCALTIME, 654
  - LOCALTIMESTAMP, 654
  - LOCATE(), 631
  - LOCK IN SHARE MODE, 744
  - LOCK TABLES, 767
  - lock-all-tables option
    - mysqldump, 523
  - lock-tables option
    - mysqldump, 523
    - mysqlexport, 531
  - locking, 464
    - external, 226, 230, 253, 343, 802
    - page-level, 446
    - row-level, 35, 446
    - table-level, 446
  - locking methods, 446
  - LockPagesInMainMemory, 945
  - log
    - changes, 1575
  - log files, 95, 361
    - maintaining, 367
    - names, 336
  - log files (MySQL Cluster), 975
  - log option
    - mysqld, 226
    - mysqld\_multi, 285
  - LOG(), 642
  - log-bin option
    - mysqld, 226
  - log-bin-index option
    - mysqld, 227
  - log-error option
    - mysqld, 227
    - mysqld\_safe, 282
  - log-isam option
    - mysqld, 227
  - log-long-format option
    - mysqld, 227
  - log-queries-not-using-indexes option
    - mysqld, 227
  - log-short-format option
    - mysqld, 227
  - log-slave-updates option
    - mysqld, 399
  - log-slow-admin-statements option
    - mysqld, 227
  - log-slow-queries option
    - mysqld, 227
  - log-update option
    - mysqld, 227
  - log-warnings option
    - mysqld, 228, 399
  - LOG10(), 642
  - LOG2(), 642
  - LogDestination, 934
  - logging commands (MySQL Cluster), 985
  - logical operators, 623
  - LogLevelCheckpoint (MySQL Cluster configuration parameter), 951
  - LogLevelConnection (MySQL Cluster configuration parameter), 951
  - LogLevelError, 951
  - LogLevelInfo, 951
  - LogLevelNodeRestart (MySQL Cluster configuration parameter), 951
  - LogLevelShutdown (MySQL Cluster configuration parameter), 951
  - LogLevelStartup (MySQL Cluster configuration parameter), 951
  - LogLevelStatistic (MySQL Cluster configuration parameter), 951
  - LONG data type, 607
  - LONGBLOB data type, 591
  - LongMessageBuffer, 942
  - LONGTEXT data type, 591
  - low-priority option
    - mysqlexport, 531
  - low-priority-updates option
    - mysqld, 228
  - LOWER(), 631
  - LPAD(), 631
  - LTRIM(), 632
- ## M
- Mac OS X, 1162
    - installation, 84
  - mailing list address, 1
  - mailing lists, 20
    - archive location, 20
    - guidelines, 23
  - main features of MySQL, 6
  - maintaining
    - log files, 367
    - tables, 353
  - MAKEDATE(), 654
  - MAKETIME(), 654
  - make\_binary\_distribution, 222
  - MAKE\_SET(), 632
  - management client (MySQL Cluster), 977
  - management node (MySQL Cluster)
    - defined, 912
  - management nodes (MySQL Cluster), 976
  - managing MySQL Cluster, 980
  - managing MySQL Cluster processes, 973
  - manual
    - available formats, 2
    - online location, 1
    - typographical conventions, 2
  - master-connect-retry option
-

- mysqld, 400
- master-data option
  - mysqldump, 523
- master-host option
  - mysqld, 400
- master-info-file option
  - mysqld, 400
- master-password option
  - mysqld, 400
- master-port option
  - mysqld, 400
- master-retry-count option
  - mysqld, 400
- master-ssl option
  - mysqld, 400
- master-ssl-ca option
  - mysqld, 400
- master-ssl-capath option
  - mysqld, 400
- master-ssl-cert option
  - mysqld, 400
- master-ssl-cipher option
  - mysqld, 400
- master-ssl-key option
  - mysqld, 400
- master-user option
  - mysqld, 400
- master/slave setup, 381
- MASTER\_POS\_WAIT(), 689, 819
- MATCH ... AGAINST(), 662
- matching
  - patterns, 190
- mathematical functions, 640
- MAX(), 693
- max-record-length option
  - myisamchk, 482
- max-relay-log-size option
  - mysqld, 401
- MAXDB SQL mode, 275
- maximum memory used, 509
- maximums
  - maximum tables per join, 1891
- MaxNoOfAttributes, 943
- MaxNoOfConcurrentIndexOperations, 941
- MaxNoOfConcurrentOperations, 940
- MaxNoOfConcurrentScans, 942
- MaxNoOfConcurrentTransactions, 939
- MaxNoOfFiredTriggers, 941
- MaxNoOfIndexes, 944
- MaxNoOfLocalOperations, 940
- MaxNoOfLocalScans, 942
- MaxNoOfOpenFiles, 943
- MaxNoOfOrderedIndexes, 944
- MaxNoOfSavedMessages, 943
- MaxNoOfTables, 944
- MaxNoOfTriggers, 944
- MaxNoOfUniqueHashIndexes, 944
- MaxScanBatchSize, 954
- max\_allowed\_packet variable, 496
- MAX\_CONNECTIONS\_PER\_HOUR, 325
- max\_join\_size variable, 496
- MAX\_QUERIES\_PER\_HOUR, 325
- MAX\_UPDATES\_PER\_HOUR, 325
- MAX\_USER\_CONNECTIONS, 325
- MBR, 1054
- MBRContains(), 1054
- MBRDisjoint(), 1054
- MBREqual(), 1055
- MBRIntersects(), 1055
- MBROverlaps(), 1055
- MBRTouches(), 1055
- MBRWithin(), 1055
- MD5(), 680
- medium-check option
  - myisamchk, 481
  - mysqlcheck, 518
- MEDIUMBLOB data type, 591
- MEDIUMINT data type, 585
- MEDIUMTEXT data type, 591
- memlock option
  - mysqld, 228
- MEMORY storage engine, 828, 897
- memory usage
  - myisamchk, 484
- memory use, 467, 509
- MERGE storage engine, 828, 894
- MERGE tables
  - defined, 894
- method option
  - mysqlhotcopy, 529
- methods
  - locking, 446
- mgm (MySQL Cluster process), 977
- mgmd (MySQL Cluster process), 976
- mgmd (MySQL Cluster)
  - defined, 912
- MICROSECOND(), 654
- Microsoft Access, 1225
- Microsoft ADO, 1227
- Microsoft Excel, 1227
- Microsoft Visual Basic, 1227
- Microsoft Visual InterDev, 1227
- MID(), 632
- MIN(), 693
- Minimum Bounding Rectangle, 1054
- minus
  - unary (-), 639
- MINUTE(), 655
- mirror sites, 54
- miscellaneous functions, 687
- MIT-pthreads, 104
- MLineFromText(), 1042
- MLineFromWKB(), 1043
- MOD (modulo), 643
- MOD(), 643
- modes

- batch, 196
- modules
  - list of, 9
- modulo (%), 643
- modulo (MOD), 643
- monitor
  - terminal, 175
- Mono, 1235
- MONTH(), 655
- MONTHNAME(), 655
- MPointFromText(), 1042
- MPointFromWKB(), 1043
- MPolyFromText(), 1042
- MPolyFromWKB(), 1043
- mSQL compatibility, 637
- MSSQL SQL mode, 275
- multi-byte character sets, 1511
- multi-byte characters, 358
- multi-column indexes, 452
- MULTILINESTRING data type, 1041
- MultiLineString(), 1044
- MultiLineStringFromText(), 1042
- MultiLineStringFromWKB(), 1043
- multiple servers, 368
- multiple-part index, 704
- multiplication (\*), 639
- MULTIPOINT data type, 1041
- MultiPoint(), 1044
- MultiPointFromText(), 1042
- MultiPointFromWKB(), 1043
- MULTIPOLYGON data type, 1041
- MultiPolygon(), 1044
- MultiPolygonFromText(), 1042
- MultiPolygonFromWKB(), 1043
- My
  - derivation, 6
- my.cnf
  - and MySQL Cluster, 920, 929, 930
  - in MySQL Cluster, 974
- my.cnf file, 392
- MyISAM
  - compressed tables, 835
  - size, 613
- MyISAM key cache, 455
- MyISAM storage engine, 828, 830
- mysam-recover option
  - mysqld, 228, 832
- mysamchk, 97, 474
  - analyze option, 483
  - backup option, 481
  - block-search option, 483
  - character-sets-dir option, 481
  - check option, 480
  - check-only-changed option, 480
  - correct-checksum option, 481
  - data-file-length option, 481
  - debug option, 478
  - description option, 483
  - example output, 348
  - extend-check option, 480, 481
  - fast option, 481
  - force option, 481, 482
  - help option, 478
  - information option, 481
  - keys-used option, 482
  - max-record-length option, 482
  - medium-check option, 481
  - no-symlinks option, 482
  - options, 478
  - parallel-recover option, 482
  - quick option, 482
  - read-only option, 481
  - recover option, 482
  - safe-recover option, 482
  - set-auto-increment[ option, 483
  - set-character-set option, 482
  - set-collation option, 483
  - silent option, 478
  - sort-index option, 483
  - sort-records option, 483
  - sort-recover option, 483
  - tmpdir option, 483
  - unpack option, 483
  - update-state option, 481
  - verbose option, 478
  - version option, 478
  - wait option, 479
- mysamlog, 474
- mysampack, 474, 717, 835
  - backup option, 486
  - character-sets-dir option, 487
  - debug option, 487
  - force option, 487
  - help option, 486
  - join option, 487
  - packlength option, 487
  - silent option, 487
  - test option, 487
  - tmpdir option, 487
  - verbose option, 487
  - version option, 487
  - wait option, 487
- mysam\_block\_size mysamchk variable, 479
- mysam\_ftdump, 474
  - count option, 476
  - dump option, 477
  - help option, 476
  - length option, 477
  - stats option, 477
  - verbose option, 477
- MyODBC, 1162
  - Borland, 1228
  - Borland Database Engine, 1228
  - reporting problems, 1234
- MySQL
  - defined, 5

- introduction, 5
- pronunciation, 6
- mysql, 474
  - auto-rehash option, 492
  - batch option, 492
  - character-sets-dir option, 492
  - column-names option, 493
  - compress option, 493
  - database option, 493
  - debug option, 493
  - debug-info option, 493
  - default-character-set option, 493
  - delimiter option, 493
  - execute option, 493
  - force option, 493
  - help option, 492
  - host option, 493
  - html option, 493
  - i-am-a-dummy option, 495
  - ignore-spaces option, 493
  - line-numbers option, 493
  - local-infile option, 493
  - named-commands option, 494
  - no-auto-rehash option, 494
  - no-beep option, 494
  - no-named-commands option, 494
  - no-pager option, 494
  - no-tee option, 494
  - one-database option, 494
  - pager option, 494
  - password option, 494
  - port option, 494
  - prompt option, 494
  - protocol option, 495
  - quick option, 495
  - raw option, 495
  - reconnect option, 495
  - safe-updates option, 495
  - secure-auth option, 495
  - sigint-ignore option, 495
  - silent option, 495
  - skip-column-names option, 495
  - skip-line-numbers option, 495
  - socket option, 495
  - SSL options, 495
  - table option, 495
  - tee option, 496
  - unbuffered option, 496
  - user option, 496
  - verbose option, 496
  - version option, 496
  - vertical option, 496
  - wait option, 496
  - xml option, 496
- MySQL AB
  - defined, 4
- MySQL binary distribution, 44
- MYSQL C type, 1065
- MySQL Cluster, 910
  - administration, 977, 977, 978, 979, 980, 980, 983, 990
  - and DNS, 916
  - and IP addressing, 916
  - and networking, 917
  - API node, 912, 953
  - arbitrator, 1025
  - available platforms, 910
  - backups, 991, 992, 992, 993, 995, 996
  - benchmarks, 1013
  - CHECKPOINT Events, 987
  - cluster logs, 984, 985
  - CLUSTERLOG commands, 985
  - CLUSTERLOG STATISTICS command, 990
  - commands, 977, 978, 979, 980, 980, 983
  - compiling from source, 927
  - concepts, 912
  - configuration, 915, 926, 927, 933, 933, 935, 953, 968, 974, 977
  - configuration (example), 930
  - configuration changes, 971
  - configuration files, 920, 929
  - configuration parameters, 958, 959, 966, 968
  - configuring, 995
  - CONNECTION Events, 986
  - connectstring, 932
  - data node, 912, 935
  - data nodes, 975
  - data types supported, 1026
  - defining node hosts, 933
  - direct connections between nodes, 955
  - ENTER SINGLE USER MODE command, 984
  - ERROR Events, 989
  - error logs, 975
  - error messages, 1023
  - event log format, 986
  - event logging thresholds, 985
  - event logs, 984, 985
  - event severity levels, 986
  - event types, 985, 986
  - EXIT command, 984
  - EXIT SINGLE USER MODE command, 984
  - FAQ, 1020
  - general description, 911
  - hardware requirements, 1021
  - HELP command, 983
  - how to obtain, 1023
  - importing existing tables, 1025
  - INFO Events, 989
  - information sources, 910
  - installation, 915, 918, 927
  - interconnects, 1009
  - log files, 975
  - logging commands, 985
  - management client (mgm), 977
  - management commands, 990
  - management node, 912, 933

- management nodes, 976
- managing, 980
- memory requirements, 1021
- mgm, 977
  - mgm client, 983
  - mgm management client, 990
  - mgm process, 977, 980
  - mgmd, 977
    - mgmd process, 976, 980
    - mysqld process, 974, 978
    - ndbd, 977
      - ndbd process, 975, 979, 990
      - ndb\_mgm, 922
      - ndb\_size.pl (utility), 1022
  - network configuration (SCI), 1010
  - network transporters, 1009, 1009
  - networking, 955, 956, 957
  - networking requirements, 1020, 1023
  - node failure (single-user mode), 991
  - node identifiers, 956, 957
  - node logs, 984
  - node types, 1021
  - NODERESTART Events, 987
  - nodes and node groups, 913
  - nodes and types, 912
  - number of computers required, 1020
  - obtaining, 918
  - partitions, 913
  - performance, 1013
  - performing queries, 922
  - platforms supported, 1021
  - process management, 973
  - quick configuration, 927
  - QUIT command, 984
  - replicas, 913
  - requirements, 917
  - resetting, 971
  - RESTART command, 983
  - restarting, 926
  - restoring backups, 993
  - roles of computers, 1020
  - runtime statistics, 990
  - SCI (Scalable Coherent Interface), 957, 1009
    - SCI drivers, 1012
    - SCI network configuration, 1010
    - SCI software installation, 1010
    - SCI software requirements, 1009
    - SCI, performance vs TCP/IP, 1014
  - shared memory transport, 956
  - SHOW command, 983
  - SHUTDOWN command, 984
  - shutting down, 926
  - single-user mode, 984, 991
  - SQL node, 912, 953
  - SQL nodes, 974
  - SQL statements, 1023
  - START BACKUP command, 992
  - START command, 983
    - starting and stopping, 1026
    - starting nodes, 921
    - starting or restarting, 980
    - starting with --initial, 927
    - STARTUP Events, 987
    - STATISTICS Events, 989
    - STATUS command, 984
    - STOP command, 983
    - storage requirements, 613
    - terminology, 1027
    - trace files, 975
    - transaction isolation levels supported, 1023
    - transactions, 937, 1023
    - transporters
      - Scalable Coherent Interface (SCI), 957
      - shared memory (SHM), 956
      - TCP/IP, 955
    - troubleshooting backups, 996
    - upgrades and downgrades, 970, 970, 971
    - using tables and data, 922
    - vs replication, 1020
- MySQL Cluster Glossary, 1027
- MySQL Cluster How-To, 915
- MySQL Cluster limitations, 1014
  - causing errors, 1015
  - database objects, 1017
  - geometry datatypes, 1015
  - implementation, 1019
  - imposed by configuration, 1016
  - multiple management servers, 1019
  - multiple MySQL servers, 1018
  - performance, 1017
  - syntax, 1015
  - transactions, 1015
  - unsupported features, 1017, 1018
- MySQL Cluster processes (types), 973
- MySQL Cluster utilities, 996
- mysql command options, 492
- mysql commands
  - list of, 497
- MySQL Dolphin name, 6
- MySQL history, 6
- mysql history file, 497
- MySQL mailing lists, 20
- MySQL name, 6
- mysql prompt command, 499
- MySQL source distribution, 44
- mysql status command, 498
- MySQL storage engines, 828
- MySQL version, 54
- MySQL++, 1159
- mysql.server, 221
- mysql.sock
  - changing location of, 96
  - protection, 1519
- MYSQL323 SQL mode, 275
- MYSQL40 SQL mode, 276
- mysqlaccess, 474



- brief option, 505
- commit option, 505
- copy option, 505
- db option, 505
- debug option, 505
- help option, 505
- host option, 505
- howto option, 505
- old\_server option, 505
- password option, 505
- plan option, 505
- preview option, 505
- relnotes option, 505
- rhost option, 505
- rollback option, 506
- spassword option, 506
- superuser option, 506
- table option, 506
- user option, 506
- version option, 506
- mysqladmin, 474, 704, 717, 804, 806, 810, 811
  - character-sets-dir option, 509
  - compress option, 509
  - count option, 510
  - debug option, 510
  - default-character-set option, 510
  - force option, 510
  - help option, 509
  - host option, 510
  - password option, 510
  - port option, 510
  - protocol option, 510
  - relative option, 510
  - silent option, 510
  - sleep option, 510
  - socket option, 510
  - SSL options, 511
  - user option, 511
  - verbose option, 511
  - version option, 511
  - vertical option, 511
  - wait option, 511
- mysqladmin command options, 509
- mysqladmin option
  - mysqld\_multi, 285
- mysqlbinlog, 475
  - character-sets-dir option, 512
  - database option, 512
  - debug option, 512
  - disable-log-bin option, 512
  - force-read option, 512
  - help option, 512
  - host option, 513
  - local-load option, 513
  - offset option, 513
  - password option, 513
  - port option, 513
  - position option, 513
  - protocol option, 513
  - read-from-remote-server option, 513
  - result-file option, 513
  - set-charset option, 513
  - short-form option, 513
  - socket option, 513
  - start-datetime option, 513
  - start-position option, 514
  - stop-datetime option, 514
  - stop-position option, 514
  - to-last-log option, 514
  - user option, 514
  - version option, 514
- mysqlbug script, 28
- mysqlcheck, 475
  - all-databases option, 517
  - all-in-1 option, 517
  - analyze option, 517
  - auto-repair option, 517
  - character-sets-dir option, 517
  - check option, 517
  - check-only-changed option, 517
  - compress option, 517
  - databases option, 517
  - debug option, 517
  - default-character-set option, 517
  - extended option, 517
  - fast option, 517
  - force option, 518
  - help option, 516
  - host option, 518
  - medium-check option, 518
  - optimize option, 518
  - password option, 518
  - port option, 518
  - protocol option, 518
  - quick option, 518
  - repair option, 518
  - silent option, 518
  - socket option, 518
  - SSL options, 518
  - tables option, 519
  - use-frm option, 519
  - user option, 519
  - verbose option, 519
  - version option, 519
- mysqlclient library, 1060
- mysqld, 221
  - allow-suspicious-udfs option, 223, 293
  - ansi option, 223
  - as MySQL Cluster process, 974, 978
  - basedir option, 223
  - bdb-home option, 901
  - bdb-lock-detect option, 901
  - bdb-logdir option, 901
  - bdb-no-recover option, 901
  - bdb-no-sync option, 901
  - bdb-shared-data option, 901

- 
- bdb-tmpdir option, 901
  - big-tables option, 223
  - bind-address option, 223
  - binlog-do-db option, 364
  - binlog-ignore-db option, 364
  - bootstrap option, 224
  - character-set-client-handshake option, 224
  - character-set-server option, 224
  - character-sets-dir option, 224
  - chroot option, 224
  - collation-server option, 224
  - command options, 222
  - console option, 224
  - core-file option, 224
  - datadir option, 224
  - debug option, 224
  - default-character-set option, 225
  - default-collation option, 225
  - default-storage-engine option, 225
  - default-table-type option, 225
  - default-time-zone option, 225
  - delay-key-write option, 225, 832
  - delay-key-write-for-all-tables option, 225
  - des-key-file option, 225
  - enable-named-pipe option, 225
  - exit-info option, 225
  - external-locking option, 226
  - flush option, 226
  - help option, 223
  - init-file option, 226
  - innodb option, 846
  - innodb-safe-binlog option, 226
  - innodb\_status\_file option, 846
  - language option, 226
  - local-infile option, 293
  - log option, 226
  - log-bin option, 226
  - log-bin-index option, 227
  - log-error option, 227
  - log-isam option, 227
  - log-long-format option, 227
  - log-queries-not-using-indexes option, 227
  - log-short-format option, 227
  - log-slave-updates option, 399
  - log-slow-admin-statements option, 227
  - log-slow-queries option, 227
  - log-update option, 227
  - log-warnings option, 228, 399
  - low-priority-updates option, 228
  - master-connect-retry option, 400
  - master-host option, 400
  - master-info-file option, 400
  - master-password option, 400
  - master-port option, 400
  - master-retry-count option, 400
  - master-ssl option, 400
  - master-ssl-ca option, 400
  - master-ssl-capath option, 400
  - master-ssl-cert option, 400
  - master-ssl-cipher option, 400
  - master-ssl-key option, 400
  - master-user option, 400
  - max-relay-log-size option, 401
  - memlock option, 228
  - myisam-recover option, 228, 832
  - ndb-connectstring option, 228
  - ndbcluster option, 229
    - new option, 229
    - old-passwords option, 229, 293
    - old-protocol option, 229
    - one-thread option, 229
    - open-files-limit option, 229
    - pid-file option, 229
    - port option, 229
    - read-only option, 401
    - relay-log option, 401
    - relay-log-index option, 401
    - relay-log-info-file option, 401
    - relay-log-purge option, 401
    - relay-log-space-limit option, 401
  - replicate-do-db option, 402
  - replicate-do-table option, 402
  - replicate-ignore-db option, 402
  - replicate-ignore-table option, 402
  - replicate-rewrite-db option, 402
  - replicate-same-server-id option, 403
  - replicate-wild-do-table option, 403
  - replicate-wild-ignore-table option, 403
  - report-host option, 403
  - report-port option, 404
  - role in MySQL Cluster, 912
  - safe-mode option, 230
  - safe-show-database option, 230, 293
  - safe-user-create option, 230, 294
  - secure-auth option, 230, 294
  - shared-memory option, 230
  - shared-memory-base-name option, 230
  - skip-bdb option, 230, 901
  - skip-concurrent-insert option, 230
  - skip-delay-key-write option, 230
  - skip-external-locking option, 230
  - skip-grant-tables option, 230, 294
  - skip-host-cache option, 231
  - skip-innodb option, 231
  - skip-isam option, 231
  - skip-merge option, 231
  - skip-name-resolve option, 231, 294
  - skip-ndbcluster option, 231
  - skip-networking option, 231, 294
  - skip-new option, 231
  - skip-safemalloc option, 232
  - skip-show-database option, 232, 294
  - skip-slave-start option, 404
  - skip-stack-trace option, 232
  - skip-symbolic-links option, 232
  - skip-symlink option, 231
-

- skip-thread-priority option, 232
- slave-load-tmpdir option, 404
- slave-net-timeout option, 404
- slave-skip-errors option, 404
- slave\_compressed\_protocol option, 404
- socket option, 232
- sql-mode option, 232
- SSL options, 231, 305
- standalone option, 232
- starting, 295
- symbolic-links option, 232
- sync-bdb-logs option, 902
- temp-pool option, 233
- tmpdir option, 233
- transaction-isolation option, 233
- user option, 233
- version option, 233
- mysqld library, 1060
- mysqld option
  - mysqld\_multi, 285
  - mysqld\_safe, 282
- mysqld options, 464
- mysqld server
  - buffer sizes, 464
- mysqld-max, 221, 277
- mysqld-version option
  - mysqld\_safe, 282
- mysqldump, 137, 475
  - add-drop-database option, 520
  - add-drop-table option, 520
  - add-locks option, 520
  - all-databases option, 520
  - allow-keywords option, 520
  - character-sets-dir option, 520
  - comments option, 520
  - compact option, 521
  - compatible option, 521
  - complete-insert option, 521
  - compress option, 521
  - create-options option, 521
  - databases option, 521
  - debug option, 521
  - default-character-set option, 521
  - delayed-insert option, 521
  - delete-master-logs option, 522
  - disable-keys option, 522
  - extended-insert option, 522
  - flush-logs option, 522
  - force option, 522
  - help option, 520
  - hex-blob option, 522
  - host option, 522
  - ignore-table option, 522
  - insert-ignore option, 523
  - lock-all-tables option, 523
  - lock-tables option, 523
  - master-data option, 523
  - no-autocommit option, 523
  - no-create-db option, 523
  - no-create-info option, 523
  - no-data option, 523
  - opt option, 523
  - order-by-primary option, 524
  - password option, 524
  - port option, 524
  - protocol option, 524
  - quick option, 524
  - quote-names option, 524
  - result-file option, 524
  - set-charset option, 524
  - single-transaction option, 525
  - skip-comments option, 525
  - skip-opt option, 525
  - socket option, 525
  - SSL options, 525
  - tab option, 525
  - tables option, 525
  - user option, 526
  - verbose option, 526
  - version option, 526
  - where option, 526
  - xml option, 526
- mysqld\_multi, 221
  - config-file option, 285
  - example option, 285
  - help option, 285
  - log option, 285
  - mysqladmin option, 285
  - mysqld option, 285
  - no-log option, 286
  - password option, 286
  - silent option, 286
  - tcp-ip option, 286
  - user option, 286
  - verbose option, 286
  - version option, 286
- mysqld\_safe, 221
  - autoclose option, 281
  - basedir option, 281
  - core-file-size option, 281
  - datadir option, 281
  - defaults-extra-file option, 281
  - defaults-file option, 282
  - ledir option, 282
  - log-error option, 282
  - mysqld option, 282
  - mysqld-version option, 282
  - nice option, 282
  - no-defaults option, 282
  - open-files-limit option, 282
  - pid-file option, 282
  - port option, 282
  - socket option, 282
  - timezone option, 283
  - user option, 283
- mysqlhotcopy, 475

- addtodest option, 528
- allowold option, 528
- checkpoint option, 528
- chroot option, 528
- debug option, 528
- dryrun option, 528
- flushlog option, 528
- help option, 528
- host option, 528
- keepold option, 528
- method option, 529
- noindices option, 529
- password option, 529
- port option, 529
- quiet option, 529
- record\_log\_pos option, 529
- regex option, 529
- resetmaster option, 529
- resetslave option, 529
- socket option, 529
- suffix option, 529
- tmpdir option, 529
- user option, 529
- mysqlimport, 137, 475, 731
  - character-sets-dir option, 530
  - columns option, 530
  - compress option, 530
  - debug option, 530
  - default-character-set option, 530
  - delete option, 530
  - force option, 531
  - help option, 530
  - host option, 531
  - ignore option, 531
  - ignore-lines option, 531
  - local option, 531
  - lock-tables option, 531
  - low-priority option, 531
  - password option, 531
  - port option, 531
  - protocol option, 531
  - replace option, 532
  - silent option, 532
  - socket option, 532
  - SSL options, 532
  - user option, 532
  - verbose option, 532
  - version option, 532
- mysqlshow, 475
  - character-sets-dir option, 533
  - compress option, 533
  - debug option, 533
  - default-character-set option, 533
  - help option, 533
  - host option, 533
  - keys option, 534
  - password option, 534
  - port option, 534
  - protocol option, 534
  - socket option, 534
  - SSL options, 534
  - status option, 534
  - user option, 534
  - verbose option, 534
  - version option, 534
- mysqltest
  - MySQL Test Suite, 1483
  - mysql\_affected\_rows(), 1074
  - mysql\_autocommit()., 1074
  - MYSQL\_BIND C type, 1121
  - mysql\_change\_user(), 1075
  - mysql\_character\_set\_name(), 1076
  - mysql\_close(), 1076
  - mysql\_commit()., 1077
  - mysql\_connect(), 1077
  - mysql\_create\_db(), 1077
  - mysql\_data\_seek(), 1078
  - MYSQL\_DEBUG environment variable, 475, 1882, 1886
  - mysql\_debug(), 1078
  - mysql\_drop\_db(), 1079
  - mysql\_dump\_debug\_info(), 1080
  - mysql\_eof(), 1080
  - mysql\_errno(), 1081
  - mysql\_error(), 1082
  - mysql\_escape\_string(), 1082
  - mysql\_explain\_log, 474
    - date option, 504
    - host option, 504
    - password option, 504
    - printerror option, 504
    - socket option, 504
    - user option, 504
  - mysql\_fetch\_field(), 1082
  - mysql\_fetch\_fields(), 1084
  - mysql\_fetch\_field\_direct(), 1083
  - mysql\_fetch\_lengths(), 1084
  - mysql\_fetch\_row(), 1085
  - MYSQL\_FIELD C type, 1066
  - mysql\_field\_count(), 1086, 1097
  - MYSQL\_FIELD\_OFFSET C type, 1066
  - mysql\_field\_seek(), 1087
  - mysql\_field\_tell(), 1087
  - mysql\_fix\_privilege\_tables, 222, 312
  - mysql\_free\_result(), 1087
  - mysql\_get\_client\_info(), 1088
  - mysql\_get\_client\_version(), 1088
  - mysql\_get\_host\_info(), 1088
  - mysql\_get\_proto\_info(), 1088
  - mysql\_get\_server\_info(), 1089
  - mysql\_get\_server\_version(), 1089
  - MYSQL\_GROUP\_SUFFIX environment variable, 1886
  - mysql\_hex\_string(), 1089
  - MYSQL\_HISTFILE environment variable, 497, 1886
  - MYSQL\_HOME environment variable, 1886
  - MYSQL\_HOST environment variable, 304, 1886

- 
- mysql\_info(), 702, 726, 738, 760, 1090
  - mysql\_init(), 1091
  - mysql\_insert\_id(), 35, 727, 1091
  - mysql\_install\_db, 222
  - mysql\_install\_db script, 114
  - mysql\_kill(), 1092
  - mysql\_library\_end(), 1093
  - mysql\_library\_init(), 1093
  - mysql\_list\_dbs(), 1093
  - mysql\_list\_fields(), 1094
  - mysql\_list\_processes(), 1095
  - mysql\_list\_tables(), 1095
  - mysql\_more\_results(), 1096
  - mysql\_next\_result(), 1096
  - mysql\_num\_fields(), 1097
  - mysql\_num\_rows(), 1098
  - mysql\_options(), 1099
  - mysql\_ping(), 1102
  - MYSQL\_PS1 environment variable, 1886
  - MYSQL\_PWD environment variable, 304, 475, 1886
  - mysql\_query(), 1102, 1153
  - mysql\_real\_connect(), 1103
  - mysql\_real\_escape\_string(), 540, 1106
  - mysql\_real\_query(), 1107
  - mysql\_refresh(), 1108
  - mysql\_reload(), 1109
  - MYSQL\_RES C type, 1066
  - mysql\_rollback(), 1110
  - MYSQL\_ROW C type, 1066
  - mysql\_row\_seek(), 1110
  - mysql\_row\_tell(), 1110
  - mysql\_select\_db(), 1111
  - mysql\_server\_end(), 1153
  - mysql\_server\_init(), 1152
  - mysql\_set\_character\_set(), 1111
  - mysql\_set\_local\_infile\_default(), 1112, 1112
  - mysql\_set\_server\_option(), 1114
  - mysql\_shutdown(), 1114
  - mysql\_sqlstate(), 1115
  - mysql\_ssl\_set(), 1116
  - mysql\_stat(), 1116
  - MYSQL\_STMT C type, 1120
  - mysql\_stmt\_affected\_rows(), 1127
  - mysql\_stmt\_attr\_get(), 1127
  - mysql\_stmt\_attr\_set(), 1128
  - mysql\_stmt\_bind\_param(), 1128
  - mysql\_stmt\_bind\_result(), 1129
  - mysql\_stmt\_close(), 1130
  - mysql\_stmt\_data\_seek(), 1131
  - mysql\_stmt\_errno(), 1131
  - mysql\_stmt\_error(), 1132
  - mysql\_stmt\_execute(), 1132
  - mysql\_stmt\_fetch(), 1135
  - mysql\_stmt\_fetch\_column(), 1139
  - mysql\_stmt\_field\_count(), 1139
  - mysql\_stmt\_free\_result(), 1140
  - mysql\_stmt\_init(), 1140
  - mysql\_stmt\_insert\_id(), 1140
  - mysql\_stmt\_num\_rows(), 1141
  - mysql\_stmt\_param\_count(), 1141
  - mysql\_stmt\_param\_metadata(), 1142
  - mysql\_stmt\_prepare(), 1142
  - mysql\_stmt\_reset(), 1143
  - mysql\_stmt\_result\_metadata(), 1144
  - mysql\_stmt\_row\_seek(), 1145
  - mysql\_stmt\_row\_tell(), 1145
  - mysql\_stmt\_send\_long\_data(), 1145
  - mysql\_stmt\_sqlstate(), 1147
  - mysql\_stmt\_store\_result(), 1147
  - mysql\_store\_result(), 1117, 1153
  - MYSQL\_TCP\_PORT environment variable, 374, 374, 475, 1886
  - mysql\_thread\_end(), 1151
  - mysql\_thread\_id(), 1118
  - mysql\_thread\_init(), 1151
  - mysql\_thread\_safe(), 1151
  - MYSQL\_TIME C type, 1122
  - MYSQL\_UNIX\_PORT environment variable, 115, 374, 374, 475, 1886
  - mysql\_use\_result(), 1118
  - mysql\_warning\_count(), 1119
  - mysql\_zap, 475
  - my\_bool C type, 1066
  - my\_bool values
    - printing, 1066
  - my\_init(), 1150
  - my\_ulonglong C type, 1066
  - my\_ulonglong values
    - printing, 1066
- N**
- named pipes, 71, 76
  - named-commands option
    - mysql, 494
  - names, 541
    - case sensitivity, 544
    - variables, 545
  - naming
    - releases of MySQL, 44
  - NATIONAL CHAR data type, 590
  - native functions
    - adding, 1495
  - native thread support, 42
  - NATURAL LEFT JOIN, 745
  - NATURAL LEFT OUTER JOIN, 745
  - NATURAL RIGHT JOIN, 745
  - NATURAL RIGHT OUTER JOIN, 745
  - NCHAR data type, 590
  - NDB, 1020
  - ndb option
    - perror, 536
  - NDB storage engine, 910
  - ndb-connectstring option
    - mysqld, 228
  - ndbcluster option
    - mysqld, 229
-

- ndbd (MySQL Cluster process), 975
- ndbd (MySQL Cluster)
  - defined, 912
- ndb\_delete\_all, 996, 998
- ndb\_desc, 996, 998
- ndb\_drop\_index, 996, 999
- ndb\_drop\_table, 996, 1000
- ndb\_error\_reporter, 996, 1000
- ndb\_mgm, 977
- ndb\_mgm (MySQL Cluster management node client), 922
- ndb\_mgmd, 976
- ndb\_mgmd (MySQL Cluster)
  - defined, 912
- ndb\_print\_backup\_file, 996, 1001
- ndb\_print\_schema\_file, 996, 1001
- ndb\_print\_sys\_file, 996, 1001
- ndb\_restore (MySQL Cluster program), 993
- ndb\_select\_all, 996, 1002
- ndb\_select\_count, 996, 1004
- ndb\_show\_tables, 996, 1004
- ndb\_size.pl, 996, 1005
- ndb\_size.pl (utility), 1022
- ndb\_waiter, 996, 1007
- negative values, 540
- nested queries, 749
- net etiquette, 23
- netmask notation
  - in mysql.user table, 305
- NetWare, 87
- net\_buffer\_length variable, 496
- new option
  - mysqld, 229
- new procedures
  - adding, 1496
- new users
  - adding, 91, 95
- newline (\n), 539
- nice option
  - mysqld\_safe, 282
- no matching rows, 1524
- no-auto-rehash option
  - mysql, 494
- no-autocommit option
  - mysqldump, 523
- no-beep option
  - mysql, 494
- no-create-db option
  - mysqldump, 523
- no-create-info option
  - mysqldump, 523
- no-data option
  - mysqldump, 523
- no-defaults option
  - mysqld\_safe, 282
- no-log option
  - mysqld\_multi, 286
- no-named-commands option
  - mysql, 494
- no-pager option
  - mysql, 494
- no-symlinks option
  - mysiamchk, 482
- no-tee option
  - mysql, 494
- node groups (MySQL Cluster), 913
- node identifiers (MySQL Cluster), 956, 957
- node logs (MySQL Cluster), 984
- NODERESTART Events (MySQL Cluster), 987
- noindices option
  - mysqlhotcopy, 529
- non-delimited strings, 597
- Non-transactional tables, 1523
- NoOfDiskPagesToDiskAfterRestartACC
  - calculating, 968
- NoOfDiskPagesToDiskAfterRestartACC (MySQL Cluster configuration parameter), 948
- NoOfDiskPagesToDiskAfterRestartTUP
  - calculating, 968
- NoOfDiskPagesToDiskAfterRestartTUP (MySQL Cluster configuration parameter), 948
- NoOfDiskPagesToDiskDuringRestartACC (MySQL Cluster configuration parameter), 949
- NoOfDiskPagesToDiskDuringRestartTUP (MySQL Cluster configuration parameter), 948
- NoOfFragmentLogFiles, 942
  - calculating, 968
- NoOfReplicas, 936
- NOT
  - logical, 624
- NOT BETWEEN, 621
- not equal (!=), 620
- not equal (<>), 620
- NOT IN, 622
- NOT LIKE, 637
- NOT NULL
  - constraint, 39
- NOT REGEXP, 637
- Novell NetWare, 87
- NOW(), 655
- NO\_AUTO\_VALUE\_ON\_ZERO SQL mode, 273
- NO\_DIR\_IN\_CREATE SQL mode, 274
- NO\_FIELD\_OPTIONS SQL mode, 274
- NO\_KEY\_OPTIONS SQL mode, 274
- NO\_TABLE\_OPTIONS SQL mode, 274
- NO\_UNSIGNED\_SUBTRACTION SQL mode, 274
- NUL, 539
- NULL, 189, 1522
  - ORDER BY, 439, 741
  - testing for null, 620, 620, 621, 626
- NULL value, 189, 541
- NULL values
  - and AUTO\_INCREMENT columns, 1523
  - and indexes, 710
  - and TIMESTAMP columns, 1523
  - vs. empty values, 1522

- NULLIF(), 626
  - numbers, 540
  - NUMERIC data type, 587
  - numeric types, 613
  - NumGeometries(), 1053
  - NumInteriorRings(), 1052
  - NumPoints(), 1050
- O**
- Obtaining MySQL Cluster, 918
  - OCT(), 632
  - OCTET\_LENGTH(), 632
  - ODBC, 1162
  - ODBC compatibility, 543, 587, 618, 621, 708, 746
  - offset option
    - mysqlbinlog, 513
  - OLAP, 693
  - old-passwords option
    - mysqld, 229, 293
  - old-protocol option
    - mysqld, 229
  - OLD\_PASSWORD(), 680
  - old\_server option
    - mysqlaccess, 505
  - ON DUPLICATE KEY, 724, 1575
  - one-database option
    - mysql, 494
  - one-thread option
    - mysqld, 229
  - online location of manual, 1
  - ONLY\_FULL\_GROUP\_BY
    - SQL mode, 696
  - ONLY\_FULL\_GROUP\_BY SQL mode, 275
  - Open Source
    - defined, 5
  - open tables, 462, 509
  - open-files-limit option
    - mysqld, 229
    - mysqld\_safe, 282
  - OpenGIS, 1032
  - opening
    - tables, 462
  - opens, 509
  - OpenSSL, 329
  - open\_files\_limit variable, 514
  - operating systems
    - file-size limits, 10
    - supported, 42
    - Windows versus Unix, 79
  - operations
    - arithmetic, 639
  - operators, 617
    - cast, 638, 673
    - logical, 623
    - precedence, 617
  - opt option
    - mysqldump, 523
  - optimization
    - tips, 443
  - optimizations, 429
  - optimize option
    - mysqlcheck, 518
  - OPTIMIZE TABLE, 782
  - optimizing
    - DISTINCT, 435
    - filesort, 437
    - GROUP BY, 439
    - LEFT JOIN, 435
    - LIMIT, 440
    - tables, 347
  - option files, 214, 312
  - options
    - command-line
      - mysql, 492
      - mysqladmin, 509
    - configure, 95
    - embedded server, 1061
    - libmysqld, 1061
    - myisamchk, 478
    - provided by MySQL, 175
    - replication, 392
  - OR, 202
    - bitwise, 676
    - logical, 624
  - Oracle compatibility, 30, 693, 762
  - ORACLE SQL mode, 276
  - ORD(), 632
  - ORDER BY, 186, 700, 741
    - NULL, 439, 741
  - order-by-primary option
    - mysqldump, 524
  - OUTFILE, 743
  - Overlaps(), 1056
  - overview, 1
- P**
- packages
    - list of, 1573
  - packlength option
    - myisampack, 487
  - pack\_isam, 474
  - page-level locking, 446
  - pager option
    - mysql, 494
  - parallel-recover option
    - myisamchk, 482
  - parameters
    - server, 464
  - parentheses ( and ), 617
  - partitions (MySQL Cluster), 913
  - password
    - root user, 121
  - password encryption
    - reversibility of, 681
  - password option
    - mysql, 494

- mysqlaccess, 505
- mysqladmin, 510
- mysqlbinlog, 513
- mysqlcheck, 518
- mysqldump, 524
- mysqld\_multi, 286
- mysqlhotcopy, 529
- mysqlimport, 531
- mysqlshow, 534
- mysql\_explain\_log, 504
- PASSWORD(), 306, 327, 680, 1510
- passwords
  - for users, 321
  - forgotten, 1514
  - lost, 1514
  - resetting, 1514
  - security, 296
  - setting, 327, 775, 779
- PATH environment variable, 90, 210, 1886
- pattern matching, 190, 1887
- performance
  - benchmarks, 419
  - disk issues, 469
  - estimating, 428
  - improving, 410, 450
- PERIOD\_ADD(), 655
- PERIOD\_DIFF(), 655
- Perl
  - installing, 170
  - installing on Windows, 172
- Perl API, 1158
- Perl DBI/DBD
  - installation problems, 172
- permission checks
  - effect on speed, 420
- perrow, 475
  - help option, 536
  - ndb option, 536
  - silent option, 536
  - verbose option, 536
  - version option, 536
- PHP API, 1157
- PI(), 643
- pid-file option
  - mysqld, 229
  - mysqld\_safe, 282
- PIPES\_AS\_CONCAT SQL mode, 275
- plan option
  - mysqlaccess, 505
- POINT data type, 1041
- point in time recovery, 341
- Point(), 1044
- PointFromText(), 1042
- PointFromWKB(), 1043
- PointN(), 1050
- PointOnSurface(), 1052
- PolyFromText(), 1042
- PolyFromWKB(), 1043
- POLYGON data type, 1041
- Polygon(), 1044
- PolygonFromText(), 1042
- PolygonFromWKB(), 1043
- port option
  - mysql, 494
  - mysqladmin, 510
  - mysqlbinlog, 513
  - mysqlcheck, 518
  - mysqld, 229
  - mysqldump, 524
  - mysqld\_safe, 282
  - mysqlhotcopy, 529
  - mysqlimport, 531
  - mysqlshow, 534
- portability, 417
  - types, 615
- porting
  - to other systems, 1876
- PortNumber, 933, 955
- position option
  - mysqlbinlog, 513
- POSITION(), 632
- post-install
  - multiple servers, 368
- post-installation
  - setup and testing, 109
- PostgreSQL compatibility, 32
- POSTGRES SQL mode, 276
- POWER(), 643
- POWER(), 643
- precedence
  - operator, 617
- PREPARE, 825, 826
- preview option
  - mysqlaccess, 505
- PRIMARY KEY, 700, 709
  - constraint, 38
- primary key
  - deleting, 700
- printerror option
  - mysql\_explain\_log, 504
- privilege
  - changes, 310
- privilege information
  - location, 300
- privilege system, 296
  - described, 296
- privileges
  - access, 296
  - adding, 322
  - default, 121
  - deleting, 325, 771
  - display, 795
  - dropping, 325, 771
  - granting, 771
  - revoking, 778
- problems



access denied errors, 1499  
 common errors, 1498  
 compiling, 101  
 DATE columns, 1521  
 date values, 597  
 installing on IBM-AIX, 156  
 installing on Solaris, 147  
 installing Perl, 172  
 linking, 1512  
 ODBC, 1234  
 reporting, 24  
 starting the server, 118  
 table locking, 448  
 time zone, 1520  
**PROCEDURE**, 744  
 procedures  
   adding, 1496  
   stored, 35  
 process management (MySQL Cluster), 973  
 process support, 42  
 processes  
   display, 798  
 processing  
   arguments, 1490  
**PROCESSLIST**, 798  
 program variables  
   setting, 219  
 programs  
   client, 474, 1155  
   crash-me, 417  
   server side, 221  
   utility, 474  
 prompt option  
   mysql, 494  
 prompts  
   meanings, 178  
 pronunciation  
   MySQL, 6  
 protocol option  
   mysql, 495  
   mysqladmin, 510  
   mysqlbinlog, 513  
   mysqlcheck, 518  
   mysqldump, 524  
   mysqlimport, 531  
   mysqlshow, 534  
**PURGE MASTER LOGS**, 814  
**PURGE STALE SESSIONS**, 981  
 Python API, 1159

## Q

**QUARTER()**, 655  
 queries  
   entering, 176  
   estimating performance, 428  
   examples, 198  
   speed of, 420  
   Twin Studies project, 205

Query Cache, 375  
 questions, 509  
   answering, 23  
 quick option  
   myisamchk, 482  
   mysql, 495  
   mysqlcheck, 518  
   mysqldump, 524  
 quiet option  
   mysqlhotcopy, 529  
**QUIT** command (MySQL Cluster), 984  
**QUOTE()**, 633  
 quote-names option  
   mysqldump, 524  
 quotes  
   in strings, 539  
 quoting, 540  
 quoting binary data, 540  
 quoting of identifiers, 542

## R

**RADIANS()**, 643  
**RAID**  
   compile errors, 102  
   table type, 714  
**RAND()**, 643  
 raw option  
   mysql, 495  
 re-creating  
   grant tables, 115  
 read-from-remote-server option  
   mysqlbinlog, 513  
 read-only option  
   myisamchk, 481  
   mysqld, 401  
 read\_buffer\_size myisamchk variable, 479  
**REAL** data type, 587  
**REAL\_AS\_FLOAT** SQL mode, 275  
**ReceiveBufferMemory**, 955  
 reconfiguring, 101, 101  
 reconnect option  
   mysql, 495  
 record\_log\_pos option  
   mysqlhotcopy, 529  
 recover option  
   myisamchk, 482  
 recovery  
   from crash, 343  
   point in time, 341  
**RedoBuffer**, 950  
 reducing  
   data size, 450  
 references, 701  
 ref\_or\_null, 434  
**REGEXP**, 637  
**REGEXP** operator, 1887  
 regexp option  
   mysqlhotcopy, 529

- regular expression syntax, 1887
- Related(), 1056
- relational databases
  - defined, 5
- relative option
  - mysqladmin, 510
- relay-log option
  - mysqld, 401
- relay-log-index option
  - mysqld, 401
- relay-log-info-file option
  - mysqld, 401
- relay-log-purge option
  - mysqld, 401
- relay-log-space-limit option
  - mysqld, 401
- release numbers, 44
- releases
  - naming scheme, 44
  - testing, 45
  - updating, 47
- RELEASE\_LOCK(), 689
- relnotes option
  - mysqlaccess, 505
- RENAME TABLE, 718
- reordering
  - columns, 1528
- repair option
  - mysqlcheck, 518
- repair options
  - myisamchk, 481
- REPAIR TABLE, 783
- repairing
  - tables, 345
- REPEAT(), 633
- replace, 475
- REPLACE, 738
- replace option
  - mysqlimport, 532
- REPLACE(), 633
- replicas (MySQL Cluster), 913
- replicate-do-db option
  - mysqld, 402
- replicate-do-table option
  - mysqld, 402
- replicate-ignore-db option
  - mysqld, 402
- replicate-ignore-table option
  - mysqld, 402
- replicate-rewrite-db option
  - mysqld, 402
- replicate-same-server-id option
  - mysqld, 403
- replicate-wild-do-table option
  - mysqld, 403
- replicate-wild-ignore-table option
  - mysqld, 403
- replication, 381
  - replication limitations, 392
  - replication masters
    - statements, 813
  - replication options, 392
  - replication slaves
    - statements, 816
  - report-host option
    - mysqld, 403
  - report-port option
    - mysqld, 404
  - reporting
    - bugs, 24
    - Connector/NET problems, 1423
    - Connector/ODBC problems, 1234
    - errors, 1, 24
    - MyODBC problems, 1234
- REQUIRE GRANT option, 776
- reserved words
  - exceptions, 547
- RESET MASTER, 814
- RESET SLAVE, 819
- resetmaster option
  - mysqlhotcopy, 529
- resetslave option
  - mysqlhotcopy, 529
- RESTART command (MySQL Cluster), 983
- restarting
  - the server, 113
- RestartOnErrorInsert, 945
- RESTORE TABLE, 784
- restoring backups
  - in MySQL Cluster, 993
- restrictions
  - subqueries, 1892
- result-file option
  - mysqlbinlog, 513
  - mysqldump, 524
- retrieving
  - data from tables, 183
- return (r), 539
- return values
  - UDFs, 1492
- REVERSE(), 633
- REVOKE, 778
- revoking
  - privileges, 778
- rhost option
  - mysqlaccess, 505
- RIGHT JOIN, 745
- RIGHT OUTER JOIN, 745
- RIGHT(), 633
- RLIKE, 637
- ROLLBACK, 33, 765
- rollback option
  - mysqlaccess, 506
- ROLLBACK TO SAVEPOINT, 767
- rolling restart (MySQL Cluster), 970
- rolling upgrades and downgrades (MySQL Cluster), 970

- ROLLUP, 693
  - root password, 121
  - root user
    - password resetting, 1514
  - ROUND(), 644
  - rounding errors, 586
  - row-level locking, 446
  - rows
    - counting, 192
    - deleting, 1524
    - locking, 35
    - matching problems, 1524
    - selecting, 184
    - sorting, 186
  - RPAD(), 633
  - RPM file, 81
  - RPM Package Manager, 81
  - RTRIM(), 634
  - RTS-threads, 1884
  - running
    - ANSI mode, 29
    - batch mode, 196
    - multiple servers, 368
    - queries, 176
  - running configure after prior invocation, 101
- S**
- safe-mode option
    - mysqld, 230
  - safe-recover option
    - myisamchk, 482
  - safe-show-database option
    - mysqld, 230, 293
  - safe-updates option, 502
    - mysql, 495
  - safe-user-create option
    - mysqld, 230, 294
  - Sakila, 6
  - SAVEPOINT, 767
  - SCI (Scalable Coherent Interface), 957, 1009
  - script files, 196
  - scripts
    - mysqlbug, 28
    - mysql\_install\_db, 114
  - searching
    - and case sensitivity, 1520
    - full-text, 662
    - MySQL Web pages, 24
    - two keys, 202
  - SECOND(), 655
  - secure-auth option
    - mysql, 495
    - mysqld, 230, 294
  - security
    - against attackers, 291
  - security system, 296
  - SEC\_TO\_TIME(), 656
  - SELECT
    - LIMIT, 739
    - optimizing, 420
    - Query Cache, 375
    - SELECT INTO TABLE, 33
    - SELECT speed, 428
    - selecting
      - databases, 180
    - select\_limit variable, 497
    - SendBufferMemory, 955
    - SendLimit, 958
    - SendSignalId, 955, 956, 958
    - SEQUENCE, 203
    - sequence emulation, 686
    - sequences, 203
    - server
      - connecting, 175, 303
      - debugging, 1877
      - disconnecting, 175
      - restart, 113
      - shutdown, 113
      - starting, 111
      - starting and stopping, 116
      - starting problems, 118
    - server variables, 234, 806
    - server-side programs, 221
    - ServerPort, 936
    - servers
      - multiple, 368
    - SESSION\_USER(), 686
    - SET, 785
      - AUTOCOMMIT, 785
      - BIG\_TABLES, 785
      - CHARACTER SET, 558, 785
      - FOREIGN\_KEY\_CHECKS, 785
      - IDENTITY, 785
      - INSERT\_ID, 785
      - LAST\_INSERT\_ID, 785
      - NAMES, 558, 785
      - ONE\_SHOT, 785
      - size, 615
      - SQL\_AUTO\_IS\_NULL, 785
      - SQL\_BIG\_SELECTS, 785
      - SQL\_BUFFER\_SELECT, 785
      - SQL\_LOG\_BIN, 785
      - SQL\_LOG\_OFF, 785
      - SQL\_LOG\_UPDATE, 785
      - SQL\_NOTES, 785
      - SQL\_QUOTE\_SHOW\_CREATE, 785
      - SQL\_SAFE\_UPDATES, 785
      - SQL\_SELECT\_LIMIT, 785
      - SQL\_WARNINGS, 785
      - TIMESTAMP, 785
      - UNIQUE\_CHECKS, 785
    - SET data type, 591, 611
    - SET GLOBAL SQL\_SLAVE\_SKIP\_COUNTER, 820
    - SET OPTION, 785
    - SET PASSWORD, 779
    - SET PASSWORD statement, 327

- SET SQL\_LOG\_BIN, 814
- SET TRANSACTION, 770
- set-auto-increment[ option
  - myisamchk, 483
- set-character-set option
  - myisamchk, 482
- set-charset option
  - mysqlbinlog, 513
  - mysqldump, 524
- set-collation option
  - myisamchk, 483
- setting
  - passwords, 327
- setting passwords, 779
- setting program variables, 219
- setup
  - post-installation, 109
- SHA(), 681
- SHA1(), 681
- shared memory transporter, 956
- shared-memory option
  - mysqld, 230
- shared-memory-base-name option
  - mysqld, 230
- SharedBufferSize, 958
- shell syntax, 4
- ShmKey, 956
- ShmSize, 956
- short-form option
  - mysqlbinlog, 513
- SHOW
  - in MySQL Cluster management client, 929
- SHOW BINARY LOGS, 815
- SHOW BINLOG EVENTS, 791, 815
- SHOW CHARACTER SET, 791
- SHOW COLLATION, 791
- SHOW COLUMNS, 790, 792
- SHOW command (MySQL Cluster), 983
- SHOW CREATE DATABASE, 790, 792
- SHOW CREATE TABLE, 790, 792
- SHOW DATABASES, 790, 793
- SHOW ENGINE, 790, 793
- SHOW ENGINES, 790, 793
- SHOW ERRORS, 790, 795
- SHOW FIELDS, 790
- SHOW GRANTS, 790, 795
- SHOW INDEX, 790, 795
- SHOW INNODB STATUS, 790
- SHOW KEYS, 790, 795
- SHOW MASTER LOGS, 791, 815
- SHOW MASTER STATUS, 791, 815
- SHOW OPEN TABLES, 790, 797
- SHOW PRIVILEGES, 790, 798
- SHOW PROCESSLIST, 790, 798
- SHOW SLAVE HOSTS, 791, 815
- SHOW SLAVE STATUS, 791, 820
- SHOW STATUS, 790
- SHOW STORAGE ENGINES, 793
- SHOW TABLE STATUS, 790
- SHOW TABLE TYPES, 790, 793
- SHOW TABLES, 790, 806
- SHOW VARIABLES, 790
- SHOW WARNINGS, 790, 807
- SHUTDOWN command (MySQL Cluster), 984
- shutdown\_timeout variable, 511
- shutting down
  - the server, 113
- sigint-ignore option
  - mysql, 495
- SIGN(), 645
- silent column changes, 716
- silent option
  - myisamchk, 478
  - mysampack, 487
  - mysql, 495
  - mysqladmin, 510
  - mysqlcheck, 518
  - mysqld\_multi, 286
  - mysqlimport, 532
  - percona, 536
- SIN(), 645
- single quote (\'), 539
- single-transaction option
  - mysqldump, 525
- single-user mode (MySQL Cluster), 984, 991
- size of tables, 10
- sizes
  - display, 584
- skip-bdb option
  - mysqld, 230, 901
- skip-column-names option
  - mysql, 495
- skip-comments option
  - mysqldump, 525
- skip-concurrent-insert option
  - mysqld, 230
- skip-delay-key-write option
  - mysqld, 230
- skip-external-locking option
  - mysqld, 230
- skip-grant-tables option
  - mysqld, 230, 294
- skip-host-cache option
  - mysqld, 231
- skip-innodb option
  - mysqld, 231
- skip-isam option
  - mysqld, 231
- skip-line-numbers option
  - mysql, 495
- skip-merge option
  - mysqld, 231
- skip-name-resolve option
  - mysqld, 231, 294
- skip-ndbcluster option
  - mysqld, 231

- skip-networking option
  - mysqld, 231, 294
- skip-new option
  - mysqld, 231
- skip-opt option
  - mysqldump, 525
- skip-safemalloc option
  - mysqld, 232
- skip-show-database option
  - mysqld, 232, 294
- skip-slave-start option
  - mysqld, 404
- skip-stack-trace option
  - mysqld, 232
- skip-symbolic-links option
  - mysqld, 232
- skip-symlink option
  - mysqld, 231
- skip-thread-priority option
  - mysqld, 232
- slave-load-tmpdir option
  - mysqld, 404
- slave-net-timeout option
  - mysqld, 404
- slave-skip-errors option
  - mysqld, 404
- slave\_compressed\_protocol option
  - mysqld, 404
- sleep option
  - mysqladmin, 510
- slow queries, 509
- slow query log, 367
- SMALLINT data type, 585
- socket location
  - changing, 96
- socket option
  - mysql, 495
  - mysqladmin, 510
  - mysqlbinlog, 513
  - mysqlcheck, 518
  - mysqld, 232
  - mysqldump, 525
  - mysqld\_safe, 282
  - mysqlhotcopy, 529
  - mysqlimport, 532
  - mysqlshow, 534
  - mysql\_explain\_log, 504
- Solaris
  - installation, 86
- Solaris installation problems, 147
- Solaris troubleshooting, 103
- Solaris x86\_64 issues, 874
- sort-index option
  - myisamchk, 483
- sort-records option
  - myisamchk, 483
- sort-recover option
  - myisamchk, 483
- sorting
  - character sets, 354
  - data, 186
  - grant tables, 307, 309
  - table rows, 186
- sort\_buffer\_size myisamchk variable, 479
- sort\_key\_blocks myisamchk variable, 479
- SOUNDEX(), 634
- SOUNDS LIKE, 634
- source distribution
  - installing, 92
- source distributions
  - on Linux, 140
- SPACE(), 634
- spassword option
  - mysqlaccess, 506
- Spatial Extensions in MySQL, 1032
- speed
  - compiling, 466
  - increasing with replication, 381
  - inserting, 441
  - linking, 466
  - of queries, 420, 428
- SQL
  - defined, 5
- SQL mode
  - ONLY\_FULL\_GROUP\_BY, 696
- SQL node (MySQL Cluster)
  - defined, 912
- SQL nodes (MySQL Cluster), 974
- SQL statements
  - replication masters, 813
  - replication slaves, 816
- SQL-92
  - extensions to, 28
- sql-mode option
  - mysqld, 232
- SQL\_BIG\_RESULT, 745
- SQL\_BUFFER\_RESULT, 745
- SQL\_CACHE, 377, 745
- SQL\_CALC\_FOUND\_ROWS, 745
- SQL\_NO\_CACHE, 377, 745
- SQL\_SMALL\_RESULT, 745
- sql\_yacc.cc problems, 102
- SQRT(), 645
- square brackets, 584
- SRID(), 1048
- SSH, 335
- SSL and X509 Basics, 329
- SSL command options, 332
- ssl option, 332
- SSL options
  - mysql, 495
  - mysqladmin, 511
  - mysqlcheck, 518
  - mysqld, 231, 305
  - mysqldump, 525
  - mysqlimport, 532

- mysqlshow, 534
- SSL related options, 776
- ssl-ca option, 332
- ssl-capath option, 332
- ssl-cert option, 332
- ssl-cipher option, 332
- ssl-key option, 332
- stability, 9
- standalone option
  - mysqld, 232
- Standard SQL
  - differences from, 32, 778
  - extensions to, 28, 29
- standards compatibility, 28
- START BACKUP command (MySQL Cluster), 992
- START command (MySQL Cluster), 983
- START SLAVE, 824
- START TRANSACTION, 765
- start-datetime option
  - mysqlbinlog, 513
- start-position option
  - mysqlbinlog, 514
- StartFailureTimeout (MySQL Cluster configuration parameter), 946
- starting
  - comments, 37
  - mysqld, 295
  - the server, 111
  - the server automatically, 116
- Starting many servers, 368
- StartPartialTimeout, 946
- StartPartitionedTimeout (MySQL Cluster configuration parameter), 946
- StartPoint(), 1050
- STARTUP Events (MySQL Cluster), 987
- startup options
  - default, 214
- startup parameters, 464
  - mysql, 492
  - mysqladmin, 509
  - tuning, 463
- statements
  - GRANT, 322
  - INSERT, 323
  - replication masters, 813
  - replication slaves, 816
- statically
  - compiling, 96
- STATISTICS Events (MySQL Cluster), 989
- stats option
  - mysam\_ftdump, 477
- stats\_method mysamchk variable, 479
- status
  - tables, 804
- status command
  - results, 508
- STATUS command (MySQL Cluster), 984
- status option
  - mysqlshow, 534
  - status variables, 266, 804
- STD(), 693
- STDDEV(), 693
- STOP command (MySQL Cluster), 983
- STOP SLAVE, 825
- stop-datetime option
  - mysqlbinlog, 514
- stop-position option
  - mysqlbinlog, 514
- StopOnError, 945
- stopping
  - the server, 116
- stopword list
  - user-defined, 671
- storage engine
  - ARCHIVE, 905
- storage engines
  - choosing, 828
- storage nodes - see data nodes, ndbd, 975
- storage of data, 450
- storage requirements
  - data type, 613
- storage space
  - minimizing, 450
- stored procedures and triggers
  - defined, 35
- STRAIGHT\_JOIN, 744, 745
- STRCMP(), 638
- string collating, 358
- string comparison functions, 636
- string comparisons
  - case sensitivity, 636
- string functions, 627
- string literal introducer, 538, 556
- string types, 605
- StringMemory, 938
- strings
  - defined, 538
  - escaping characters, 538
  - non-delimited, 597
- striping
  - defined, 470
- STR\_TO\_DATE(), 656
- SUBDATE(), 656
- subqueries, 749
- subquery, 749
- subquery restrictions, 1892
- subselects, 749
- SUBSTR(), 634
- SUBSTRING(), 634
- SUBSTRING\_INDEX(), 635
- SUBTIME(), 657
- subtraction (-), 639
- suffix option
  - mysqlhotcopy, 529
- SUM(), 693
- superuser, 121

- superuser option
    - mysqlaccess, 506
  - support
    - for operating systems, 42
  - suppression
    - default values, 39
  - Sybase compatibility, 765
  - symbolic links, 470, 473
  - symbolic-links option
    - mysqld, 232
  - SymDifference(), 1054
  - sync-bdb-logs option
    - mysqld, 902
  - syntax
    - regular expression, 1887
  - SYSDATE(), 657
  - system
    - privilege, 296
    - security, 289
  - system optimization, 463
  - system table, 422
  - system variables, 234, 257, 806
  - SYSTEM\_USER(), 686
- T**
- tab (\t), 539
  - tab option
    - mysqldump, 525
  - table
    - changing, 698, 701, 1528
    - deleting, 718
  - table aliases, 741
  - table cache, 462
  - table is full, 787, 1509
  - table names
    - case sensitivity, 544
    - case-sensitivity, 30
  - table option
    - mysql, 495
    - mysqlaccess, 506
  - Table scans
    - avoiding, 440
  - table types
    - choosing, 828
  - table-level locking, 446
  - tables
    - BDB, 899
    - Berkeley DB, 899
    - BLACKHOLE, 907
    - changing column order, 1528
    - checking, 480
    - closing, 462
    - compressed format, 835
    - constant, 422, 430
    - copying, 714
    - counting rows, 192
    - creating, 181
    - CSV, 906
    - defragment, 835
    - defragmenting, 354, 782
    - deleting rows, 1524
    - displaying status, 804
    - dynamic, 834
    - error checking, 344
    - EXAMPLE, 905
    - flush, 509
    - fragmentation, 782
    - grant, 310
    - HEAP, 897
    - host, 310
    - improving performance, 450
    - information, 348
    - information about, 196
    - InnoDB, 838
    - ISAM, 908
    - loading data, 182
    - maintenance schedule, 353
    - maximum size, 10
    - MEMORY, 897
    - MERGE, 894
    - merging, 894
    - multiple, 194
    - MyISAM, 830
    - names, 541
    - open, 462
    - opening, 462
    - optimizing, 347
    - partitioning, 894
    - RAID, 714
    - repairing, 345
    - retrieving data, 183
    - selecting columns, 185
    - selecting rows, 184
    - sorting rows, 186
    - symbolic links, 471
    - system, 422
    - too many, 463
    - unique ID for last row, 1154
    - updating, 33
  - tables option
    - mysqlcheck, 519
    - mysqldump, 525
  - table\_cache, 462
  - TAN(), 645
  - tar
    - problems on Solaris, 86, 147
  - Tcl API, 1159
  - tcp-ip option
    - mysqld\_multi, 286
  - TCP/IP, 71, 76
  - tee option
    - mysql, 496
  - temp-pool option
    - mysqld, 233
  - temporary file
    - write access, 115

- temporary tables
    - problems, 1529
  - terminal monitor
    - defined, 175
  - test option
    - mysampack, 487
  - testing
    - connection to the server, 305
    - installation, 111
    - of MySQL releases, 45
    - post-installation, 109
  - testing mysqld
    - mysqltest, 1483
  - TEXT
    - size, 614
  - TEXT columns
    - default values, 608
    - indexing, 451, 710
  - TEXT data type, 591, 607
  - thread packages
    - differences between, 1885
  - thread support, 42
    - non-native, 104
  - threaded clients, 1155
  - threads, 509, 798, 1483
    - display, 798
    - RTS, 1884
  - TIME data type, 588, 603
  - time types, 613
  - time zone problems, 1520
  - TIME(), 657
  - TimeBetweenGlobalCheckpoints (MySQL Cluster configuration parameter), 947
  - TimeBetweenInactiveTransactionAbortCheck (MySQL Cluster configuration parameter), 947
  - TimeBetweenLocalCheckpoints (MySQL Cluster configuration parameter), 947
  - TimeBetweenWatchDogCheck, 945
  - TIMEDIFF(), 657
  - timeout, 237, 687, 729
    - connect\_timeout variable, 496, 511
    - shutdown\_timeout variable, 511
  - TIMESTAMP
    - and NULL values, 1523
  - TIMESTAMP data type, 588, 596
  - TIMESTAMP(), 657
  - timezone option
    - mysqld\_safe, 283
  - TIME\_FORMAT(), 657
  - TIME\_TO\_SEC(), 658
  - TINYBLOB data type, 591
  - TINYINT data type, 584
  - TINYTEXT data type, 591
  - tips
    - optimization, 443
  - TMPDIR environment variable, 115, 1886
  - tmpdir option
    - mysampack, 487
  - mysqld, 233
  - mysqlhotcopy, 529
  - to-last-log option
    - mysqlbinlog, 514
  - TODO
    - symlinks, 472
  - tools
    - list of, 1573
  - Touches(), 1056
  - TO\_DAYS(), 658
  - trace DBI method, 1880
  - trace files (MySQL Cluster), 975
  - transaction-isolation option
    - mysqld, 233
  - transaction-safe tables, 33, 838
  - TransactionBufferMemory, 941
  - TransactionDeadlockDetectionTimeout (MySQL Cluster configuration parameter), 948
  - TransactionInactiveTimeout (MySQL Cluster configuration parameter), 947
  - transactions
    - support, 33, 838
  - Translators
    - list of, 1570
  - triggers, 35
  - TRIM(), 635
  - troubleshooting
    - FreeBSD, 103
    - Solaris, 103
  - TRUE, 540, 541
  - TRUNCATE, 759
  - TRUNCATE(), 645
  - tutorial, 175
  - Twin Studies
    - queries, 205
  - type conversions, 618, 619
  - types
    - column, 584
    - columns, 615
    - data, 584
    - date, 613
    - Date and Time, 595
    - numeric, 613
    - of tables, 828
    - portability, 615
    - strings, 605
    - time, 613
  - typographical conventions, 2
  - TZ environment variable, 1520, 1886
- ## U
- UCASE(), 635
  - UCS-2, 551
  - UDFs, 1485, 1486
    - compiling, 1492
    - defined, 1484
    - return values, 1492



- ulimit, 1512
  - UMASK environment variable, 1513, 1886
  - UMASK\_DIR environment variable, 1513, 1886
  - unary minus (-), 639
  - unbuffered option
    - mysql, 496
  - UNCOMPRESS(), 681
  - UNCOMPRESSED\_LENGTH(), 681
  - UndoDataBuffer, 950
  - UndoIndexBuffer, 949
  - UNHEX(), 635
  - Unicode, 551
  - Unicode Collation Algorithm, 574
  - UNION, 202, 747
  - Union(), 1054
  - UNIQUE, 700
    - constraint, 38
  - unique ID, 1154
  - Unix, 1162, 1235
  - UNIX\_TIMESTAMP(), 658
  - unloading
    - tables, 183
  - UNLOCK TABLES, 767
  - unnamed views, 754
  - unpack option
    - myisamchk, 483
  - UPDATE, 760
  - update log, 362
  - update-state option
    - myisamchk, 481
  - updating
    - releases of MySQL, 47
    - tables, 33
  - upgrades
    - MySQL Cluster, 970, 970, 971
  - upgrades and downgrades (MySQL Cluster)
    - compatibility between versions, 971
  - upgrading, 124
    - 3.23 to 4.0, 132
    - 4.0 to 4.1, 125
    - different architecture, 137
  - UPPER(), 636
  - uptime, 509
  - URLs for downloading MySQL, 54
  - USE, 764
  - USE INDEX, 740, 746
  - USE KEY, 740, 747
  - use-frm option
    - mysqlcheck, 519
  - USER environment variable, 304, 1886
  - user option
    - mysql, 496
    - mysqlaccess, 506
    - mysqladmin, 511
    - mysqlbinlog, 514
    - mysqlcheck, 519
    - mysqld, 233
    - mysqldump, 526
    - mysqld\_multi, 286
    - mysqld\_safe, 283
    - mysqlhotcopy, 529
    - mysqlimport, 532
    - mysqlshow, 534
    - mysql\_explain\_log, 504
  - user privileges
    - adding, 322
    - deleting, 325, 771
    - dropping, 325, 771
  - user table
    - sorting, 307
  - user variables, 545
  - USER(), 687
  - user-defined functions
    - adding, 1484, 1486
  - User-defined functions, 1485, 1486
  - usernames
    - and passwords, 321
  - users
    - adding, 91, 95
    - deleting, 325, 771
    - root, 121
  - uses
    - of MySQL, 418
  - using multiple disks to start data, 473
  - UTC\_DATE(), 659
  - UTC\_TIME(), 659
  - UTC\_TIMESTAMP(), 659
  - UTF-8, 551
  - utility programs, 474
  - UUID(), 689
- ## V
- valid numbers
    - examples, 540
  - VALUES(), 690
  - VARBINARY data type, 591, 606
  - VARCHAR
    - size, 614
  - VARCHAR data type, 590, 605
  - VARCHARACTER data type, 590
  - variables
    - mysqld, 464
    - server, 234, 806
    - status, 266, 804
    - system, 234, 257, 806
    - user, 545
  - VARIANCE(), 693
  - verbose option
    - myisamchk, 478
    - myisampack, 487
    - myisam\_ftdump, 477
    - mysql, 496
    - mysqladmin, 511
    - mysqlcheck, 519
    - mysqldump, 526
    - mysqld\_multi, 286

- mysqlimport, 532
  - mysqlshow, 534
  - peror, 536
  - version
    - choosing, 44
    - latest, 54
  - version option
    - myisamchk, 478
    - myisampack, 487
    - mysql, 496
    - mysqlaccess, 506
    - mysqladmin, 511
    - mysqlbinlog, 514
    - mysqlcheck, 519
    - mysqld, 233
    - mysqldump, 526
    - mysqld\_multi, 286
    - mysqlimport, 532
    - mysqlshow, 534
    - peror, 536
  - VERSION(), 687
  - vertical option
    - mysql, 496
    - mysqladmin, 511
  - views, 37
    - updatable, 37
  - virtual memory
    - problems while compiling, 102
  - Vision, 1230
  - Visual Objects, 1227
- W**
- wait option
    - myisamchk, 479
    - myisampack, 487
    - mysql, 496
    - mysqladmin, 511
  - WEEK(), 659
  - WEEKDAY(), 661
  - WEEKOFYEAR(), 661
  - Well-Known Binary format, 1040
  - Well-Known Text format, 1040
  - WHERE, 429
  - where option
    - mysqldump, 526
  - widths
    - display, 584
  - Wildcard character (%), 539
  - Wildcard character (\_, 539
  - wildcards
    - and LIKE, 454
    - in mysql.columns\_priv table, 309
    - in mysql.db table, 309
    - in mysql.host table, 309
    - in mysql.procs\_priv table, 309
    - in mysql.tables\_priv table, 309
    - in mysql.user table, 305
  - Windows, 1162, 1235
    - compiling on, 109
    - open issues, 81
    - upgrading, 78
    - versus Unix, 79
  - with-big-tables option, 95
  - Within(), 1056
  - without-server option, 95
  - WKB format, 1040
  - WKT format, 1040
  - wrappers
    - Eiffel, 1159
  - write access
    - tmp, 115
  - write\_buffer\_size myisamchk variable, 479
- X**
- X(), 1049
  - X509/Certificate, 330
  - xml option
    - mysql, 496
    - mysqldump, 526
  - XOR
    - bitwise, 676
    - logical, 624
- Y**
- Y(), 1049
  - Year 2000 compliance, 11
  - Year 2000 issues, 604
  - YEAR data type, 588, 604
  - YEAR(), 661
  - YEARWEEK(), 661