
MySQL 5.1 Release Notes

Abstract

This document contains release notes for the changes in each release of MySQL 5.1, up through MySQL 5.1.73. For information about changes in a different MySQL series, see the release notes for that series.

For additional MySQL 5.1 documentation, see the [MySQL 5.1 Reference Manual](#), which includes an overview of features added in MySQL 5.1 ([What Is New in MySQL 5.1](#)), and discussion of upgrade issues that you may encounter for upgrades from MySQL 5.0 to MySQL 5.1 ([Changes Affecting Upgrades to 5.1](#)).

Updates to these notes occur as new product features are added, so that everybody can follow the development process. If a recent version is listed here that you cannot find on the download page (<http://dev.mysql.com/downloads/>), the version has not yet been released.

The documentation included in source and binary distributions may not be fully up to date with respect to release note entries because integration of the documentation occurs at release build time. For the most up-to-date release notes, please refer to the online documentation instead.

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

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

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

Document generated on: 2016-06-03 (revision: 9034)

Table of Contents

Preface and Legal Notices	3
Changes in MySQL 5.1.73 (2013-12-03)	4
Changes in MySQL 5.1.72 (2013-09-20)	5
Changes in MySQL 5.1.71 (2013-08-01)	6
Changes in MySQL 5.1.70 (2013-06-03)	7
Changes in MySQL 5.1.69 (2013-04-18)	9
Changes in MySQL 5.1.68 (2013-02-05)	11
Changes in MySQL 5.1.67 (2012-12-21)	13
Changes in MySQL 5.1.66 (2012-09-28)	15
Changes in MySQL 5.1.65 (2012-08-09)	17
Changes in MySQL 5.1.64 (Not released)	17
Changes in MySQL 5.1.63 (2012-05-07)	19
Changes in MySQL 5.1.62 (2012-03-21)	20
Changes in MySQL 5.1.61 (2012-01-10)	24
Changes in MySQL 5.1.60 (2011-11-16)	25
Changes in MySQL 5.1.59 (2011-09-15)	27
Changes in MySQL 5.1.58 (2011-07-05)	28
Changes in MySQL 5.1.57 (2011-05-05)	30
Changes in MySQL 5.1.56 (2011-03-01)	32
Changes in MySQL 5.1.55 (2011-02-07)	34
Changes in MySQL 5.1.54 (2010-11-26)	38

Changes in MySQL 5.1.53 (2010-11-03)	39
Changes in MySQL Enterprise 5.1.52sp1 [QSP] (2011-02-21)	41
Changes in MySQL 5.1.52 (2010-10-11)	42
Changes in MySQL 5.1.51 (2010-09-10)	44
Changes in MySQL 5.1.50 (2010-08-03)	48
Changes in MySQL Enterprise 5.1.49sp1 [QSP] (2010-09-28)	50
Changes in MySQL 5.1.49 (2010-07-09)	50
Changes in MySQL 5.1.48 (2010-06-02)	53
Changes in MySQL 5.1.47 (2010-05-06)	56
Changes in MySQL Enterprise 5.1.46sp1 [QSP] (2010-06-23)	58
Changes in MySQL 5.1.46 (2010-04-06)	59
Changes in MySQL 5.1.45 (2010-03-01)	64
Changes in MySQL 5.1.44 (2010-02-04)	67
Changes in MySQL Enterprise 5.1.43sp1 [QSP] (2010-03-25)	69
Changes in MySQL 5.1.43 (2010-01-15)	69
Changes in MySQL 5.1.42 (2009-12-15)	73
Changes in MySQL 5.1.41 (2009-11-05)	77
Changes in MySQL Enterprise 5.1.40sp1 [QSP] (2009-11-25)	83
Changes in MySQL 5.1.40 (2009-10-06)	84
Changes in MySQL 5.1.39 (2009-09-04)	88
Changes in MySQL 5.1.38 (2009-09-01)	91
Changes in MySQL Enterprise 5.1.37sp1 [QSP] (2009-10-10)	96
Changes in MySQL 5.1.37 (2009-07-13)	97
Changes in MySQL 5.1.36 (2009-06-16)	101
Changes in MySQL 5.1.35 (2009-05-13)	106
Changes in MySQL Enterprise 5.1.34sp1 [QSP] (2009-06-25)	110
Changes in MySQL 5.1.34 (2009-04-02)	111
Changes in MySQL 5.1.33 (2009-03-13)	113
Changes in MySQL 5.1.32 (2009-02-14)	119
Changes in MySQL Enterprise 5.1.31sp1 [QSP] (2009-03-19)	122
Changes in MySQL 5.1.31 (2009-01-19)	123
Changes in MySQL 5.1.30 (2008-11-14, General Availability)	132
Changes in MySQL 5.1.29 (2008-10-11)	133
Changes in MySQL 5.1.28 (2008-08-28)	137
Changes in MySQL 5.1.27 (Not released)	142
Changes in MySQL 5.1.26 (2008-06-30)	143
Changes in MySQL 5.1.25 (2008-05-28)	144
Changes in MySQL 5.1.24 (2008-04-08)	147
Changes in MySQL 5.1.23 (2008-01-29)	165
Changes in MySQL 5.1.22 (2007-09-24, Release Candidate)	193
Changes in MySQL 5.1.21 (2007-08-16)	195
Changes in MySQL 5.1.20 (2007-06-25)	208
Changes in MySQL 5.1.19 (2007-05-25)	219
Changes in MySQL 5.1.18 (2007-05-08)	223
Changes in MySQL 5.1.17 (2007-04-04)	236
Changes in MySQL 5.1.16 (2007-02-26)	246
Changes in MySQL 5.1.15 (2007-01-25)	252
Changes in MySQL 5.1.14 (2006-12-05)	261
Changes in MySQL 5.1.13 (Not released)	267
Changes in MySQL 5.1.12 (2006-10-24)	273
Changes in MySQL 5.1.11 (2006-05-26)	306
Changes in MySQL 5.1.10 (Not released)	310
Changes in MySQL 5.1.9 (2006-04-12)	318
Changes in MySQL 5.1.8 (Not released)	320

Changes in MySQL 5.1.7 (2006-02-27)	329
Changes in MySQL 5.1.6 (2006-02-01)	334
Changes in MySQL 5.1.5 (2006-01-10)	338
Changes in MySQL 5.1.4 (2005-12-21)	340
Changes in MySQL 5.1.3 (2005-11-29)	342
Changes in MySQL 5.1.2 (Not released)	343
Changes in MySQL 5.1.1 (Not released)	344

Preface and Legal Notices

This document contains release notes for the changes in each release of MySQL 5.1, up through MySQL 5.1.73.

Legal Notices

Copyright © 1997, 2016, Oracle and/or its affiliates. All rights reserved.

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

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

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

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

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

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and

expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc>.

Access to Oracle Support

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> or visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> if you are hearing impaired.

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 Oracle disseminates it (that is, electronically for download on a Web site 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 Oracle. Oracle and/or its affiliates reserve any and all rights to this documentation not expressly granted above.

Changes in MySQL 5.1.73 (2013-12-03)

- [Packaging Notes](#)
- [Bugs Fixed](#)

Packaging Notes

- Previously, MySQL Server distributions included the MySQL Reference Manual in Info format (the Docs/mysql.info file). Because the license for the manual restricts redistribution, its inclusion in Community packages caused problems for downstream redistributors, such as those who create Linux distributions. Community distributions of MySQL Server no longer include the mysql.info file, to make the repackaging and redistribution process easier (for example, the source tarball and its checksum can be used directly). This change applies to all source and binary Community packaging formats. Commercial (Enterprise) distributions are unchanged.

For those who wish to continue using the MySQL Reference Manual in Info format, we have made it available at <http://dev.mysql.com/doc/>.

Bugs Fixed

- **InnoDB:** In debug builds, test case failures would occur due to `ibuf_contract_ext` performing merges and `dict_stats_update` returning evicted pages back into the buffer pool while `ibuf_change_buffering_debug` is enabled. (Bug #17446090)

- **InnoDB:** [InnoDB](#) failed to return an error when attempting to run a query after discarding the tablespace. (Bug #17431533)
- **InnoDB:** When the change buffer is enabled, [InnoDB](#) failed to write a transaction log record when merging a record from the insert buffer to a secondary index page if the insert was performed as an “update-in-place”. (Bug #16752251, Bug #69122)
- **Partitioning:** The storage engine was set incorrectly during a rebuild of a partition; the table storage engine was ignored and the default storage engine used instead. Thus, in MySQL 5.1, it was possible for [REBUILD PARTITION](#) to change the partition storage engine from [InnoDB](#) to [MyISAM](#), and for the reverse (rebuilding partitions of [MyISAM](#) tables causing the partitions to use [InnoDB](#)) to occur in MySQL 5.5 and later. Now, when rebuilding partitions, the storage engine actually used by the table is checked and used by the handler for the rebuild operation, so that the partition storage engine is not inadvertently changed. (Bug #17559867)
- **Replication:** When an error encountered by the dump thread while reading events from the active binary log file was a temporary error, so that the dump thread tried to read the event, it was possible for the dump thread to seek the wrong position, which could cause one or more events to be resent. To prevent this, the thread's position is obtained after each correct read of an event.

In addition, with this fix, only binary logs that are not closed normally are marked as possibly being corrupted.

Finally, two warnings are added; these are now returned when a dump thread encounters a temporary error. (Bug #17402313)

- **Replication:** The value of [LAST_INSERT_ID\(\)](#) was not correctly replicated when filtering rules were used on the slave. (Bug #17234370, Bug #69861)
- Enabling Index Merge optimizer switches and setting a small [sort_buffer_size](#) value could lead to a server exit. (Bug #17617945)
- The [filesort](#) implementation sometimes failed to allocate enough buffer space, leading to a server exit. (Bug #17326567)
- The [mysql_options\(\)](#) C API function could leak memory if called more than once with the [MYSQL_SET_CLIENT_IP](#) option. (Bug #17297012)
- An error array in the SSL code was missing a comma, leading to implicit concatenation of adjacent messages and a resulting off-by-one error in the relationship between error numbers and messages. (Bug #17294150)
- Very long database names in queries could cause the server to exit. (Bug #15912213, Bug #16900358)
- The [my_b_vprintf\(\)](#) function could produce incorrect results for long integers on 64-bit systems. (Bug #67386, Bug #16978278)
- Host names in grant tables are stored in lowercase, but [mysql_install_db](#) could fail to observe this convention, leading to accounts that could not be dropped with [DROP USER](#). (Bug #62255, Bug #12917164, Bug #62254, Bug #12917151)

Changes in MySQL 5.1.72 (2013-09-20)

Bugs Fixed

- **InnoDB:** The [row_sel_sec_rec_is_for_clust_rec](#) function would incorrectly prepare to compare a NULL column prefix in a secondary index with a non-NULL column in a clustered index. (Bug #17312846)

- **InnoDB:** An incorrect purge would occur when rolling back an update to a delete-marked record. (Bug #17302896)
- **InnoDB:** `InnoDB` would rename a user-defined foreign key constraint containing the string “_ibfk_” in its name, resulting in a duplicate constraint. (Bug #17076737, Bug #69693, Bug #17076718, Bug #69707)
- **InnoDB:** Rolling back an `INSERT` after a failed `BLOB` write would result in an assertion failure. The assertion has been modified to allow NULL `BLOB` pointers if an error occurs during a `BLOB` write. (Bug #16971045)
- **InnoDB:** The `srv_master_thread` background thread, which monitors server activity and performs activities such as page flushing when the server is inactive or in a shutdown state, runs on a one second delay loop. `srv_master_thread` failed to check if the server is in a shutdown state before sleeping. (Bug #13417564, Bug #63276)
- **InnoDB:** An infinite loop could occur in `buf_page_get_gen` when handling compressed-only pages. (Bug #12560151, Bug #61132)
- Within a stored program, comparison of the value of a scalar subquery with an `IN` clause resulted in an error for the first execution and raised an assertion for the second execution. (Bug #17029399)
- The `my_strtoll10()` function could incorrectly convert some long string-format numbers to numeric values and fail to set the overflow flag. (Bug #16997513)
- For queries that accessed an `INFORMATION_SCHEMA` table in a subquery, an attempt to lock a mutex that had already been locked could cause a server crash. (Bug #11765744)
- `mysqldump` wrote `SET` statements as `SET OPTION`, which failed when reloaded because the deprecated `OPTION` keyword has been removed from `SET` syntax. (Bug #67507, Bug #15844882)
- For `DIV` expressions, assignment of the result to multiple variables could cause a server crash. (Bug #59241, Bug #11766191)

References: See also: Bug #8457.

- If one connection changed its default database and simultaneously another connection executed `SHOW PROCESSLIST`, the second connection could access invalid memory when attempting to display the first connection's default database. memory. (Bug #58198, Bug #11765252)

Changes in MySQL 5.1.71 (2013-08-01)

- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

Functionality Added or Changed

- `comp_err` now checks to make sure that new errors are not being added to MySQL 5.1 or 5.5 because the set of errors for these series is frozen. (Bug #16807394)

Bugs Fixed

- **InnoDB:** During an insert buffer merge, InnoDB would invoke `lock_rec_restore_from_page_infimum()` on a potentially invalid record pointer. (Bug #16806366)
- **InnoDB:** The `page_zip_validate()` consistency check failed after compressing a page, in `page_zip_compress()`. This problem was caused by `page_zip_decompress()`, which failed to set

`heap_no` correctly when a record contained no user data bytes. A record with no user data bytes occurs when, for example, a primary key is an empty string and all secondary index fields are NULL or an empty string. (Bug #16736929)

- **InnoDB:** The `pthread_mutex`, `commit_threads_m`, which was initialized but never used, has been removed from the code base. (Bug #60225, Bug #11829813)
- **Partitioning:** When dropping a partitioned table, the table's `.par` file was deleted first, before the table definition or data. This meant that, if the server failed during the drop operation, the table could be left in an inconsistent state in which it could neither be accessed nor dropped.

The fix for this problem makes the following changes:

- Now, when dropping a partitioned table, the table's `.par` file is not removed until all table data has been deleted.
- When executing `DROP TABLE` of a partitioned table, in the event that its `.par` file is determined to be missing, the table's `.frm` file is now immediately deleted, in effect forcing the drop to complete.

(Bug #13548704, Bug #63884)

- Shared-compatibility conflict errors occurred for RPM install operations, even if no shared-compatibility RPMs were already installed. (Bug #16678122)
- A user variable referenced during execution of a prepared statement is set to memory that is freed at the end of execution. A second execution of the statement could result in Valgrind warnings when accessing this memory. (Bug #16119355)
- Misoptimization of left expressions in prepared statements could cause a server exit. (Bug #16095534)
- Subsequent to `Prepared statement needs to be re-prepared` errors, inserts into `DECIMAL` columns caused a server exit. (Bug #12608543)
- Assigning the result of a subquery to a user variable raised an assertion when the outer query included `DISTINCT` and `GROUP BY`. (Bug #57196, Bug #11764371)

Changes in MySQL 5.1.70 (2013-06-03)

Bugs Fixed

- **Important Change; Replication:** When the server was running with `--binlog-ignore-db` and `SELECT DATABASE()` returned `NULL` (that is, there was no currently selected database), statements using fully qualified table names in `dbname.tblname` format were not written to the binary log. This was because the lack of a currently selected database in such cases was treated as a match for any possible ignore option rather than for no such option; this meant that these statements were always ignored.

Now, if there is no current database, a statement using fully qualified table names is always written to the binary log. (Bug #11829838, Bug #60188)

- **InnoDB:** The fix for Bug #16722314 resulted in a linker error. (Bug #16798595)
- **InnoDB:** Valgrind testing returned memory leak errors which resulted from a regression introduced by the fix for Bug #11753153. The `dict_create_add_foreign_to_dictionary` function would call `pars_info_create` but failed to call `pars_info_free`. (Bug #16754901)
- **InnoDB:** Some characters in the identifier for a `foreign key constraint` are modified during table exports. (Bug #16722314, Bug #69062)

- **InnoDB:** Crash recovery failed with a `!recv_no_log_write` assertion when reading a page. (Bug #16405422)
- **InnoDB:** When tables are linked by foreign key constraints, loading one table would open other linked tables recursively. When numerous tables are linked by foreign key constraints, this would sometimes lead to a thread stack overflow causing the server to exit. Tables linked by foreign key constraints are now loaded iteratively. Cascade operations, which were also performed in a recursive manner, are now performed iteratively using an explicit stack. (Bug #16244691, Bug #65384)
- **Replication:** Using the `--replicate-*` options (see [Replication Slave Options and Variables](#)) could in some cases lead to a memory leak on the slave. (Bug #16056813, Bug #67983)
- **Replication:** The binary log contents got corrupted sometimes, because the function `MYSQL_BIN_LOG::write_cache` always thought it had reached the end-of-cache when the function `my_b_fill()` reported a '0,' while that could also mean an error had occurred. This fix makes sure that whenever `my_b_fill()` returns a '0,' an error check is performed on `info->error`. (Bug #14324766, Bug #60173)
- The WKB reader for spatial operations could fail and cause a server exit. (Bug #16451878)
- For debug builds, `GROUP_CONCAT(... ORDER BY)` within an `ORDER BY` clause could cause a server exit. (Bug #16347426)
- A `GROUP_CONCAT()` invocation containing subquery having an outer reference caused the server to exit. (Bug #16347343)
- If loose index scan was used on a query that used `MIN()`, a segmentation fault could occur. (Bug #16222245)
- A prepared statement that used `GROUP_CONCAT()` and an `ORDER BY` clause that named multiple columns could cause the server to exit. (Bug #16075310)
- `ORDER BY MATCH ... AGAINST` could cause a server exit. (Bug #16073689)
- When a partition is missing, code in `ha_innodb.cc` would retry 10 times and sleep for a microsecond each time while holding `LOCK_open`. The retry logic for partitioned tables was introduced as a fix for Bug#33349 but did not include a test case to validate it. This fix removes the retry logic for partitioned tables. If the problem reported in Bug#33349 reappears, a different solution will be explored. (Bug #15973904)
- The `mysql.server` script exited with an error if the `status` command was executed with multiple servers running. (Bug #15852074)
- When processing row-based-replication events in the old binary log format from prior to MySQL 5.1 GA builds, `mysqlbinlog` could result in out-of-bounds heap buffer reads and undefined behaviour. (Bug #14771299)
- The `mysql` client allocated but did not free a string after reading each line in interactive mode, resulting in a memory leak. (Bug #14685362)
- Grouping by an outer `BLOB` column in a subquery caused a server exit. (Bug #13966809, Bug #14700180)
- The test for stack overrun did not work for recent `gcc` versions and could lead to server exit. (Bug #62856, Bug #13243248)

References: See also: Bug #42213.

- The `url` columns in the `mysql` database help tables were too short to hold some of the URLs in the help content. For new installations, these columns are now created as type `TEXT` to accommodate longer URLs.

For upgrades, `mysql_upgrade` does *not* update the columns. Modify them manually using these statements:

```
ALTER TABLE mysql.help_category MODIFY url TEXT NOT NULL;  
ALTER TABLE mysql.help_topic MODIFY url TEXT NOT NULL;
```

(Bug #61520, Bug #12671635)

- `IF()` function evaluations could produce different results when executed in a prepared versus nonprepared statement. (Bug #45370, Bug #11753852)

Changes in MySQL 5.1.69 (2013-04-18)

- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

Functionality Added or Changed

- MySQL no longer uses the default OpenSSL compression. (Bug #16235681)

Bugs Fixed

- **Performance; InnoDB:** The `DROP TABLE` statement for a table using `compression` could be slower than necessary, causing a stall for several seconds. MySQL was unnecessarily decompressing `pages` in the `buffer pool` related to the table as part of the `DROP` operation. (Bug #16067973)
- **Important Note; Replication:** Using row-based logging to replicate from a table to a same-named view led to a failure on the slave. Now, when using row-based logging, the target object type is checked prior to performing any DML, and an error is given if the target on the slave is not actually a table.



Note

It remains possible to replicate from a table to a same-named view using statement-based logging.

(Bug #11752707, Bug #43975)

- **InnoDB:** The `page_zip_available` function would count some fields twice. (Bug #16463505)
- **InnoDB:** For `InnoDB` tables, if a `PRIMARY KEY` on a `VARCHAR` column (or prefix) was empty, index page compression could fail. (Bug #16400920)
- **InnoDB:** For debug builds, `InnoDB` status exporting was subject to a race condition that could cause a server exit. (Bug #16292043)
- **InnoDB:** Arithmetic underflow during page compression for `CREATE TABLE` on an `InnoDB` table could cause a server exit. (Bug #16089381)
- **InnoDB:** This fix makes MySQL more responsive to `KILL QUERY` statements when the query is accessing an `InnoDB` table. (Bug #14704286)
- **InnoDB:** When printing out long semaphore wait diagnostics, `sync_array_cell_print()` ran into a segmentation violation (SEGV) caused by a race condition. This fix addresses the race condition by allowing the cell to be freed while it is being printed. (Bug #13997024)

- **InnoDB:** Killing a query caused an InnoDB assertion failure when the same table (cursor) instance was used again. This is the result of a regression error introduced by the fix for Bug#14704286. The fix introduced a check to handle kill signals for long running queries but the cursor was not restored to the proper state. (Bug #68051, Bug #16088883)
- **InnoDB:** The length of internally generated foreign key names was not checked. If internally generated foreign key names were over the 64 character limit, this resulted in invalid DDL from SHOW CREATE TABLE. This fix checks the length of internally generated foreign key names and reports an error message if the limit is exceeded. (Bug #44541, Bug #11753153)
- **Partitioning:** A query on a table partitioned by range and using `TO_DAYS()` as a partitioning function always included the first partition of the table when pruning. This happened regardless of the range employed in the `BETWEEN` clause of such a query. (Bug #15843818, Bug #49754)
- **Replication:** A zero-length name for a user variable (such as `@``) was incorrectly considered to be a sign of data or network corruption when reading from the binary log. (Bug #16200555, Bug #68135)
- **Replication:** Backtick (```) characters were not always handled correctly in internally generated SQL statements, which could sometimes lead to errors on the slave. (Bug #16084594, Bug #68045)

References: This issue is a regression of: Bug #14548159, Bug #66550.

- **Replication:** It was possible in certain cases—immediately after detecting an EOF in the dump thread read event loop, and before deciding whether to change to a new binary log file—for new events to be written to the binary log before this decision was made. If log rotation occurred at this time, any events that occurred following EOF detection were dropped, resulting in loss of data. Now in such cases, steps are taken to make sure that all events are processed before allowing the log rotation to take place. (Bug #13545447, Bug #67929)

References: See also: Bug #16016886.

- A long database name in a `GRANT` statement could cause the server to exit. (Bug #16372927)
- Incorrect results were returned if a query contained a subquery in an `IN` clause which contained an `XOR` operation in the `WHERE` clause. (Bug #16311231)
- yaSSL did not perform proper padding checks, but instead examined only the last byte of cleartext and used it to determine how many bytes to remove. (Bug #16218104)
- Invocation of the range optimizer for a `NULL` select caused the server to exit. (Bug #16192219)
- `SHOW COLUMNS` on a view defined as a `UNION` of `Geometry` columns could cause the server to exit. (Bug #14362617)
- A `LIKE` pattern with too many `'%'` wildcards could cause a segmentation fault. (Bug #14303860)
- `SET var_name = VALUES(col_name)` could cause the server to exit. This syntax is now prohibited because in `SET` context there is no column name and the statement returns `ER_BAD_FIELD_ERROR`. (Bug #14211565)
- The `COM_CHANGE_USER` command in the client/server protocol did not properly use the character set number in the command packet, leading to incorrect character set conversion of other values in the packet. (Bug #14163155)
- Subqueries with `OUTER JOIN` could return incorrect results if the subquery referred to a column from another `SELECT`. (Bug #13068506)
- `Field_geom::reset()` failed to reset its base `Field_blob`. The range optimizer used the uninitialized field during optimization and execution, causing the server to exit. (Bug #11908153)

- `mysql_install_db` did not escape `'_'` in the host name for statements written to the grant tables. (Bug #11746817)
- Adjusted MySQL configuration to account for change in Automake 1.12 that produced `sql_yacc.hh` rather than `sql_yacc.h` as expected by `sql/Makefile.am`. (Bug #67177, Bug #15967374)
- If a dump file contained a view with one character set and collation defined on a view with a different character set and collation, attempts to restore the dump file failed with an “illegal mix of collations” error. (Bug #65382, Bug #14117025)
- Incorrect metadata could be produced for columns returned from some views. (Bug #65379, Bug #14096619)
- For debug builds, some queries with `SELECT ... FROM DUAL` nested subqueries raised an assertion. (Bug #60305, Bug #11827369)
- `PARTITION BY KEY` on a `utf32 ENUM` column raised a debugging assertion. (Bug #52121, Bug #11759782)
- `UNION ALL` on `BLOB` columns could produce incorrect results. (Bug #50136, Bug #11758009)
- The `REPLACE()` function produced incorrect results when a user variable was supplied as an argument and the operation was performed on multiple rows. (Bug #49271, Bug #11757250)
- Setting `max_connections` to a value less than the current number of open connections caused the server to exit. (Bug #44100, Bug #11752803)
- The optimizer used loose index scan for some queries for which this access method is inapplicable. (Bug #42785, Bug #11751794)
- View access in low memory conditions could raise a debugging assertion. (Bug #39307, Bug #11749556)

Changes in MySQL 5.1.68 (2013-02-05)

Bugs Fixed

- **Performance; InnoDB:** Some data structures related to undo logging could be initialized unnecessarily during a query, although they were only needed under specific conditions. (Bug #14676084)
- **Performance; InnoDB:** Optimized read operations for `compressed` tables by skipping redundant tests. The check for whether any related changes needed to be merged from the `insert buffer` was being called more often than necessary. (Bug #14329288, Bug #65886)
- **Performance; InnoDB:** Immediately after a table was created, a query against it would not use a `loose index scan`. The same query might use a loose index scan following an `ALTER TABLE` on the table. The fix improves the accuracy of the cost estimate for queries involving the grouping functions `min()` and `max()`, and prevents the query plan from being changed by the `ALTER TABLE` statement. (The more stable query plan might or might not use a loose index scan.) (Bug #14200010)
- **InnoDB; Partitioning:** Previously, when attempting to optimize one or more partitions of a partitioned table that used a storage engine that does not support partition-level `OPTIMIZE`, such as `InnoDB`, MySQL reported `Table does not support optimize, doing recreate + analyze instead`, then re-created the entire table, but did not actually analyze it. Now in such cases, the warning message is, `Table does not support optimize on partitions. All partitions will be rebuilt and analyzed`. In addition, the entire table is analyzed after first being rebuilt. (Bug #11751825, Bug #42822)

- **InnoDB:** Creating an index on a [CHAR](#) column could fail for a table with a character set with varying length, such as [UTF-8](#), if the table was created with the [ROW_FORMAT=REDUNDANT](#) clause. (Bug #15874001)
- **InnoDB:** The status variable [Innodb_buffer_pool_read_ahead_evicted](#) could show an inaccurate value, higher than expected, because some pages in the [buffer pool](#) were incorrectly considered as being brought in by [read-ahead](#) requests. (Bug #15859402, Bug #67476)
- **InnoDB:** A regression introduced by the fix for Bug#14100254 would result in a “!BPAGE->FILE_PAGE_WAS_FREED” assertion. (Bug #14676249)
- **InnoDB:** If the server crashed at a precise moment during an [ALTER TABLE](#) operation that rebuilt the [clustered index](#) for an [InnoDB](#) table, the original table could be inaccessible afterward. An example of such an operation is [ALTER TABLE ... ADD PRIMARY KEY](#). The fix preserves the original table if the server halts during this operation. You might still need to rename the [.ibd](#) file manually to restore the original table contents: in MySQL 5.6 and higher, rename from [#sql-ib\\$new_table_id.ibd](#) to [table_name.ibd](#) within the database directory; prior to MySQL 5.6, the temporary file to rename is [table_name#1](#) or [#2](#). (Bug #14669848)
- **InnoDB:** An error at the filesystem level, such as too many open files, could cause an unhandled error during an [ALTER TABLE](#) operation. The error could be accompanied by Valgrind warnings, and by this assertion message:

```
Assertion '! is_set()' failed.  
mysqld got signal 6 ;
```

(Bug #14628410, Bug #16000909)

- **InnoDB:** During shutdown, with the [innodb_purge_threads](#) configuration option set greater than 1, the server could halt prematurely with this error:

```
mysqld got signal 11
```

A workaround was to increase [innodb_log_file_size](#) and set [innodb_purge_threads=1](#). The fix was backported to MySQL 5.5 and 5.1, although those versions do not have the [innodb_purge_threads](#) configuration option so the error was unlikely to occur. (Bug #14234028)

- **InnoDB:** The value of the [innodb_version](#) variable was not updated consistently for all server releases for the InnoDB Plugin in MySQL 5.1, and the integrated [InnoDB](#) component in MySQL 5.5, 5.6, and higher. Since [InnoDB](#) and MySQL Server development cycles are fully integrated and synchronized, now the value returned by the [innodb_version](#) variable is the same as for the [version](#) variable. (Bug #13463493, Bug #63435)
- **Partitioning:** When used with a table having multiple columns in its primary key, but partitioned by [KEY](#) using a column that was not part of the primary key as the partitioning column, a query using an aggregate function and [DISTINCT](#) such as [SELECT SUM\(DISTINCT pk_column_1\) FROM table WHERE pk_column_2 = constant](#) was not handled correctly. (Bug #14845133)

References: See also: Bug #14495351. This issue is a regression of: Bug #13025132.

- **Partitioning:** Concurrent [ALTER TABLE ... REBUILD PARTITION](#) operations could interfere with one another, even when not running against the same table, because they both used global memory for storage. Now each partition rebuild operation stores intermediate data in memory that is local to that process. (Bug #14589559, Bug #66645)

- **Replication:** Repeated execution of `CHANGE MASTER TO` statements using invalid `MASTER_LOG_POS` values could lead to errors and possibly a crash on the slave. Now in such cases, the statement fails with a clear error message. (Bug #11764602, Bug #57454)
 - **Replication:** If the disk becomes full while writing to the binary log, the server hangs until space is freed up manually. It was possible after this was done for the MySQL server to fail, due to an internal status value being set when not needed. Now in such cases, rather than trying to set this status, a warning is written in the error log instead. (Bug #11753923, Bug #45449)
 - **Microsoft Windows:** Dynamic file names (with colons) are no longer allowed. Static file names using the Alternate Data Stream (ADS) NTFS functionality of Microsoft Windows may continue to be used. (Bug #11761752)
 - Directory name manipulation could result in stack overflow on Mac OS X and Windows. (Bug #16066243)
 - A buffer-handling problem in yaSSL was fixed. (Bug #15965288)
 - Metadata locking and table definition cache routines did not always check length of names passed to them. (Bug #15954872)
 - It was possible in theory for `UpdateXML()` to return `NULL` incorrectly. (Bug #15948580)
- References: See also: Bug #13007062.
- Enabling the query cache during high client contention could cause the server to exit. (Bug #14727815)
 - The server sometimes failed to respect `MAX_CONNECTIONS_PER_HOUR` limits on user connections. (Bug #14627287)
 - Passing an unknown time zone specification to `CONVERT_TZ()` resulted in a memory leak. (Bug #12347040)
 - `mysqld_safe` used the nonportable `-e` test construct. (Bug #67976, Bug #16046140)
 - For dumps of the `mysql` database, `mysqldump` skips the `event` table unless the `--events` option is given. `mysqldump` now prints a warning if invoked without `--events` that the `mysql.event` table is not dumped without that option. (Bug #55587, Bug #11762933)
 - For `MEMORY` tables with `HASH` indexes, `DELETE` sometimes failed to delete all applicable rows. (Bug #51763, Bug #11759445)
 - `UNION` type conversion could incorrectly turn unsigned values into signed values. (Bug #49003, Bug #11757005)
 - During the startup process, `mysqld` could incorrectly remove the PID file of an already running `mysqld`. (Bug #23790, Bug #11746142)

References: See also: Bug #14726272.

Changes in MySQL 5.1.67 (2012-12-21)

Bugs Fixed

- **Performance; InnoDB:** The timing values for low-level `InnoDB` read operations were adjusted for better performance with fast storage devices, such as `SSD`. This enhancement primarily affects read operations for `BLOB` columns in `compressed` tables. (Bug #13702112, Bug #64258)

- **Incompatible Change:** `LAST_INSERT_ID(expr)` did not work for `expr` values greater than the largest signed `BIGINT` value. Such arguments now are accepted, with some consequences for compatibility with previous versions:

- `LAST_INSERT_ID()` now returns a `BIGINT UNSIGNED` value, not a `BIGINT` (signed) value.
- `LAST_INSERT_ID(expr)` now returns an unsigned integer value, not a signed integer value.
- For `AUTO_INCREMENT` columns, negative values are no longer supported.

(Bug #20964, Bug #11745891)

- **InnoDB:** An online DDL operation for an `InnoDB` table incorrectly reported an empty value (``) instead of the correct key value when it reported a duplicate key error for a unique index using an index prefix. (Bug #14729221)
- **InnoDB:** If a `CREATE TABLE` statement failed due to a disk full error, some memory allocated during the operation was not freed properly. (Bug #14708715)
- **InnoDB:** If the server crashed at the specific point when a `change buffer` entry was being merged into a buffer pool page, the transaction log and the change buffer were left in an inconsistent state. After a restart, MySQL could crash after reading the corresponding secondary index page. The problem was more likely to occur in MySQL 5.5 or later, where the original `insert buffering` mechanism was generalized to cover other operations. (Bug #14636528, Bug #66819, Bug #58571, Bug #61104, Bug #65443)
- **InnoDB:** In rare circumstances, MySQL could apply `InnoDB undo` records out of order during a `ROLLBACK` of an operation that modified a BLOB column. This issue could cause an assertion error in debug builds:

```
!bpage->file_page_was_freed
```

(Bug #13249921)

- **InnoDB:** In debug builds, a mismatch in the `InnoDB PAGE_FREE` list would cause an assertion. (Bug #12701488)
- **Replication:** Updates writing user variables whose values were never set on a slave while using `--replicate-ignore-table` could cause the slave to fail. (Bug #14597605)

References: This issue is a regression of: Bug #14275000.

- **Replication:** Following an insert into a nontransactional table that failed due to insufficient disk space, the server did not properly clean up all pending events, leading to an assert or possibly to other errors. (Bug #11750014)
- **Replication:** Backtick (`) characters were not always handled correctly in internally generated SQL statements, which could sometimes lead to errors on replication slaves or cause failure of restore operations from binary log files. (Bug #66550, Bug #14548159, Bug #29422, Bug #11746883)
- Within a stored procedure, executing a multiple-table `DELETE` statement that used a very long table alias could cause the server to exit. (Bug #15954896)
- Very long database names in queries could cause the server to exit. (Bug #15912213, Bug #16900358)
- Attempting to create an `auto-increment` column in an `InnoDB` table with a `NULL` type attribute could cause a serious error. (Bug #14758479)

- A `DELETE` statement for an `InnoDB` table could write incorrect transaction metadata into a record, causing the server to halt with an error. To work around this issue, reduce the specified length of the primary key to less than 1K bytes. (Bug #14731482)
- Repeated execution of a query containing a subquery that used `MAX()` could result in increasing memory consumption. (Bug #14683676)
- `USE dbname` could fail with `Unknown database` when `dbname` contained multiple backtick (```) characters. (Bug #14645196)
- `SHOW PROFILE` could be used to cause excessive server memory consumption. (Bug #14629232)
- The thread cache implementation worked in LIFO rather than FIFO fashion and could result in a thread being denied service (although this was a remote possibility). (Bug #14621627)
- `CREATE USER` and `DROP USER` could fail to flush the privileges, requiring `FLUSH PRIVILEGES` to be used explicitly. (Bug #13864642)
- A memory leak could occur for queries containing a subquery that used `GROUP BY` on an outer column. (Bug #13724099)
- A “buffer too small” error message from the `myisamchk` command referred to the `myisam_sort_buffer_size` configuration option, when it should have referred to `sort_buffer_size`.

`myisamchk` now has a `myisam_sort_buffer_size` variable available as an alternative name to `sort_buffer_size`. `myisam_sort_buffer_size` is preferable to `sort_buffer_size` because its name corresponds to the `myisam_sort_buffer_size` server system variable that has a similar meaning. `sort_buffer_size` should be considered deprecated. (Bug #11754894, Bug #46578)
- The number of connection errors from a given host as counted by the server was periodically reset, with the result that `max_connect_errors` was never reached and invalid hosts were never blocked from trying to connect. (Bug #11753779)

References: See also: Bug #38247, Bug #43006, Bug #45584, Bug #45606.

- On Windows, the Perl version of `mysql_install_db` created system tables in the `mysql` database that were not populated properly. (Bug #65584, Bug #14181049)
- `mysqld_safe` ignored the value of the `UMASK` environment variable, leading to behavior different from `mysqld` with respect to the access mode of created files. Now `mysqld_safe` (and `mysqld_multi`) attempt to approximate the same behavior as `mysqld`. (Bug #57406, Bug #11764559)
- During optimization, `ZEROFILL` values may be converted to string constants. However, `CASE` expressions did not handle switching data types after the planning stage, leading to `CASE` finding a null pointer instead of its argument. (Bug #57135, Bug #11764313)

Changes in MySQL 5.1.66 (2012-09-28)

Bugs Fixed

- **InnoDB:** Inserting data of varying record lengths into an `InnoDB` table that used `compression` could cause the server to halt with an error. (Bug #14554000, Bug #13523839, Bug #63815, Bug #12845774, Bug #61456, Bug #12595091, Bug #61208)
- **InnoDB:** Under heavy load of concurrent `DML` and queries, an `InnoDB` table with a unique index could return nonexistent duplicate rows to a query. (Bug #14399148, Bug #66134)

- **InnoDB:** Deleting from an [InnoDB](#) table containing a [prefix](#) index, and subsequently dropping the index, could cause a crash with an assertion error. (Bug #13807811)
- **InnoDB:** The error message was improved for the case where an [UPDATE](#) failed because the row included several BLOB values greater than 768 bytes each, causing the size of a row to exceed half the [page size](#). The old message, was misleading; it suggested using BLOBs, when the 768-byte prefix for each BLOB column was the cause of the limit error:

```
Error Code 1118: Row size too large. The maximum row size for the used table
type, not counting BLOBs, is 8126. You have to change some columns to TEXT or
BLOBs
```

A workaround for the problem was to create the table with the [ROW_FORMAT=DYNAMIC](#) or [ROW_FORMAT=COMPRESSED](#) clause, which is now suggested in the message. (Bug #13453036, Bug #63507)

- **InnoDB:** Certain [INFORMATION_SCHEMA](#) tables originally introduced in MySQL 5.6 are now also available in MySQL 5.5 and MySQL 5.1: [INNODB_BUFFER_PAGE](#), [INNODB_BUFFER_PAGE_LRU](#), and [INNODB_BUFFER_POOL_STATS](#). (Bug #13113026)
- **InnoDB:** When a [SELECT ... FOR UPDATE](#), [UPDATE](#), or other SQL statement scanned rows in an [InnoDB](#) table using a [<](#) or [<=](#) operator in a [WHERE](#) clause, the next row after the affected range could also be locked. This issue could cause a lock wait timeout for a row that was not expected to be locked. The issue occurred under various isolation levels, such as [READ COMMITTED](#) and [REPEATABLE READ](#). (Bug #11765218)
- **Partitioning:** The buffer for the row currently read from each partition used for sorted reads was allocated on open and freed only when the partitioning handler was closed or destroyed. For [SELECT](#) statements on tables with many partitions and large rows, this could cause the server to use excessive amounts of memory.

This issue has been addressed by allocating buffers for reads from partitioned tables only when they are needed and freeing them immediately once they are no longer needed. As part of this fix, memory is now allocated for reading from rows only in partitions that have not been pruned (see [Partition Pruning](#)). (Bug #13025132)

References: See also: Bug #11764622, Bug #14537277.

- **Replication:** In master-master replication with [--log-slave-updates](#) enabled, setting a user variable and then performing inserts using this variable caused the [Exec_master_log_position](#) column in the output of [SHOW SLAVE STATUS](#) not to be updated. (Bug #13596613)
- When resolving outer fields, [Item_field::fix_outer_fields\(\)](#) creates new [Item_refs](#) for each execution of a prepared statement, so these must be allocated in the runtime memroot. The memroot switching before resolving [JOIN::having](#) caused these to be allocated in the statement root, leaking memory for each prepared statement execution. (Bug #14409015)
- Small [sort_buffer_size](#) values could result in a server crash. (Bug #14111180)
- The [libmysqlclient_r](#) client library exported symbols from [yaSSL](#) that conflict with [OpenSSL](#). If a program linked against that library and [libcurl](#), it could crash with a segmentation fault. (Bug #14068244, Bug #65055, Bug #14072299)
- The argument for [LIMIT](#) must be an integer, but if the argument was given by a placeholder in a prepared statement, the server did not reject noninteger values such as ['5'](#). (Bug #13868860)
- Access to [INFORMATION_SCHEMA](#) tables through a view could leak memory. (Bug #13734987)

- A query for a [FEDERATED](#) table could return incorrect results when the underlying table had a compound index on two columns and the query included an [AND](#) condition on the columns. (Bug #12876932)
- In debug builds, an [InnoDB](#) assertion was overly aggressive about prohibiting an open range. (Bug #66513, Bug #14547952)
- The argument to the `--ssl-key` option was not verified to exist and be a valid key. The resulting connection used SSL, but the key was not used. (Bug #62743, Bug #13115401)
- Adding a [LIMIT](#) clause to a query containing [GROUP BY](#) and [ORDER BY](#) could cause the optimizer to choose an incorrect index for processing the query, and return more rows than required. (Bug #54599, Bug #11762052)
- [mysqlbinlog](#) did not accept input on the standard input when the standard input was a pipe. (Bug #49336, Bug #11757312)

Changes in MySQL 5.1.65 (2012-08-09)

- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

Functionality Added or Changed

- **Important Change:** The [YEAR\(2 \)](#) data type is now deprecated because it is problematic. Support for [YEAR\(2 \)](#) will be removed in a future MySQL release. For more information, see [YEAR\(2\) Limitations and Migrating to YEAR\(4\)](#).

Bugs Fixed

- The server did not build with [gcc](#) 4.7. (Bug #14238406)

Changes in MySQL 5.1.64 (Not released)

- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

Functionality Added or Changed

- **Important Change; Replication:** The [SHOW BINARY LOGS](#) statement (and its equivalent [SHOW MASTER LOGS](#)) may now be executed by a user with the [REPLICATION CLIENT](#) privilege. (Formerly, the [SUPER](#) privilege was necessary to use either form of this statement.)

Bugs Fixed

- **InnoDB:** If a row was deleted from an [InnoDB](#) table, then another row was re-inserted with the same primary key value, an attempt by a concurrent transaction to lock the row could succeed when it should have waited. This issue occurred if the locking select used a [WHERE](#) clause that performed an index scan using a secondary index. (Bug #14100254, Bug #65389)
- **InnoDB:** In a transaction using the [REPEATABLE READ](#) isolation level, an [UPDATE](#) or [DELETE](#) statement for an [InnoDB](#) table could sometimes overlook rows recently committed by other transactions. As explained in [Consistent Nonlocking Reads](#), DML statements within a [REPEATABLE READ](#) transaction apply to rows committed by other transactions, even if a query could not see those rows. (Bug #14007649, Bug #65111)

- **InnoDB:** Performing an `ALTER TABLE` operation on an `InnoDB` could cause the server to halt with an error, if the `tablespace` for that table was already removed by an `ALTER TABLE ... DISCARD TABLESPACE` statement. (Bug #13943231)
- **InnoDB:** Using the `KILL` statement to terminate a query could cause an unnecessary message in the error log:

```
[ERROR] Got error -1 when reading table table_name
```

(Bug #13933132)

- **InnoDB:** For an `InnoDB` table with a trigger, under the setting `innodb_autoinc_lock_mode=1`, sometimes auto-increment values could be interleaved when inserting into the table from two sessions concurrently. The sequence of auto-increment values could vary depending on timing, leading to data inconsistency in systems using replication. (Bug #12752572, Bug #61579)
- **InnoDB:** The `CHECK TABLE` statement could fail for a large `InnoDB` table due to a timeout value of 2 hours. For typical storage devices, the issue could occur for tables that exceeded approximately 200 or 350 GB, depending on I/O speed. The fix relaxes the locking performed on the table being checked, which makes the timeout less likely. It also makes `InnoDB` recognize the syntax `CHECK TABLE QUICK`, which avoids the possibility of the timeout entirely. (Bug #11758510, Bug #50723)
- **Replication:** It was theoretically possible for concurrent execution of more than one instance of `SHOW BINLOG EVENTS` to crash the MySQL Server. (Bug #13979418)
- **Replication:** An event whose length exceeded the size of the master dump thread's `max_allowed_packet` caused replication to fail. This could occur when updating many large rows and using row-based replication.

As part of this fix, a new server option `--slave-max-allowed-packet` is added, which permits `max_allowed_packet` to be exceeded by the slave SQL and I/O threads. Now the size of a packet transmitted from the master to the slave is checked only against this value (available as the value of the `slave_max_allowed_packet` server system variable), and not against the value of `max_allowed_packet`. (Bug #12400221, Bug #60926)

- **Replication:** Statements using `AUTO_INCREMENT`, `LAST_INSERT_ID()`, `RAND()`, or user variables could be applied in the wrong context on the slave when using statement-based replication and replication filtering server options (see [How Servers Evaluate Replication Filtering Rules](#)). (Bug #11761686, Bug #54201)

References: See also: Bug #11754117, Bug #45670, Bug #11746146, Bug #23894.

- **Replication:** An `INSERT` into a table that has a composite primary key that includes an `AUTO_INCREMENT` column that is not the first column of this composite key is not safe for statement-based binary logging or replication. Such statements are now marked as unsafe and fail with an error when using the `STATEMENT` binary logging format. For more information, see [Determination of Safe and Unsafe Statements in Binary Logging](#), as well as [Replication and AUTO_INCREMENT](#).



Note

This issue does not affect tables using the `InnoDB` storage engine, since an `InnoDB` table with an `AUTO_INCREMENT` column requires at least one key where the auto-increment column is the only or leftmost column.

(Bug #11754117, Bug #45670)

References: See also: Bug #11761686, Bug #54201, Bug #11746146, Bug #23894.

- **Replication:** After upgrading a replication slave to MySQL 5.5.60 or later, enabling the query cache eventually caused the slave to fail. (Bug #64624, Bug #14005409)
- When the index enforcing a foreign key constraint was dropped while `foreign_key_checks=0`, further operations involving the foreign key column could cause a serious error after the `foreign_key_checks` option was re-enabled. (Bug #14025221)
- Incorrect stored program caching could cause statements within a stored program that included a `GROUP BY` clause to return different results across multiple program invocations. (Bug #13805127)
- For queries with `ORDER BY COUNT(*)` and `LIMIT`, the optimizer could choose an execution plan that produced incorrect results. (Bug #12713907)
- `SHOW TABLES` was very slow unless the required information was already in the disk cache. (Bug #60961, Bug #12427262)
- Sessions could end up deadlocked when executing a combination of `SELECT`, `DROP TABLE`, `KILL`, and `SHOW ENGINE INNODB STATUS`. (Bug #60682, Bug #12636001)
- `mysqlbinlog` exited with no error code if file write errors occurred. (Bug #55289, Bug #11762667)
- `yaSSL` rejected valid SSL certificates that `OpenSSL` accepts. (Bug #54348, Bug #11761822)
- When dumping the `mysql` database, `mysqldump` did not include the `general_log` and `slow_query_log` tables because they cannot be locked. This caused a problem after reloading the dump file if that file contained a `DROP DATABASE` statement for the `mysql` database: The database no longer contained the log tables and attempts to log to them failed. Now `mysqldump` includes statements to re-create the `general_log` and `slow_query_log` tables so that they exist after loading the dump file. Log table contents still are not dumped. (Bug #45740, Bug #11754178)

Changes in MySQL 5.1.63 (2012-05-07)

Bugs Fixed

- **Security Fix:** A security bug was fixed. (Bug #64884)
- **Security Fix:** A security bug was fixed. (Bug #59387)
- **Important Change; Partitioning:** The query cache did not always function correctly with partitioned tables in a transactional context. For this reason, the query cache is now disabled for any queries using partitioned tables, and such queries can no longer be cached. For more information, see [Restrictions and Limitations on Partitioning](#). (Bug #11761296, Bug #53775)
- **InnoDB:** Deleting a huge amount of data from `InnoDB` tables within a short time could cause the purge operation that removes delete-marked records to stall. This issue could result in unnecessary disk space use, but does not cause any problems with data integrity. If this issue causes a disk space shortage, restart the server to work around it. This issue is only likely to occur on 32-bit platforms. (Bug #13847885)
- **InnoDB:** The server could halt with an assertion error when the schema contained a large number of foreign key constraints, frequently being created and dropped, concurrent with other DML and DDL operations:

```
InnoDB: Assertion failure in thread thread_num in file data0data.ic line 344
InnoDB: Failing assertion: n < tuple->n_fields
```

(Bug #13635833)

- **InnoDB:** If the server crashed during a `TRUNCATE TABLE` or `CREATE INDEX` statement for an `InnoDB` table, or a `DROP DATABASE` statement for a database containing `InnoDB` tables, an index could be corrupted, causing an error message when accessing the table after restart:

```
InnoDB: Error: trying to load index index_name for table table_name
InnoDB: but the index tree has been freed!
```

In MySQL 5.1, this fix applies to the InnoDB Plugin, but not the built-in InnoDB storage engine. (Bug #12861864, Bug #11766019)

- **InnoDB:** When data was removed from an `InnoDB` table, newly inserted data might not reuse the freed disk blocks, leading to an unexpected size increase for the system tablespace or `.ibd` file (depending on the setting of `innodb_file_per_table`). The `OPTIMIZE TABLE` could compact a `.ibd` file in some cases but not others. The freed disk blocks would eventually be reused as additional data was inserted. (Bug #11766634, Bug #59783)
- **Replication:** The `--relay-log-space-limit` option was sometimes ignored.

More specifically, when the SQL thread went to sleep, it allowed the I/O thread to queue additional events in such a way that the relay log space limit was bypassed, and the number of events in the queue could grow well past the point where the relay logs needed to be rotated. Now in such cases, the SQL thread checks to see whether the I/O thread should rotate and provide the SQL thread a chance to purge the logs (thus freeing space).

Note that, when the SQL thread is in the middle of a transaction, it cannot purge the logs; it can only ask for more events until the transaction is complete. Once the transaction is finished, the SQL thread can immediately instruct the I/O thread to rotate. (Bug #12400313, Bug #64503)

References: See also: Bug #13806492.

- Mishandling of `NO_BACKSLASH_ESCAPES` SQL mode within stored procedures on slave servers could cause replication failures. (Bug #12601974)
- If the system time was adjusted backward during query execution, the apparent execution time could be negative. But in some cases these queries would be written to the slow query log, with the negative execution time written as a large unsigned number. Now statements with apparent negative execution time are not written to the slow query log. (Bug #63524, Bug #13454045)

References: See also: Bug #27208.

- `mysql_store_result()` and `mysql_use_result()` are not for use with prepared statements and are not intended to be called following `mysql_stmt_execute()`, but failed to return an error when invoked that way in `libmysqld`. (Bug #62136, Bug #13738989)

References: See also: Bug #47485.

- `SHOW` statements treated stored procedure, stored function, and event names as case sensitive. (Bug #56224, Bug #11763507)
- On Windows, `mysqlslap` crashed for attempts to connect using shared memory. (Bug #31173, Bug #11747181, Bug #59107, Bug #11766072)

Changes in MySQL 5.1.62 (2012-03-21)

- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

Functionality Added or Changed

- yaSSL was upgraded from version 1.7.2 to 2.2.0. (Bug #13706828)

References: See also: Bug #13713205.

- New `utf8_general_mysql500_ci` and `ucs2_general_mysql500_ci` collations have been added that preserve the behavior of `utf8_general_ci` and `ucs2_general_ci` from versions of MySQL previous to 5.1.24. Bug #27877 corrected an error in the original collations but introduced an incompatibility for columns that contain German 'ß' LATIN SMALL LETTER SHARP S. (As a result of the fix, that character compares equal to characters with which it previously compared different.) A symptom of the problem after upgrading to MySQL 5.1.24 or newer from a version older than 5.1.24 is that `CHECK TABLE` produces this error:

```
Table upgrade required.
Please do "REPAIR TABLE `t`" or dump/reload to fix it!
```

Unfortunately, `REPAIR TABLE` could not fix the problem. The new collations permit older tables created before MySQL 5.1.24 to be upgraded to current versions of MySQL.

To convert an affected table after a binary upgrade that leaves the table files in place, alter the table to use the new collation. Suppose that the table `t1` contains one or more problematic `utf8` columns. To convert the table at the table level, use a statement like this:

```
ALTER TABLE t1
CONVERT TO CHARACTER SET utf8 COLLATE utf8_general_mysql500_ci;
```

To apply the change on a column-specific basis, use a statement like this (be sure to repeat the column definition as originally specified except for the `COLLATE` clause):

```
ALTER TABLE t1
MODIFY c1 CHAR(N) CHARACTER SET utf8 COLLATE utf8_general_mysql500_ci;
```

To upgrade the table using a dump and reload procedure, dump the table using `mysqldump`, modify the `CREATE TABLE` statement in the dump file to use the new collation, and reload the table.

After making the appropriate changes, `CHECK TABLE` should report no error.

For more information, see [Checking Whether Tables or Indexes Must Be Rebuilt](#), and [Rebuilding or Repairing Tables or Indexes](#). (Bug #43593, Bug #11752408)

References: See also: Bug #27877.

Bugs Fixed

- **Security Fix:** A security bug was fixed. (Bug #63775)
- **Incompatible Change:** An earlier change (in MySQL 5.1.59 and 5.5.16) was found to modify date-handling behavior in General Availability-status series (MySQL 5.1 and 5.5). This change has been reverted.

The change was that several functions became more strict when passed a `DATE()` function value as their argument, thus they rejected incomplete dates with a day part of zero. These functions were affected: `CONVERT_TZ()`, `DATE_ADD()`, `DATE_SUB()`, `DAYOFYEAR()`, `LAST_DAY()`, `TIMESTAMPDIFF()`, `TO_DAYS()`, `TO_SECONDS()`, `WEEK()`, `WEEKDAY()`, `WEEKOFYEAR()`, `YEARWEEK()`. The previous behavior has been restored. (Bug #13458237)

- **Important Change; InnoDB:** When a row grew in size due to an `UPDATE` operation, other (non-updated) columns could be moved to off-page storage so that information about the row still fit within the constraints of the `InnoDB` page size. The pointer to the new allocated off-page data was not set up until the pages were allocated and written, potentially leading to lost data if the system crashed while the column was being moved out of the page. The problem was more common with tables using `ROW_FORMAT=DYNAMIC` or `ROW_FORMAT=COMPRESSED` along with the Barracuda file format, particularly with the `innodb_file_per_table` setting enabled, because page allocation operations are more common as the `.ibd` tablespace files are extended. Still, the problem could occur with any combination of InnoDB version, file format, and row format.

A related issue was that during such an `UPDATE` operation, or an `INSERT` operation that reused a delete-marked record, other transactions could see invalid data for the affected column, regardless of isolation level.

The fix corrects the order of operations for moving the column data off the original page and replacing it with a pointer. Now if a crash occurs at the precise moment when the column data is being transferred, the transfer will not be re-run during crash recovery.

In MySQL 5.1, this fix applies to the InnoDB Plugin, but not the built-in InnoDB storage engine. (Bug #13721257, Bug #12612184, Bug #12704861)

- **InnoDB:** An erroneous assertion could occur, in debug builds only, when creating an index on a column containing zero-length values (that is, `' '`). (Bug #13654923)
- **InnoDB:** A DDL operation such as `ALTER TABLE ... ADD COLUMN` could stall, eventually timing out with an `Error 1005: Can't create table` message referring to `fil_rename_tablespace`. (Bug #13636122, Bug #62100, Bug #63553)
- **InnoDB:** References to C preprocessor symbols and macros `HAVE_purify`, `UNIV_INIT_MEM_TO_ZERO`, and `UNIV_SET_MEM_TO_ZERO` were removed from the `InnoDB` source code. They were only used in debug builds instrumented for Valgrind. They are replaced by calls to the `UNIV_MEM_INVALID()` macro. (Bug #13418934)
- **InnoDB:** A DDL operation for an `InnoDB` table could cause a busy MySQL server to halt with an assertion error:

```
InnoDB: Failing assertion: trx->error_state == DB_SUCCESS
```

The error occurred if the DDL operation was run while all 1023 undo slots were in use by concurrent transactions. This error was less likely to occur in MySQL 5.5 and 5.6, because raising the number of `InnoDB` undo slots increased the number of simultaneous transactions (corresponding to the number of undo slots) from 1K to 128K. (Bug #12739098, Bug #62401)

- **InnoDB:** With 1024 concurrent `InnoDB` transactions running concurrently and the `innodb_file_per_table` setting enabled, a `CREATE TABLE` operation for an `InnoDB` table could fail. The `.ibd` file from the failed `CREATE TABLE` was left behind, preventing the table from being created later, after the load had dropped.

The fix adds error handling to delete the erroneous `.ibd` file. This error was less likely to occur in MySQL 5.5 and 5.6, because raising the number of `InnoDB` undo slots increased the number of simultaneous transactions needed to trigger the bug, from 1K to 128K. (Bug #12400341)

- **InnoDB:** When copying a partitioned `InnoDB` table from a Linux system to a Windows system, you could encounter this error:


```
101115 14:19:53 [ERROR] Table .\test\d has no primary key in InnoDB data dictionary, but has one in MySQL!
```

Normally, the solution to copy [InnoDB](#) tables from Linux to Windows is to create the tables on Linux with the [lower_case_table_names](#) option enabled. Partitioned tables, with [#P#](#) appended to the filename, were not covered by that solution. (Bug #11765438, Bug #58406)

- **InnoDB:** Server startup could produce an error for temporary tables using the [InnoDB](#) storage engine, if the path in the [\\$TMPDIR](#) variable ended with a `/` character. The error log would look like:

```
120202 19:21:26 InnoDB: Operating system error number 2 in a file operation.
InnoDB: The error means the system cannot find the path specified.
InnoDB: If you are installing InnoDB, remember that you must create
InnoDB: directories yourself, InnoDB does not create them.
120202 19:21:26 InnoDB: Error: trying to open a table, but could not
InnoDB: open the tablespace file './t/#sql7750_1_0.ibd'!
InnoDB: Have you moved InnoDB .ibd files around without using the
InnoDB: commands DISCARD TABLESPACE and IMPORT TABLESPACE?
InnoDB: It is also possible that this is a temporary table #sql...,
InnoDB: and MySQL removed the .ibd file for this.
```

The workaround for the problem was to create a similar temporary table again, copy its [.frm](#) file to [tmpdir](#) under the name mentioned in the error message (for example, [#sql123.frm](#)) and restart [mysqld](#) with [tmpdir](#) set to its normal value without a trailing slash, for example `/var/tmp`. On startup, MySQL would see the [.frm](#) file and issue `DROP TABLE` for the orphaned temporary table. (Bug #11754376, Bug #45976)

- **Replication:** It was possible on replication slaves where [FEDERATED](#) tables were in use to get timeouts on long-running operations, such as Error 1160 `Got an error writing communication packets`. The [FEDERATED](#) tables did not need to be replicated for the issue to occur. (Bug #11758931, Bug #51196)

References: See also: Bug #12896628, Bug #61790.

- [yaSSL](#) fixes previously applied to MySQL 5.5 were backported to 5.0 and 5.1. (Bug #13706621)
- A query that used an index on a [CHAR](#) column referenced in a [BETWEEN](#) clause could return invalid results. (Bug #13463488, Bug #63437)
- When the optimizer performed conversion of [DECIMAL](#) values while evaluating range conditions, it could produce incorrect results. (Bug #13453382)
- When used with the `--xml` option, `mysqldump --routines` failed to dump any stored routines, triggers, or events. (Bug #11760384, Bug #52792)
- If an attempt to initiate a statement failed, the issue could not be reported to the client because it was not prepared to receive any error messages prior to the execution of any statement. Since the user could not execute any queries, they were simply disconnected without providing a clear error.

After the fix for this issue, the client is prepared for an error as soon as it attempts to initiate a statement, so that the error can be reported prior to disconnecting the user. (Bug #11755281, Bug #47032)

- Using [myisamchk](#) with the sort recover method to repair a table having fixed-width row format could cause the row pointer size to be reduced, effectively resulting in a smaller maximum data file size. (Bug #48848, Bug #11756869)
- Under some circumstances, the result of [SUBSTRING_INDEX\(\)](#) incorrectly depended on the contents of the previous row. (Bug #42404, Bug #11751514)

- Due to improper locking, concurrent inserts into an [ARCHIVE](#) table at the same time as repair and check operations on the table resulted in table corruption. (Bug #37280, Bug #11748748)

Changes in MySQL 5.1.61 (2012-01-10)

Bugs Fixed

- **Important Change; InnoDB:** If an [ALTER TABLE](#) statement failed for an [InnoDB](#) table due to an error code from an underlying file-renaming system call, [InnoDB](#) could lose track of the [.ibd](#) file for the table. This issue only occurred when the [innodb_file_per_table](#) configuration option was enabled, and when the low-level error persisted through thousands of retry attempts. In MySQL 5.1, this issue applied to the InnoDB Plugin but not the built-in InnoDB storage engine.

For example, if you encounter an error like the following:

```
mysql> alter table sb2 add column d2 int;
ERROR 1025 (HY000): Error on rename of './sctest/#sql-1eb9_1' to
'./sctest/sb2' (errno: -1)
```

you might be able to access the [#sql*](#) table by copying a [.frm](#) file from a table with an identical schema. The table name to use for the [.frm](#) file would be ``sctest.#mysql50##sql-1eb9_1`` in the preceding example. (Bug #12884631, Bug #62146)

- **InnoDB:** Issuing [INSERT...ON DUPLICATE KEY](#) statements for [InnoDB](#) tables from concurrent threads could cause a [deadlock](#), particularly with the [INSERT...ON DUPLICATE KEY UPDATE](#) form. The problem could also be triggered by issuing multiple [INSERT IGNORE](#) statements. The fix avoids deadlocks caused by the same row being accessed by more than one transaction. Deadlocks could still occur when multiple rows are inserted and updated simultaneously by different transactions in inconsistent order; those types of deadlocks require the standard error handling on the application side, of re-trying the transaction. (Bug #11759688, Bug #52020, Bug #12842206)
- An incorrect [InnoDB](#) assertion could cause the server to halt. This issue only affected debug builds. The assertion referenced the source file [btr0pcur.ic](#) and the variable [cursor->pos_state](#). (Bug #13358468)
- Writes to [MyISAM](#) temporary tables could include uninitialized data, which could contain sensitive information. Now only bytes containing initialized data are copied, which also improves performance. (Bug #12997905)
- Passing a user variable as an argument to [GROUP_CONCAT\(\)](#) could cause a server exit if the variable value changed during query execution. (Bug #12408412)
- [LOAD INDEX INTO CACHE](#) could cause a server exit if the index cache was too small. (Bug #12361113)
- The [handle_segfault\(\)](#) signal-handler code in [mysqld](#) could itself crash due to calling unsafe functions. (Bug #54082, Bug #11761576)
- [ARCHIVE](#) tables with [NULL](#) columns could cause server crashes or become corrupt under concurrent load. (Bug #51252, Bug #11758979)
- Enabling [myisam_use_mmap](#) could cause the server to crash. (Bug #48726, Bug #11756764)
- Concurrent access to [ARCHIVE](#) tables could cause corruption. (Bug #42784, Bug #11751793)

Changes in MySQL 5.1.60 (2011-11-16)

- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

Functionality Added or Changed

- Upgrading from an [Advanced GPL](#) RPM package to an [Advanced](#) RPM package did not work. Now on Linux it is possible to use `rpm -U` to replace any installed MySQL product by any other of the same release family. It is not necessary to remove the old produce with `rpm -e` first. (Bug #11886309)
- [MEMORY](#) table creation time is now available in the `CREATE_TIME` column of the `INFORMATION_SCHEMA.TABLES` table and the `Create_time` column of `SHOW TABLE STATUS` output. (Bug #51655, Bug #11759349)

Bugs Fixed

- **Performance; InnoDB:** The process of deallocating the [InnoDB Adaptive Hash Index](#) was made faster, during shutdown or when turning off the AHI with the statement:

```
SET GLOBAL innodb_adaptive_hash_index=OFF;
```

(Bug #13006367, Bug #62487)

- **Performance; InnoDB:** This fix improves the performance of instrumentation code for [InnoDB](#) buffer pool operations. (Bug #12950803, Bug #62294)
- **Performance; InnoDB:** The [InnoDB buffer pool](#) management code was optimized for handling pages from [compressed](#) tables. This fixes a slowdown that could occur particularly during the [warmup](#) period for the buffer pool. (Bug #12610930, Bug #61341)
- **InnoDB:** Fixed a compilation problem that affected the [InnoDB](#) source code with `gcc` 4.6.1. The affected [InnoDB](#) source file was `btr/btr0cur.c`. (Bug #13116045)
- **InnoDB:** Lookups using secondary indexes could give incorrect matches under a specific set of conditions. The conditions involve an index defined on a column prefix, for a BLOB or other long column stored outside the index page, with a table using the Barracuda file format. (Bug #12601439, Bug #12543666)
- **InnoDB:** An `UPDATE` statement for an [InnoDB](#) table could hang. The issue affects tables using the [Barracuda](#) file format and having multiple indexes on [column prefixes](#). The size of an [undo log](#) record could exceed the [page size](#), even though the total size of the column prefixes was less than the page size (usually 16KB). In MySQL 5.5 and higher, this error is now reported using the new code `ER_UNDO_RECORD_TOO_BIG`. In MySQL 5.1 with the InnoDB Plugin, this error is reported using the existing code `ER_TOO_BIG_ROW_SIZE`. (Bug #12547647)
- **InnoDB:** This fix corrects cases where the MySQL server could hang or abort with a [long semaphore wait](#) message. (This is a different issue than when these symptoms occurred during a `CHECK TABLE` statement.) (Bug #11766591, Bug #59733)
- **InnoDB:** A timing issue could cause a crash while processing the `SHOW ENGINE INNODB STATUS` command. This issue only affected debug builds of the server. (Bug #11766546, Bug #59682)
- **Replication:** Issuing the following statements, in the order shown, could cause a deadlock between the user thread and I/O thread:

```
START SLAVE;
```

```
STOP SLAVE SQL_THREAD;  
START SLAVE;
```

(Bug #11878104)

References: See also: Bug #44312, Bug #11752963, Bug #38715, Bug #38716.

- Internal conversion of zero to binary and back could yield a result with incorrect precision. (Bug #12911710)
- Valgrind warnings generated by `filesort` operations were fixed. (Bug #12856915)
- Several improvements were made to the `libedit` library bundled with MySQL distributions, and that is available for all platforms that MySQL supports except Windows.
 - Navigation keys did not work for UTF-8 input.
 - Word navigation and delete operations did not work for UTF-8 input with Cyrillic characters.
 - Nonlatin characters were corrupted in overwrite mode for UTF-8 input.
 - Long queries caused the statement history file to become corrupted.
 - The Alt key caused history operations to fail.

(Bug #12605400, Bug #12613725, Bug #12618092, Bug #12624155, Bug #12617651, Bug #12605388)

- `decimal_round()` could cause a server exit when processing long numeric strings. (Bug #12563865)
- With Valgrind enabled, `InnoDB` semaphore wait timeouts were too low and could expire. (Bug #11765460)
- The help message for `mysql_install_db` did not indicate that it supports the `--defaults-file`, `--defaults-extra-file` and `--no-defaults` options. (Bug #58898, Bug #11765888)
- An assertion designed to detect zero-length sort keys also was raised when the entire key set fit in memory. (Bug #58200, Bug #11765254)
- `myisampack` could create corrupt `FULLTEXT` indexes when compressing tables. (Bug #53646, Bug #11761180)
- `OPTIMIZE TABLE` could corrupt `MyISAM` tables if `myisam_use_mmap` was enabled. (Bug #49030, Bug #11757032)
- If MySQL was configured with `--without-plugin-innobase` and `--with-plugin-innodb_plugin`, to suppress building the built-in `InnoDB` storage engine and build the `InnoDB Plugin` instead, the `innochecksum` utility was not built. (Bug #47337, Bug #11755544)
- A linking problem prevented the `FEDERATED` storage engine plugin from loading. (Bug #40942, Bug #11750417)
- On Fedora, certain accesses to `/var/lib/mysql/HOSTNAME.err` were blocked by SELinux policy, which made the server fail at startup with the message: `Manager of pid-file quit without updating file` (Bug #37165, Bug #12927740)
- For `FEDERATED` tables, loss of connection to the remote table during some insert operations could cause a server crash. (Bug #34660, Bug #11747970)

Changes in MySQL 5.1.59 (2011-09-15)

- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

Functionality Added or Changed

- If the `--with-ndbcluster` option is given to the `configure` script, it now produces a warning that the version of MySQL Cluster included in 5.1 is no longer maintained. (The separate MySQL Cluster distribution should be used instead.) (Bug #49093, Bug #11757091)

Bugs Fixed

- **Performance; InnoDB:** This fix improves the performance of operations on `VARCHAR(N)` columns in InnoDB tables, where `N` is declared as a large value but the actual string values in the table are short. (Bug #12835650)
- **Performance; InnoDB:** The `DROP TABLE` command for an InnoDB table could be very slow, in a configuration with a combination of table compression, partitioning, and a large buffer pool. (Bug #12635227, Bug #61188)
- **Performance; InnoDB:** The “random `read-ahead`” feature that was removed from the InnoDB Plugin is now available again. Because it is only helpful for certain workloads, it is turned off by default. To turn it on, enable the `innodb_random_read_ahead` configuration option. Because this feature can improve performance in some cases and reduce performance in others, before relying on this setting, benchmark both with and without the setting enabled. (Bug #12356373)

- **Incompatible Change:** Handling of a date-related assertion was modified.

However, a consequence of this change is that several functions become more strict when passed a `DATE()` function value as their argument and reject incomplete dates with a day part of zero. These functions are affected: `CONVERT_TZ()`, `DATE_ADD()`, `DATE_SUB()`, `DAYOFYEAR()`, `LAST_DAY()`, `TIMESTAMPDIFF()`, `TO_DAYS()`, `TO_SECONDS()`, `WEEK()`, `WEEKDAY()`, `WEEKOFYEAR()`, `YEARWEEK()`. Because this changes date-handling behavior in General Availability-status series (MySQL 5.1 and 5.5), it was reverted in 5.1.62 and 5.5.21. The change is retained in MySQL 5.6.

References: See also: Bug #13458237.

- **InnoDB:** The `DATA_LENGTH` column in the `INFORMATION_SCHEMA.TABLES` table now correctly reports the on-disk sizes of tablespaces for InnoDB compressed tables. (Bug #12770537)
- **InnoDB:** With the configuration settings `innodb_file_per_table=1` and `innodb_file_format=Barracuda`, inserting a column value greater than half the page size, and including that column in a secondary index, could cause a crash when that column value was updated. (Bug #12637786)
- **InnoDB:** Unused functions were removed from the internal InnoDB code related to mini-transactions, to clarify the logic. (Bug #12626794, Bug #61240)
- **InnoDB:** A `DROP TABLE` or `DROP INDEX` statement for an InnoDB table on a busy server could cause a crash or corrupt data in the `buffer pool`, if the buffer pool contained data from an InnoDB compressed table that was being accessed at the same time. (The crash could occur whether or not the table being dropped used compression.) (Bug #11765566, Bug #58549)
- **Partitioning:** Auto-increment columns of partitioned tables were checked even when they were not being written to. In debug builds, this could lead to a server crash. (Bug #11765667, Bug #58655)

- **Replication:** Processing of corrupted table map events could cause the server to crash. This was especially likely if the events mapped different tables to the same identifier, such as could happen due to Bug #56226.

Now, before applying a table map event, the server checks whether the table has already been mapped with different settings, and if so, an error is raised and the slave SQL thread stops. If it has been mapped with the same settings, or if the table is set to be ignored by filtering rules, there is no change in behavior: the event is skipped and IDs are not checked. (Bug #44360, Bug #11753004)

References: See also: Bug #56226, Bug #11763509.

- The option-parsing code for empty strings leaked memory. (Bug #12589928)
- `ALTER TABLE {MODIFY|CHANGE} ... FIRST` did nothing except rename columns if the old and new versions of the table had exactly the same structure with respect to column data types. As a result, the mapping of column name to column data was incorrect. The same thing happened for `ALTER TABLE DROP COLUMN ... ADD COLUMN` statements intended to produce a new version of the table with exactly the same structure as the old version. (Bug #61493, Bug #12652385)
- For a `lower_case_table_names` value of 1 or 2 and a database having a mixed-case name, calling a stored function using a fully qualified name including the database name failed. (Bug #60347, Bug #11840395)
- `SELECT DISTINCT` with a deterministic stored function in the `WHERE` clause could produce incorrect results. (Bug #59736, Bug #11766594)
- For `MyISAM` tables, attempts to insert incorrect data into an indexed `GEOMETRY` column could result in table corruption. (Bug #57323, Bug #11764487)
- `CREATE TABLE` without an `ENGINE` option determined the default engine at parse rather than execution time. This led to incorrect results if the statement was executed within a stored program and the default engine had been changed in the meantime. (Bug #50614, Bug #11758414)
- A race condition between loading a stored routine using the name qualified by the database name and dropping that database resulted in a spurious error message: `The table mysql.proc is missing, corrupt, or contains bad data` (Bug #47870, Bug #11756013)
- Upgrades using an RPM package recreated the `test` database, which is undesirable when the DBA had removed it. (Bug #45415, Bug #11753896)

Changes in MySQL 5.1.58 (2011-07-05)

Bugs Fixed

- **InnoDB:** If the `column prefix` in an index key came from an `off-page column`, a serious error could occur during a background operation (`purge` of delete-marked records). (Bug #12429576)
- **InnoDB:** If the server crashed while an XA transaction was prepared but not yet committed, the transaction could remain in the system after restart, and cause a subsequent shutdown to hang. (Bug #11766513, Bug #59641)
- **Partitioning:** When executing a row-ordered retrieval index merge, the partitioning handler used memory from that allocated for the table, rather than that allocated to the query, causing table object memory not to be freed until the table was closed. (Bug #11766249, Bug #59316)
- **Replication:** When `mysqlbinlog` was invoked using `--base64-output=decode-row` and `--start-position=pos`, (where `pos` is a point in the binary log past the format description log event), a spurious error of the type shown here was generated:


```
malformed binlog: it does not contain any Format_description_log_event...
```

However, since there is nothing unsafe about not printing the format description log event, the error has been removed for this case. (Bug #12354268)

- **Replication:** Typographical errors appeared in the text of several replication error messages. (The word “position” was misspelled as “postion”.) (Bug #11762616, Bug #55229)
- Large values passed to `FORMAT()` caused a buffer overflow and a server exit. (Bug #12406055)
- After the fix for Bug #11889186, `MAKEDATE()` arguments with a year part greater than 9999 raised an assertion. (Bug #12403504)

References: This issue is a regression of: Bug #11889186.

- An assertion could be raised due to a missing `NULL` value check in `Item_func_round::fix_length_and_dec()`. (Bug #12392636)
- In debug builds on Solaris, an assertion was raised if a reverse IP lookup with `gethostbyaddr_r()` failed. (Bug #12377872)
- MySQL did not build if configured with both `--with-debug` and `--with-libedit`. (Bug #12329909)
- A problem introduced in MySQL 5.1.57 caused very old (MySQL 4.0) clients to be unable to connect to the server. (Bug #61222, Bug #12563279)
- Using `CREATE EVENT IF NOT EXISTS` for an event that already existed and was enabled caused multiple instances of the event to run. (Bug #61005, Bug #12546938)
- An incorrect `max_length` value for `YEAR` values could be used in temporary result tables for `UNION`, leading to incorrect results. (Bug #59343, Bug #11766270)
- In `Item_func_in::fix_length_and_dec()`, a Valgrind warning for uninitialized values was corrected. (Bug #59270, Bug #11766212)
- In `ROUND()` calculations, a Valgrind warning for uninitialized memory was corrected. (Bug #58937, Bug #11765923)

References: This issue is a regression of: Bug #33143.

- Valgrind warnings caused by comparing index values to an uninitialized field were corrected. (Bug #58705, Bug #11765713)
- `LOAD DATA INFILE` errors could leak I/O cache memory. (Bug #58072, Bug #11765141)
- For `LOAD DATA INFILE`, multibyte character sequences could be pushed onto a stack too small to accommodate them. (Bug #58069, Bug #11765139)
- An embedded client aborted rather than issuing an error message if it issued a `TEE` command (`\T file_name`) and the directory containing the file did not exist. This occurred because the wrong error handler was called. (Bug #57491, Bug #11764633)
- `ALTER EVENT` could change the event status. (Bug #57156, Bug #11764334)
- Compilation failed on AIX due to a missing `bzero()` prototype. (Bug #55721, Bug #11763056)
- In debug builds, `Field_new_decimal::store_value()` was subject to buffer overflows. (Bug #55436, Bug #11762799)

- On Linux, the `mysql` client built using the bundled `libedit` did not read `~/.editrc`. (Bug #49967, Bug #11757855)
- The optimizer sometimes incorrectly processed `HAVING` clauses for queries that did not also have an `ORDER BY` clause. (Bug #48916, Bug #11756928)
- `PROCEDURE ANALYSE()` could leak memory for `NULL` results, and could return incorrect results if used with a `LIMIT` clause. (Bug #48137, Bug #11756242)
- On some platforms, the `Incorrect value: xxx for column yyy at row zzz` error produced by `LOAD DATA INFILE` could have an incorrect value of `zzz`. (Bug #46895, Bug #11755168)
- In MySQL 5.1 and up, if a table had triggers that used syntax supported in 5.0 but not 5.1, the table became unavailable. Now the table is marked as having broken triggers. These should be dropped and recreated manually. (Bug #45235, Bug #11753738)
- An attempt to install nonexistent files during installation was corrected. (Bug #43247, Bug #11752142)
- On FreeBSD 64-bit builds of the embedded server, exceptions were not prevented from propagating into the embedded application. (Bug #38965, Bug #11749418)

Changes in MySQL 5.1.57 (2011-05-05)

- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

Functionality Added or Changed

- When invoked with the `--auto-generate-sql` option, `mysqlslap` dropped the schema specified with the `--create-schema` option at the end of the test run, which may have been unexpected by the user. `mysqlslap` now has a `--no-drop` option that prevents any schema created during the test run from being dropped. (Bug #58090, Bug #11765157)
- A new system variable, `max_long_data_size`, now controls the maximum size of parameter values that can be sent with the `mysql_stmt_send_long_data()` C API function. If not set at server startup, the default is the value of the `max_allowed_packet` system variable. This variable is deprecated. In MySQL 5.6, it is removed and the maximum parameter size is controlled by `max_allowed_packet`.

Bugs Fixed

- **InnoDB; Replication:** Trying to update a column, previously set to `NULL`, of an `InnoDB` table with no primary key caused replication to fail on the slave with `Can't find record in 'table'`.



Note

This issue was inadvertently reintroduced in MySQL 5.6.6, and fixed again in MySQL 5.6.12.

(Bug #11766865, Bug #60091)

References: See also: Bug #16566658.

- **InnoDB:** The server could halt if `InnoDB` interpreted a very heavy I/O load for 15 minutes or more as an indication that the server was hung. This change fixes the logic that measures how long `InnoDB` threads were waiting, which formerly could produce false positives. (Bug #11877216, Bug #11755413, Bug #47183)

- **InnoDB:** The MySQL server could hang during `CREATE TABLE`, `OPTIMIZE TABLE`, or `ALTER TABLE` or other DDL operation that performs a table copy for an InnoDB table, if such operations were performed by multiple sessions simultaneously. The error was reported as:

```
InnoDB: Error: semaphore wait has lasted > 600 seconds
```

(Bug #11760042, Bug #52409)

- **Partitioning:** A problem with a previous fix for poor performance of `INSERT ON DUPLICATE KEY UPDATE` statements on tables having many partitions caused the handler function for reading a row from a specific index to fail to store the ID of the partition last used. This caused some statements to fail with `Can't find record` errors. (Bug #59297, Bug #11766232)

References: This issue is a regression of: Bug #52455.

- **Replication:** Using the `--server-id` option with `mysqlbinlog` could cause format description log events to be filtered from the binary log, leaving `mysqlbinlog` unable to read the remainder of the log. Now such events are always read without regard to the value of this option.

As part of the fix for this problem, `mysqlbinlog` now also reads rotate log events without regard to the value of `--server-id`. (Bug #59530, Bug #11766427)

- InnoDB invoked some `zlib` functions without proper initialization. (Bug #11849231)
- Two unused test files in `storage/ndb/test/sql` contained incorrect versions of the GNU Lesser General Public License. The files and the directory containing them have been removed. (Bug #11810224)

References: See also: Bug #11810156.

- Selecting from a view for which the definition included a `HAVING` clause failed with an error:

```
1356: View '...' references invalid table(s) or column(s)
or function(s) or definer/invoke of view lack rights to use them
```

(Bug #60295, Bug #11829681)

- The server permitted `max_allowed_packet` to be set lower than `net_buffer_length`, which does not make sense because `max_allowed_packet` is the upper limit on `net_buffer_length` values. Now a warning occurs and the value remains unchanged. (Bug #59959, Bug #11766769)
- The server read one byte too many when trying to process an XML string lacking a closing single quote (') or double quote (") character used as an argument for `UpdateXML()` or `ExtractValue()`. (Bug #59901, Bug #11766725)

References: See also: Bug #44332, Bug #11752979.

- Attempting to create a spatial index on a `CHAR` column longer than 31 bytes led to an assertion failure if the server was compiled with `safemutex` support. (Bug #59888, Bug #11766714)
- Aggregation followed by a subquery could produce an incorrect result. (Bug #59839, Bug #11766675)
- An incorrect character set pointer passed to `my_strtoll10_mb2()` caused an assertion to be raised. (Bug #59648, Bug #11766519)
- A missing variable initialization for `Item_func_set_user_var` objects could raise an assertion. (Bug #59527, Bug #11766424)

- `mysqldump` did not quote database names in `ALTER DATABASE` statements in its output, which could cause an error at reload time for database names containing a dash. (Bug #59398, Bug #11766310)
- In `Item_func_month::val_str()`, a Valgrind warning for a too-late `NULL` value check was corrected. (Bug #59166, Bug #11766126)
- In `Item::get_date`, a Valgrind warning for a missing `NULL` value check was corrected. (Bug #59164, Bug #11766124)
- In `extract_date_time()`, a Valgrind warning for a missing end-of-string check was corrected. (Bug #59151, Bug #11766112)
- In string context, the `MIN()` and `MAX()` functions did not take into account the unsignedness of a `BIGINT UNSIGNED` argument. (Bug #59132, Bug #11766094)
- In `Item_func::val_decimal`, a Valgrind warning for a missing `NULL` value check was corrected. (Bug #59125, Bug #11766087)
- In `Item_func_str_to_date::val_str`, a Valgrind warning for an uninitialized variable was corrected. (Bug #58154, Bug #11765216)
- The code for `PROCEDURE ANALYSE()` had a missing `DEBUG_RETURN` statement, which could cause a server crash in debug builds. (Bug #58140, Bug #11765202)
- An assertion could be raised in `Item_func_int_val::fix_num_length_and_dec()` due to overflow for geometry functions. (Bug #57900, Bug #11764994)
- An assertion could be raised if a statement that required a name lock on a table (for example, `DROP TRIGGER`) executed concurrently with an `INFORMATION_SCHEMA` query that also used the table. (Bug #56541, Bug #11763784)
- For a client connected using SSL, the `Ssl_cipher_list` status variable was empty and did not show the possible cipher types. (Bug #52596, Bug #11760210)
- With `lower_case_table_names=2`, resolution of objects qualified by database names could fail. (Bug #50924, Bug #11758687)
- A potential invalid memory access discovered by Valgrind was fixed. (Bug #48053, Bug #11756169)
- Bitmap functions used in one thread could change bitmaps used by other threads, raising an assertion. (Bug #43152, Bug #11752069)
- `SHOW EVENTS` did not always show events from the correct database. (Bug #41907, Bug #11751148)

Changes in MySQL 5.1.56 (2011-03-01)

- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

Functionality Added or Changed

- `mysqldump --xml` now displays comments from column definitions. (Bug #13618, Bug #11745324)

Bugs Fixed

- **Security Fix:** A security bug was fixed. (Bug #36544)
- **InnoDB:** InnoDB returned values for “rows examined” in the query plan that were higher than expected. `NULL` values were treated in an inconsistent way. The inaccurate statistics could trigger “false positives”

in combination with the `max_join_size` setting, because the queries did not really examine as many rows as reported.

A new configuration option `innodb_stats_method` lets you specify how `NULL` values are treated when calculating index statistics. Allowed values are `nulls_equal` (the default), `nulls_unequal` and `null_ignored`. The meanings of these values are similar to those of the `myisam_stats_method` option. (Bug #30423)

- **Partitioning:** Trying to use the same column more than once in the partitioning key when partitioning a table by `KEY` caused `mysqld` to crash. Such duplication of key columns is now expressly disallowed, and fails with an appropriate error. (Bug #53354, Bug #57924)
- **Replication:** When using the statement-based logging format, `INSERT ON DUPLICATE KEY UPDATE` and `INSERT IGNORE` statements affecting transactional tables that did not fail were not written to the binary log if they did not insert any rows. (With statement-based logging, all successful statements should be logged, whether they do or do not cause any rows to be changed.) (Bug #59338, Bug #11766266)
- **Replication:** Formerly, `STOP SLAVE` stopped the slave I/O thread first and then stopped the slave SQL thread; thus, it was possible for the I/O thread to stop after replicating only part of a transaction which the SQL thread was executing, in which case—if the transaction could not be rolled back safely—the SQL thread could hang.

Now, `STOP SLAVE` stops the slave SQL thread first and then stops the I/O thread; this guarantees that the I/O thread can fetch any remaining events in the transaction that the SQL thread is executing, so that the SQL thread can finish the transaction if it cannot be rolled back safely. (Bug #58546, Bug #11765563)

- A query of the following form returned an incorrect result, where the values for `col_name` in the result set were entirely replaced with `NULL` values:

```
SELECT DISTINCT col_name ... ORDER BY col_name DESC;
```

(Bug #59308, Bug #11766241)

- `DELETE` or `UPDATE` statements could fail if they used `DATE` or `DATETIME` values with a year, month, or day part of zero. (Bug #59173)
- The `ESCAPE` clause for the `LIKE` operator permits only expressions that evaluate to a constant at execution time, but aggregate functions were not being rejected. (Bug #59149, Bug #11766110)
- Memory leaks detected by Valgrind, some of which could cause incorrect query results, were corrected. (Bug #59110, Bug #11766075)
- `mysqlslap` failed to check for a `NULL` return from `mysql_store_result()` and crashed trying to process the result set. (Bug #59109, Bug #11766074)
- In debug builds, `SUBSTRING_INDEX(FORMAT(...), FORMAT(...))` could cause a server crash. (Bug #58371, Bug #11765406)
- When `mysqladmin` was run with the `--sleep` and `--count` options, it went into an infinite loop executing the specified command. (Bug #58221, Bug #11765270)
- Some string-manipulating SQL functions use a shared string object intended to contain an immutable empty string. This object was used by the SQL function `SUBSTRING_INDEX()` to return an empty string when one argument was of the wrong data type. If the string object was then modified by the SQL function `INSERT()`, undefined behavior ensued. (Bug #58165, Bug #11765225)

- Parsing nested regular expressions could lead to recursion resulting in a stack overflow crash. (Bug #58026, Bug #11765099)
- The `mysql` client went into an infinite loop if the standard input was a directory. (Bug #57450, Bug #11764598)
- The expression `const1 BETWEEN const2 AND field` was optimized incorrectly and produced incorrect results. (Bug #57030, Bug #11764215)
- Some RPM installation scripts used a hardcoded value for the data directory, which could result in a failed installation for users who have a nonstandard data directory location. The same was true for other configuration values such as the PID file name. (Bug #56581, Bug #11763817)
- On FreeBSD and OpenBSD, the server incorrectly checked the range of the system date, causing legal values to be rejected. (Bug #55755, Bug #11763089)
- When using `ExtractValue()` or `UpdateXML()`, if the XML to be read contained an incomplete XML comment, MySQL read beyond the end of the XML string when processing, leading to a crash of the server. (Bug #44332, Bug #11752979)

Changes in MySQL 5.1.55 (2011-02-07)

- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

Functionality Added or Changed

- The time zone tables available at <http://dev.mysql.com/downloads/timezones.html> have been updated. These tables can be used on systems such as Windows or HP-UX that do not include zoneinfo files. (Bug #40230)

Bugs Fixed

- **Security Fix:** A security bug was fixed. (Bug #57952)
- **Performance; InnoDB:** An `UPDATE` statement for an `InnoDB` table could be slower than necessary if it changed a column covered by a prefix index, but did not change the prefix portion of the value. The fix improves performance for InnoDB 1.1 in MySQL 5.5 and higher, and the InnoDB Plugin for MySQL 5.1. (Bug #58912, Bug #11765900)
- **Performance:** Queries involving `InnoDB` tables in the `INFORMATION_SCHEMA` tables `TABLE_CONSTRAINTS`, `KEY_COLUMN_USAGE`, or `REFERENTIAL_CONSTRAINTS` were slower than necessary because statistics were rechecked more often than required, even more so when many foreign keys were present. The improvement to this may be of particular benefit to users of MySQL Enterprise Monitor with many monitored servers or tens of thousands of tables. (Bug #43818, Bug #11752585)
- **Incompatible Change:** When `auto_increment_increment` is greater than one, values generated by a bulk insert that reaches the maximum column value could wrap around rather producing an overflow error.

As a consequence of the fix, it is no longer possible for an auto-generated value to be equal to the maximum `BIGINT UNSIGNED` value. It is still possible to store that value manually, if the column can accept it. (Bug #39828, Bug #11749800)

- **Important Change; Partitioning:** Date and time functions used as partitioning functions now have the types of their operands checked; use of a value of the wrong type is now disallowed in such cases. In

addition, `EXTRACT(WEEK FROM col_name)`, where `col_name` is a `DATE` or `DATETIME` column, is now disallowed altogether because its return value depends on the value of the `default_week_format` system variable. (Bug #54483, Bug #11761948)

References: See also: Bug #57071, Bug #11764255.

- **InnoDB:** A compilation problem affected the `InnoDB` source code on NetBSD/sparc64. (Bug #59327)

References: See also: Bug #53916.

- **InnoDB:** The presence of a double quotation mark inside the `COMMENT` field for a column could prevent a foreign key constraint from being created properly. (Bug #59197, Bug #11766154)
- **InnoDB:** In `InnoDB` status output, the value for `I/O sum[]` could be incorrect, displayed as a very large number. (Bug #57600)
- **InnoDB:** The server could crash with an assertion error, if a stored procedure, stored function, or trigger modified one `InnoDB` table containing an `auto-increment` column, and dropped another `InnoDB` table containing an `auto-increment` column. (Bug #56228)
- **InnoDB:** It was not possible to query the `information_schema.INNODB_TRX` table while other connections were running queries involving `BLOB` types. (Bug #55397, Bug #11762763)
- **InnoDB:** The `OPTIMIZE TABLE` statement reset the `auto-increment` counter for an `InnoDB` table. Now the `auto-increment` value is preserved across this operation. (Bug #18274)
- **Partitioning:** Failed `ALTER TABLE ... PARTITION` statements could cause memory leaks. (Bug #56380, Bug #11763641)

References: See also: Bug #46949, Bug #11755209, Bug #56996, Bug #11764187.

- **Replication:** When closing a session that used temporary tables, binary logging could sometimes fail with a spurious `Failed to write the DROP statement for temporary tables to binary log`. (Bug #57288)
- **Replication:** By default, a value is generated for an `AUTO_INCREMENT` column by inserting either `NULL` or 0 into the column. Setting the `NO_AUTO_VALUE_ON_ZERO` server SQL mode suppresses this behavior for 0, so that it occurs only when `NULL` is inserted into the column.

This behavior is also followed on a replication slave (by the slave SQL thread) when applying events that have been logged on the master using the statement-based format. However, when applying events that had been logged using the row-based format, `NO_AUTO_VALUE_ON_ZERO` was ignored, which could lead to an assertion.

To fix this issue, the value of an `AUTO_INCREMENT` column is no longer generated when applying an event that was logged using the row-based row format, as this value is already contained in the changes applied on the slave. (Bug #56662)

- **Replication:** The `BINLOG` statement modified the values of session variables, which could lead to problems with operations such as point-in-time recovery. One such case occurred when replaying a row-based binary log which relied on setting `foreign_key_checks = OFF` at the session level to create and populate a set of `InnoDB` tables having foreign key constraints. (Bug #54903)
- **Replication:** `mysqlbinlog` printed `USE` statements to its output only when the default database changed between events. To illustrate how this could cause problems, suppose that a user issued the following sequence of statements:

```
CREATE DATABASE mydb;
```



```
USE mydb;
CREATE TABLE mytable (column_definitions);
DROP DATABASE mydb;
CREATE DATABASE mydb;
USE mydb;
CREATE TABLE mytable (column_definitions);
```

When played back using `mysqlbinlog`, the second `CREATE TABLE` statement failed with `Error: No Database Selected` because the second `USE` statement was not played back, due to the fact that a database other than `mydb` was never selected.

This fix ensures that `mysqlbinlog` outputs a `USE` statement whenever it reads one from the binary log. (Bug #50914, Bug #11758677)

- **Replication:** Previously, when a statement failed with a different error on the slave than on the master, the slave SQL thread displayed a message containing:
 - The error message for the master error code
 - The master error code
 - The error message for the slaves error code
 - The slave error code

However, the slave has no information with which to fill in any print format specifiers for the master message, so it actually displayed the message format string. To make it clearer that the slave is not displaying the actual message as it appears on the master, the slave now indicates that the master part of the output is the message format, not the actual message. For example, previously the slave displayed information like this:

```
Error: "Query caused different errors on master and slave. Error
on master: 'Duplicate entry '%-.192s' for key %d' (1062), Error on
slave: 'no error' (0). Default database: 'test'. Query: 'insert
into t1 values(1),(2)'" (expected different error codes on master
and slave)
```

Now the slave displays this:

```
Error: "Query caused different errors on master and slave. Error
on master: message format='Duplicate entry '%-.192s' for key %d'
error code=1062 ; Error on slave: actual message='no error', error
code=0. Default database: 'test'. Query: 'insert into t1 values(1),(2)'"
(expected different error codes on master and slave)
```

(Bug #46697)

- **Replication:** When an error occurred in the generation of the name for a new binary log file, the error was logged but not shown to the user. (Bug #46166)

References: See also: Bug #37148, Bug #11748696, Bug #40611, Bug #11750196, Bug #43929, Bug #51019.

- `MIN(year_col)` could return an incorrect result in some cases. (Bug #59211, Bug #11766165)
- If `max_allowed_packet` was set larger than 16MB, the server failed to reject too-large packets with "Packet too large" errors. (Bug #58887, Bug #11765878)

- Issuing `EXPLAIN EXTENDED` for a query that would use condition pushdown could cause `mysqld` to crash. (Bug #58553, Bug #11765570)
- `EXPLAIN` could crash for queries that used `GROUP_CONCAT()`. (Bug #58396)
- Configuration with maintainer mode enabled resulted in errors when compiling with `icc`. (Bug #57991, Bug #58871)
- Unnecessary subquery evaluation in contexts such as statement preparation or view creation could cause a server crash. (Bug #57703)
- View creation could produce Valgrind warnings. (Bug #57352)
- `NULL` geometry values could cause a crash in `Item_func_spatial_collection::fix_length_and_dec`. (Bug #57321)
- The `cp1251` character set did not properly support the Euro sign (`0x88`). For example, converting a string containing this character to `utf8` resulted in `'?'` rather than the `utf8` Euro sign. (Bug #56639)
- Some unsigned system variables could be displayed with negative values. (Bug #55794)
- `CREATE DATABASE` and `DROP DATABASE` caused `mysql --one-database` to lose track of the statement-filtering context. (Bug #54899)
- An assertion could be raised during concurrent execution of `DROP DATABASE` and `REPAIR TABLE` if the drop deleted a table's `.TMD` file at the same time the repair tried to read details from the old file that was just removed.

A problem could also occur when `DROP TABLE` tried to remove all files belonging to a table at the same time `REPAIR TABLE` had just deleted the table's `.TMD` file. (Bug #54486)

- When `mysqld` printed crash dump information, it incorrectly indicated that some valid pointers were invalid. (Bug #51817)
- On FreeBSD, if `mysqld` was killed with a `SIGHUP` signal, it could corrupt InnoDB `.ibd` files. (Bug #51023, Bug #11758773)
- An assertion could be raised if `-1` was inserted into an `AUTO_INCREMENT` column by a statement writing more than one row. (Bug #50619, Bug #11758417)
- If a client supplied a user name longer than the maximum 16 characters permitted for names stored in the MySQL grant tables, all characters were being considered significant when checking for a match. Historically, only the first 16 characters were used for matching; this behavior was restored. (Bug #49752)
- The `my_seek()` and `my_tell()` functions ignored the `MY_WME` flag when they returned an error, which could cause client programs to hang. (Bug #48451)
- During assignment of values to system variables, legality checks on the value range occurred too late, preventing proper error checking. (Bug #43233)
- On Solaris, time-related functions such as `NOW()` or `SYSDATE()` could return a constant value. (Bug #42054)
- If the remote server for a `FEDERATED` table could not be accessed, queries for the `INFORMATION_SCHEMA.TABLES` table failed. (Bug #35333)

Changes in MySQL 5.1.54 (2010-11-26)

- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

Functionality Added or Changed

- Support for the `IBMDB2I` storage engine has been removed. (Bug #58079)
- The `pstack` library was nonfunctional and has been removed, along with the `--with-pstack` option for `configure`. The `--enable-pstack` option for `mysqld` is deprecated and will be removed in MySQL 5.5. (Bug #57210)

Bugs Fixed

- **Security Fix:** A security bug was fixed. (Bug #58005)
- **Performance; InnoDB:** Improved concurrency when several `ANALYZE TABLE` or `SHOW TABLE STATUS` statements are run simultaneously for `InnoDB` tables. (Bug #53046)
- **InnoDB:** Dropping an `InnoDB` index used by a foreign key constraint, while `foreign_key_checks` was set to 0, could cause the server to crash with an assertion. This operation now does not cause a crash. The foreign key constraint can no longer be enforced once the associated index is removed, so do not rely on it for referential integrity in this case. (Bug #11762483, Bug #55084)
- **InnoDB:** For an `InnoDB` table created with `ROW_FORMAT=COMPRESSED` or `ROW_FORMAT=DYNAMIC`, a query using the `READ UNCOMMITTED` isolation level could cause the server to stop with an assertion error, if `BLOB` or other large columns that use off-page storage were being inserted at the same time. (Bug #57799)
- **InnoDB:** If the MySQL Server crashed immediately after creating an `InnoDB` table, the server could quit with a `signal 11` during the subsequent restart. The issue could occur if the server halted after `InnoDB` created the primary index for the table, but before the index definition was recorded in the MySQL metadata. (Bug #57616)

References: This issue is a regression of: Bug #54582.

- **InnoDB:** An existing `InnoDB` table could be switched to `ROW_FORMAT=COMPRESSED` implicitly by a `KEY_BLOCK_SIZE` clause in an `ALTER TABLE` statement. Now, the row format is only switched to compressed if there is an explicit `ROW_FORMAT=COMPRESSED` clause. on the `ALTER TABLE` statement.

Any valid, nondefault `ROW_FORMAT` parameter takes precedence over `KEY_BLOCK_SIZE` when both are specified. `KEY_BLOCK_SIZE` only enables `ROW_FORMAT=COMPRESSED` if `ROW_FORMAT` is not specified on either the `CREATE TABLE` or `ALTER TABLE` statement, or is specified as `DEFAULT`. In case of a conflict between `KEY_BLOCK_SIZE` and `ROW_FORMAT` clauses, the `KEY_BLOCK_SIZE` is ignored if `innodb_strict_mode` is off, and the statement causes an error if `innodb_strict_mode` is on. (Bug #56632)

- **InnoDB:** The clause `KEY_BLOCK_SIZE=0` is now permitted on `CREATE TABLE` and `ALTER TABLE` statements for `InnoDB` tables, regardless of the setting of `innodb_strict_mode`. The zero value has the effect of resetting the `KEY_BLOCK_SIZE` table parameter to its default value, depending on the `ROW_FORMAT` parameter, as if it had not been specified. That default is 8 if `ROW_FORMAT=COMPRESSED`. Otherwise, `KEY_BLOCK_SIZE` is not used or stored with the table parameters.

As a consequence of this fix, `ROW_FORMAT=FIXED` is not permitted when `innodb_strict_mode` is enabled. (Bug #56628)

- **InnoDB:** `InnoDB` startup messages now include the start and end times for buffer pool initialization, and the total buffer pool size. (Bug #48026)
- **Partitioning:** In debug builds, an `INSERT ... ON DUPLICATE KEY UPDATE col_name = 0` statement on an `AUTO_INCREMENT` column caused the server to crash. (Bug #57890)
- **Partitioning:** An `ALTER TABLE` statement acting on table partitions that failed while the affected table was locked could cause the server to crash. (Bug #56172)
- Several compilation problems were fixed. (Bug #57992, Bug #57993, Bug #57994, Bug #57995, Bug #57996, Bug #57997, Bug #58057)
- Passing a string that was not null-terminated to `UpdateXML()` or `ExtractValue()` caused the server to fail with an assertion. (Bug #57279, Bug #11764447)
- Queries executed using the Index Merge access method and a temporary file could return incorrect results. (Bug #56862)
- Valgrind warnings about overlapping memory when double-assigning the same variable were corrected. (Bug #56138)
- An error in a stored procedure could leave the session in a different default database. (Bug #54375)
- Grouping by a `TIME_TO_SEC()` function result could cause a server crash or incorrect results. Grouping by a function returning a `BLOB` could cause an unexpected “Duplicate entry” error and incorrect result. (Bug #52160)
- The `find_files()` function used by `SHOW` statements performed redundant and unnecessary memory allocation. (Bug #51208)
- The Windows sample option files contained values more appropriate for Linux. (Bug #50021)
- A failed `RENAME TABLE` operation could prevent a `FLUSH TABLES WITH READ LOCK` from completing. (Bug #47924)

Changes in MySQL 5.1.53 (2010-11-03)

Bugs Fixed

- **Security Fix; InnoDB:** A failed `CREATE TABLE` statement for an `InnoDB` table could allocate memory that was never freed. (Bug #56947)
- **Security Fix:** A security bug was fixed. (Bug #57477)
- **Security Fix:** A security bug was fixed. (Bug #57272)
- **Security Fix:** A security bug was fixed. (Bug #56814)
- **Security Fix:** A security bug was fixed. (Bug #55146, Bug #56287)
- **Security Fix:** A security bug was fixed. (Bug #54484)
- **InnoDB:** A followup fix to bug #54678. `TRUNCATE TABLE` could still cause a crash (assertion error) in the debug version of the server. (Bug #57700)

References: See also: Bug #54678.

- **InnoDB:** The `InnoDB` system tablespace could grow continually for a server under heavy load. (Bug #57611)

- **InnoDB:** Turning off the `innodb_stats_on_metadata` option could prevent the `ANALYZE TABLE` statement from updating the cardinality statistics of `InnoDB` tables. (Bug #57252)
- **InnoDB:** If the server crashed during an `ALTER TABLE` operation on an `InnoDB` table, examining the table through `SHOW CREATE TABLE` or querying the `INFORMATION_SCHEMA` tables could cause the server to stop with an assertion error. (Bug #56982)
- **InnoDB:** A query for an `InnoDB` table could return the wrong value if a column value was changed to a different case, and the column had a case-insensitive index. (Bug #56680, Bug #11763909)
- **InnoDB:** A large number of foreign key declarations could cause the output of the `SHOW CREATE STATEMENT` statement to be truncated. (Bug #56143)
- **InnoDB:** A compilation problem affected the `InnoDB` source code on NetBSD/sparc64. (Bug #53916)

References: See also: Bug #59327.

- **Replication:** `SET PASSWORD` caused failure of row-based replication between a MySQL 5.1 master and a MySQL 5.5 slave.

This fix makes it possible to replicate `SET PASSWORD` correctly, using row-based replication between a master running MySQL 5.1.53 or a later MySQL 5.1 release to a slave running MySQL 5.5.7 or a later MySQL 5.5 release. (Bug #57098)

References: See also: Bug #55452, Bug #57357.

- **Replication:** An `ALTER TABLE` statement that altered a column of a `MyISAM` table without setting the column's size caused the binary log to become corrupted when the table map was unexpectedly set to 0 by updates (including deletes) on multiple tables, leading to replication failure when more than one table received the same table map ID. (Bug #56226, Bug #11763509)
- **Replication:** When `STOP SLAVE` is issued, the slave SQL thread rolls back the current transaction and stops immediately if the transaction updates only tables which use transactional storage engines. Previously, this occurred even when the transaction contained `CREATE TEMPORARY TABLE` statements, `DROP TEMPORARY TABLE` statements, or both, although these statements cannot be rolled back. Because temporary tables persist for the lifetime of a user session (in the case, the replication user), they remain until the slave is stopped or reset. When the transaction is restarted following a subsequent `START SLAVE` statement, the SQL thread aborts with an error that a temporary table to be created (or dropped) already exists (or does not exist, in the latter case).

Following this fix, if an ongoing transaction contains `CREATE TEMPORARY TABLE` statements, `DROP TEMPORARY TABLE` statements, or both, the SQL thread now waits until the transaction ends, then stops. (Bug #56118, Bug #11763416)

- **Replication:** If there exist both a temporary table and a nontemporary table having the same name, updates normally apply only to the temporary table, with the exception of a `CREATE TABLE ... SELECT` statement that creates a nontemporary table having the same name as an existing temporary table. When such a statement was replicated using the `MIXED` logging format, and the statement was unsafe for row-based logging, updates were misapplied to the temporary table. (Bug #55478)

References: See also: Bug #47899, Bug #55709.

- **Replication:** When a slave tried to execute a transaction larger than the slave's value for `max_binlog_cache_size`, it crashed. This was caused by an assertion that the server should roll back only the statement but not the entire transaction when the error `ER_TRANS_CACHE_FULL` occurred. However, the slave SQL thread always rolled back the entire transaction whenever any error occurred, regardless of the type of error. (Bug #55375)

- **Replication:** When making changes to relay log settings using `CHANGE MASTER TO`, the I/O cache was not cleared. This could result in replication failure when the slave attempted to read stale data from the cache and then stopped with an assertion. (Bug #55263)
- **Replication:** Trying to read from a binary log containing a log event of an invalid type caused the slave to crash. (Bug #38718)
- **Replication:** When replicating the `mysql.tables_priv` table, the `Grantor` column was not replicated, and was thus left empty on the slave. (Bug #27606)
- `SET GLOBAL debug` could cause a crash on Solaris if the server failed to open the trace file. (Bug #57274)
- A `SELECT` statement could produce a number of rows different from a `CREATE TABLE ... SELECT` that was supposed to select the same rows. (Bug #56423)

References: This issue is a regression of: Bug #38999.

- On file systems with case insensitive file names, and `lower_case_table_names=2`, the server could crash due to a table definition cache inconsistency. (Bug #46941)
- Handling of host name lettercase in `GRANT` statements was inconsistent. (Bug #36742)

Changes in MySQL Enterprise 5.1.52sp1 [QSP] (2011-02-21)

This is a *Service Pack* release of the MySQL Enterprise Server 5.1.

- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

Functionality Added or Changed

- Support for the `IBMDB2I` storage engine has been removed. (Bug #58079)

Bugs Fixed

- **Incompatible Change:** When `auto_increment_increment` is greater than one, values generated by a bulk insert that reaches the maximum column value could wrap around rather producing an overflow error.

As a consequence of the fix, it is no longer possible for an auto-generated value to be equal to the maximum `BIGINT UNSIGNED` value. It is still possible to store that value manually, if the column can accept it. (Bug #39828, Bug #11749800)

- **Important Change; Partitioning:** Date and time functions used as partitioning functions now have the types of their operands checked; use of a value of the wrong type is now disallowed in such cases. In addition, `EXTRACT(WEEK FROM col_name)`, where `col_name` is a `DATE` or `DATETIME` column, is now disallowed altogether because its return value depends on the value of the `default_week_format` system variable. (Bug #54483, Bug #11761948)

References: See also: Bug #57071, Bug #11764255.

- **InnoDB:** The `InnoDB` system tablespace could grow continually for a server under heavy load. (Bug #57611)
- **InnoDB:** If the server crashed during an `ALTER TABLE` operation on an `InnoDB` table, examining the table through `SHOW CREATE TABLE` or querying the `INFORMATION_SCHEMA` tables could cause the server to stop with an assertion error. (Bug #56982)

- **Partitioning:** Failed `ALTER TABLE ... PARTITION` statements could cause memory leaks. (Bug #56380, Bug #11763641)

References: See also: Bug #46949, Bug #11755209, Bug #56996, Bug #11764187.

- `MIN(year_col)` could return an incorrect result in some cases. (Bug #59211, Bug #11766165)
- `EXPLAIN` could crash for queries that used `GROUP_CONCAT()`. (Bug #58396)
- Unnecessary subquery evaluation in contexts such as statement preparation or view creation could cause a server crash. (Bug #57703)
- View creation could produce Valgrind warnings. (Bug #57352)
- `NULL` geometry values could cause a crash in `Item_func_spatial_collection::fix_length_and_dec`. (Bug #57321)
- `SET GLOBAL debug` could cause a crash on Solaris if the server failed to open the trace file. (Bug #57274)
- Valgrind warnings about overlapping memory when double-assigning the same variable were corrected. (Bug #56138)
- On FreeBSD, if `mysqld` was killed with a `SIGHUP` signal, it could corrupt InnoDB `.ibd` files. (Bug #51023, Bug #11758773)

Changes in MySQL 5.1.52 (2010-10-11)

Bugs Fixed

- **Security Fix; InnoDB:** Issuing `TRUNCATE TABLE` and examining the same table's information in the `INFORMATION_SCHEMA` database at the same time could cause a crash in the debug version of the server. (Bug #54678)
- **Security Fix:** The server crashed for assignment of values of types other than `Geometry` to items of type `GeometryCollection` (`MultiPoint`, `MultiCurve`, `MultiSurface`). Now the server checks the value type and fails with `bad geometry value` if it detects incorrect parameters. (Bug #55531)
- **Security Fix:** `EXPLAIN EXTENDED` caused a server crash with some prepared statements. (Bug #54494)
- **Security Fix:** In prepared-statement mode, `EXPLAIN` for a `SELECT` from a derived table caused a server crash. (Bug #54488)
- **Important Change; Replication:** The `LOAD DATA INFILE` statement is now considered unsafe for statement-based replication. When using statement-based logging mode, the statement now produces a warning; when using mixed-format logging, the statement is made using the row-based format. (Bug #34283)
- **InnoDB:** The server could crash with a high volume of concurrent `LOCK TABLES` and `UNLOCK TABLES` statements. (Bug #57345)
- **InnoDB:** InnoDB incorrectly reported an error when a cascading foreign key constraint deleted more than 250 rows. (Bug #57255)
- **InnoDB:** For debug builds, a `SELECT ... FOR UPDATE` statement affecting a range of rows in an InnoDB table could cause a server crash. (Bug #56716)

- **InnoDB:** Improved the performance of `UPDATE` operations on `InnoDB` tables, when only non-indexed columns are changed. (Bug #56340)
- **InnoDB:** The server could crash on shutdown, if started with `--innodb-use-system-malloc=0`. (Bug #55627)
- **InnoDB:** For an `InnoDB` table with an auto-increment column, the server could crash if the first statement that references the table after a server restart is a `SHOW CREATE TABLE` statement. (Bug #55277)
- **InnoDB:** Setting the `PACK_KEYS=0` table option for an `InnoDB` table prevented new indexes from being added to the table. (Bug #54606)
- **InnoDB:** Changed the locking mechanism for the `InnoDB` data dictionary during `ROLLBACK` operations, to improve concurrency for `REPLACE` statements. (Bug #54538)
- **InnoDB:** `InnoDB` transactions could be incorrectly committed during recovery, rather than rolled back, if the server crashed and was restarted after performing `ALTER TABLE ... ADD PRIMARY KEY` on an `InnoDB` table, or some other operation that involves copying the entire table. (Bug #53756)
- **Partitioning; Replication:** Attempting to execute `LOAD DATA` on a partitioned `MyISAM` table while using statement-based logging mode caused the master to hang or crash. (Bug #51851)
- **Partitioning:** Multiple-table `UPDATE` statements involving a partitioned `MyISAM` table could cause this table to become corrupted. Not all tables affected by the `UPDATE` needed to be partitioned for this issue to be observed. (Bug #55458)
- **Partitioning:** `EXPLAIN PARTITIONS` returned bad estimates for range queries on partitioned `MyISAM` tables. In addition, values in the `rows` column of `EXPLAIN PARTITIONS` output did not take partition pruning into account. (Bug #53806, Bug #46754)
- **Replication:** Backticks used to enclose identifiers for savepoints were not preserved in the binary log, which could lead to replication failure when the identifier, stripped of backticks, could be misinterpreted, causing a syntax or other error.

This could cause problems with MySQL application programs making use of generated savepoint IDs. If, for instance, `java.sql.Connection.setSavepoint()` is called without any parameters, Connector/J automatically generates a savepoint identifier consisting of a string of hexadecimal digits `0-F` encased in backtick (```) characters. If such an ID took the form ``NeN`` (where `N` represents a string of the decimal digits `0-9`, and `e` is a literal uppercase or lowercase “E” character). Removing the backticks when writing the identifier into the binary log left behind a substring which the slave MySQL server tried to interpret as a floating point number, rather than as an identifier. The resulting syntax error caused loss of replication. (Bug #55961)

References: See also: Bug #55962.

- When `mysqld` was started as a service on Windows and `mysqld` was writing the error log to a file (for example, if it was started with the `--log-error` option), the server reassigned the file descriptors of the `stdout` and `stderr` streams to the file descriptor of the log file. On Windows, if `stdout` or `stderr` is not associated with an output stream, the file descriptor returns a negative value. Previously, this caused the file descriptor reassignment to fail and the server to abort. To avoid this problem on Windows, the server now first assigns the `stdout` and `stderr` streams to the log file stream by opening this file. This causes the `stdout` and `stderr` file descriptors to be nonzero and the server can successfully reassign them to the file descriptor of the log file. (Bug #56821)

References: This issue is a regression of: Bug #29751.

- Memory leaks detected by Valgrind were corrected. (Bug #56709)

- If a query specified a `DATE` or `DATETIME` value in a format different from `'YYYY-MM-DD HH:MM:SS'`, a greater-than-or-equal (`>=`) condition matched only greater-than values in an indexed `TIMESTAMP` column. (Bug #55779, Bug #50774, Bug #11758558)
- If there was an active `SELECT` statement, an error arising during trigger execution could cause a server crash. (Bug #55421)
- With an `UPDATE IGNORE` statement including a subquery that was evaluated using a temporary table, an error transferring the data from the temporary was ignored, causing an assertion to be raised. (Bug #54543)
- Row subqueries producing no rows were not handled as `UNKNOWN` values in row-comparison expressions. (Bug #54190)
- The `max_length` metadata value of `MEDIUMBLOB` types was reported as 1 byte greater than the correct value. (Bug #53296)
- In some cases, when the left part of a `NOT IN` subquery predicate was a row and contained `NULL` values, the query result was incorrect. (Bug #51070)
- For some queries, the optimizer produced incorrect results using the Index Merge access method with `InnoDB` tables. (Bug #50402)
- `EXPLAIN` produced an incorrect `rows` value for queries evaluated using an index scan and that included `LIMIT`, `GROUP BY`, and `ORDER BY` on a computed column. (Bug #50394)
- `mysql_store_result()` and `mysql_use_result()` are not for use with prepared statements and are not intended to be called following `mysql_stmt_execute()`, but failed to return an error when invoked that way. (Bug #47485)
- Using `REPAIR TABLE tbl_name USE_FRM` on a `MERGE` table caused the server to crash. (Bug #46339)
- A malformed packet sent by the server when the query cache was in use resulted in lost-connection errors. (Bug #42503)
- `CREATE TABLE` failed if a column referred to in an index definition and foreign key definition had different lettercases in the two definitions. (Bug #39932)

Changes in MySQL 5.1.51 (2010-09-10)

- [InnoDB Plugin Notes](#)
- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

InnoDB Plugin Notes

- `InnoDB Plugin` has been upgraded to version 1.0.12. This version is considered of General Availability (GA) quality.

In this release, the `InnoDB Plugin` is included in source and binary distributions, except RHEL3, RHEL4, SuSE 9 (x86, x86_64, ia64), and generic Linux RPM packages. It also does not work for FreeBSD 6 and HP-UX or for Linux on generic ia64.

Functionality Added or Changed

- **Incompatible Change:** Previously, if you flushed the logs using `FLUSH LOGS` or `mysqladmin flush-logs` and `mysqld` was writing the error log to a file (for example, if it was started with the `--log-error` option), it renamed the current log file with the suffix `-old`, then created a new empty log file. This had the problem that a second log-flushing operation thus caused the original error log file to be lost unless you saved it under a different name. For example, you could use the following commands to save the file:

```
shell> mysqladmin flush-logs
shell> mv host_name.err-old backup-directory
```

To avoid the preceding file-loss problem, renaming no longer occurs. The server merely closes and reopens the log file. To rename the file, you can do so manually before flushing. Then flushing the logs reopens a new file with the original file name. For example, you can rename the file and create a new one using the following commands:

```
shell> mv host_name.err host_name.err-old
shell> mysqladmin flush-logs
shell> mv host_name.err-old backup-directory
```

(Bug #29751)

References: See also: Bug #56821.

Bugs Fixed

- **Security Fix:** During evaluation of arguments to extreme-value functions such as `LEAST()` and `GREATEST()`, type errors did not propagate properly, causing the server to crash. (Bug #55826, CVE-2010-3833)
- **Security Fix:** The server could crash after materializing a derived table that required a temporary table for grouping. (Bug #55568, CVE-2010-3834)
- **Security Fix:** A user-variable assignment expression that is evaluated in a logical expression context can be precalculated in a temporary table for `GROUP BY`. However, when the expression value is used after creation of the temporary table, it was re-evaluated, not read from the table, and a server crash resulted. (Bug #55564, CVE-2010-3835)
- **Security Fix:** The `CONVERT_TZ()` function crashed the server when the timezone argument was an empty `SET` column value. (Bug #55424)
- **Security Fix:** Pre-evaluation of `LIKE` predicates during view preparation could cause a server crash. (Bug #54568, Bug #11762026, CVE-2010-3836)
- **Security Fix:** `GROUP_CONCAT()` and `WITH ROLLUP` together could cause a server crash. (Bug #54476, CVE-2010-3837)
- **Security Fix:** Queries could cause a server crash if the `GREATEST()` or `LEAST()` function had a mixed list of numeric and `LONGBLOB` arguments, and the result of such a function was processed using an intermediate temporary table. (Bug #54461, CVE-2010-3838)
- **Security Fix:** Queries with nested joins could cause an infinite loop in the server when used from stored procedures and prepared statements. (Bug #53544, CVE-2010-3839)
- **Security Fix:** The `PolyFromWKB()` function could crash the server when improper WKB data was passed to the function. (Bug #51875, Bug #11759554, CVE-2010-3840)

- **Incompatible Change; Replication:** As of MySQL 5.5.6, handling of `CREATE TABLE IF NOT EXISTS ... SELECT` statements has been changed for the case that the destination table already exists:
 - Previously, for `CREATE TABLE IF NOT EXISTS ... SELECT`, MySQL produced a warning that the table exists, but inserted the rows and wrote the statement to the binary log anyway. By contrast, `CREATE TABLE ... SELECT` (without `IF NOT EXISTS`) failed with an error, but MySQL inserted no rows and did not write the statement to the binary log.
 - MySQL now handles both statements the same way when the destination table exists, in that neither statement inserts rows or is written to the binary log. The difference between them is that MySQL produces a warning when `IF NOT EXISTS` is present and an error when it is not.

This change in handling of `IF NOT EXISTS` results in an incompatibility for statement-based replication from a MySQL 5.1 master with the original behavior and a MySQL 5.5 slave with the new behavior. Suppose that `CREATE TABLE IF NOT EXISTS ... SELECT` is executed on the master and the destination table exists. The result is that rows are inserted on the master but not on the slave. (Row-based replication does not have this problem.)

To address this issue, statement-based binary logging for `CREATE TABLE IF NOT EXISTS ... SELECT` is changed in MySQL 5.1 as of 5.1.51:

- If the destination table does not exist, there is no change: The statement is logged as is.
- If the destination table does exist, the statement is logged as the equivalent pair of `CREATE TABLE IF NOT EXISTS` and `INSERT ... SELECT` statements. (If the `SELECT` in the original statement is preceded by `IGNORE` or `REPLACE`, the `INSERT` becomes `INSERT IGNORE` or `REPLACE`, respectively.)

This change provides forward compatibility for statement-based replication from MySQL 5.1 to 5.5 because when the destination table exists, the rows will be inserted on both the master and slave. To take advantage of this compatibility measure, the 5.1 server must be at least 5.1.51 and the 5.5 server must be at least 5.5.6.

To upgrade an existing 5.1-to-5.5 replication scenario, upgrade the master first to 5.1.51 or higher. Note that this differs from the usual replication upgrade advice of upgrading the slave first.

A workaround for applications that wish to achieve the original effect (rows inserted regardless of whether the destination table exists) is to use `CREATE TABLE IF NOT EXISTS` and `INSERT ... SELECT` statements rather than `CREATE TABLE IF NOT EXISTS ... SELECT` statements.

Along with the change just described, the following related change was made: Previously, if an existing view was named as the destination table for `CREATE TABLE IF NOT EXISTS ... SELECT`, rows were inserted into the underlying base table and the statement was written to the binary log. As of MySQL 5.1.51 and 5.5.6, nothing is inserted or logged. (Bug #47442, Bug #47132, Bug #48814, Bug #49494)

- **Important Change; InnoDB:** The server could crash with an assertion, possibly leading to data corruption, while updating the primary key of an InnoDB table containing `BLOB` or other columns requiring off-page storage. This fix applies to the InnoDB Plugin in MySQL 5.1, and to InnoDB 1.1 in MySQL 5.5. (Bug #55543)
- **InnoDB:** When MySQL was restarted after a crash with the option `innodb_force_recovery=6`, certain queries against InnoDB tables could fail, depending on `WHERE` or `ORDER BY` clauses.

Usually in such a disaster recovery situation, you dump the entire table using a query without these clauses. During advanced troubleshooting, you might use queries with these clauses to diagnose the position of the corrupted data, or to recover data following the corrupted part. (Bug #55832)

- **InnoDB:** The `CHECK TABLE` command could cause a time-consuming verification of the `InnoDB` adaptive hash index memory structure. Now this extra checking is only performed in binaries built for debugging. (Bug #55716)
- **InnoDB:** A heavy workload with a large number of threads could cause a crash in the debug version of the server. (Bug #55699)
- **InnoDB:** If the server crashed during a `RENAME TABLE` operation on an `InnoDB` table, subsequent crash recovery could fail. This problem could also affect an `ALTER TABLE` statement that caused a rename operation internally. (Bug #55027)
- **InnoDB:** The server could crash when opening an `InnoDB` table linked through foreign keys to a long chain of child tables. (Bug #54582, Bug #11762038)
- **Partitioning:** When the storage engine used to create a partitioned table was disabled, attempting to drop the table caused the server to crash. (Bug #46086)
- If a view was named as the destination table for `CREATE TABLE ... SELECT`, the server produced a warning whether or not `IF NOT EXISTS` was used. Now it produces a warning only when `IF NOT EXISTS` is used, and an error otherwise. (Bug #55777)
- After the fix for Bug #39653, the shortest available secondary index was used for full table scans. The primary clustered key was used only if no secondary index could be used. However, when the chosen secondary index includes all columns of the table being scanned, it is better to use the primary index because the amount of data to scan is the same but the primary index is clustered. This is now taken into account. (Bug #55656)

References: See also: Bug #39653.

- The server was not checking for errors generated during the execution of `Item::val_xxx()` methods when copying data to a group, order, or distinct temp table's row. (Bug #55580)
- `ORDER BY` clauses that included user-variable expressions could raise a debug assertion. (Bug #55565)
- Assignment of `InnoDB` scalar subquery results to a variable resulted in unexpected `S` locks in `READ COMMITTED` transaction isolation level. (Bug #55382)
- Queries involving predicates of the form `const NOT BETWEEN not_indexed_column AND indexed_column` could return incorrect data due to incorrect handling by the range optimizer. (Bug #54802)
- `MIN()` or `MAX()` with a subquery argument could raise a debug assertion for debug builds or return incorrect data for nondebug builds. (Bug #54465)
- `INFORMATION_SCHEMA` plugins with no `deinit()` method resulted in a memory leak. (Bug #54253)
- After `ALTER TABLE` was used on a temporary transactional table locked by `LOCK TABLES`, any later attempts to execute `LOCK TABLES` or `UNLOCK TABLES` caused a server crash. (Bug #54117)
- `INSERT IGNORE INTO ... SELECT` statements could raise a debug assertion. (Bug #54106)
- `INFORMATION_SCHEMA.COLUMNS` reported incorrect precision for `BIGINT UNSIGNED` columns. (Bug #53814)

- The fix for Bug #30234 caused the server to reject the `DELETE tbl_name.* ...` Access compatibility syntax for multiple-table `DELETE` statements. (Bug #53034)

References: See also: Bug #30234.

- `XA START` had a race condition that could cause a server crash. (Bug #51855)
- Enumeration plugin variables were subject to a type-casting error, causing inconsistent results between different platforms. (Bug #42144)
- A PKG install on Solaris put some files in incorrect locations. (Bug #31058)
- Problems in the atomic operations implementation could lead to server crashes. (Bug #22320, Bug #52261)

Changes in MySQL 5.1.50 (2010-08-03)

- [icc Notes](#)
- [InnoDB Plugin Notes](#)
- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

icc Notes

- This is the final release of MySQL 5.1 for which Generic Linux MySQL binary packages built with the `icc` compiler on x86 and x86_64 will be offered. These were previously produced as an alternative to our main packages built using `gcc`, as they provided noticeable performance benefits. In recent times the performance differences have diminished and build and runtime problems have surfaced, thus it is no longer viable to continue producing them.

We continue to use the `icc` compiler to produce our distribution-specific RPM packages on ia64.

InnoDB Plugin Notes

- `InnoDB Plugin` has been upgraded to version 1.0.11. This version is considered of General Availability (GA) quality.

In this release, the `InnoDB Plugin` is included in source and binary distributions, except RHEL3, RHEL4, SuSE 9 (x86, x86_64, ia64), generic Linux RPM packages, and any builds produced with the `icc` compiler. It also does not work for FreeBSD 6 and HP-UX or for Linux on generic ia64.

Functionality Added or Changed

- `REPAIR TABLE` and `OPTIMIZE TABLE` table now catch and throw any errors that occur while copying table statistics from the old corrupted file to newly created file. For example, if the user ID of the owner of the `.frm`, `.MYD`, or `.MYI` file is different from the user ID of the `mysqld` process, `REPAIR TABLE` and `OPTIMIZE TABLE` generate a "cannot change ownership of the file" error unless `mysqld` is started by the `root` user. (Bug #61598, Bug #13600058)

Bugs Fixed

- **Security Fix:** A security bug was fixed. (Bug #49124)
- **InnoDB:** The server could crash on shutdown, if started with `--innodb-use-system-malloc=0`. (Bug #55581, Bug #11762927)

- **InnoDB:** For debug builds, the database server could crash when renaming a table that had active transactions. (Bug #54453)
- **InnoDB:** The server could crash during the recovery phase of startup, if it previously crashed while inserting `BLOB` or other large columns that use off-page storage into an `InnoDB` table created with `ROW_FORMAT=REDUNDANT` or `ROW_FORMAT=COMPACT`. (Bug #54408)
- **InnoDB:** For an `InnoDB` table created with `ROW_FORMAT=COMPRESSED` or `ROW_FORMAT=DYNAMIC`, a query using the `READ UNCOMMITTED` isolation level could cause the server to stop with an assertion error, if `BLOB` or other large columns that use off-page storage were being inserted at the same time. (Bug #54358)
- **Partitioning:** `UPDATE` and `INSERT` statements affecting partitioned tables performed poorly when using row-based replication. (Bug #52517)

References: This issue is a regression of: Bug #39084.

- **Partitioning:** `INSERT ON DUPLICATE KEY UPDATE` statements performed poorly on tables having many partitions. The handler function for reading a row from a specific index was not optimized in the partitioning handler. (Bug #52455)
- **Replication:** When using the row-based logging format, a failed `CREATE TABLE ... SELECT` statement was written to the binary log, causing replication to break if the failed statement was later re-run on the master. In such cases, a `DROP TABLE ... IF EXIST` statement is now logged in the event that a `CREATE TABLE ... SELECT` fails. (Bug #55625)
- `GROUP BY` operations used `max_sort_length` inconsistently. (Bug #55188)
- Building MySQL on Solaris 8 x86 failed when using Sun Studio due to `gcc` inline assembly code. (Bug #55061)
- In debug builds, an assertion could be raised when the server tried to send an OK packet to the client after having failed to detect errors during processing of the `WHERE` condition of an `UPDATE` statement. (Bug #54734)
- A join with an aggregated function and impossible `WHERE` condition returned an extra row. (Bug #54416)
- A client could supply data in chunks to a prepared statement parameter other than of type `TEXT` or `BLOB` using the `mysql_stmt_send_long_data()` C API function (or `COM_STMT_SEND_LONG_DATA` command). This led to a crash because other data types are not valid for long data. (Bug #54041)
- `mysql_secure_installation` did not properly identify local accounts and could incorrectly remove nonlocal `root` accounts. (Bug #54004)
- Portability problems in `SHOW STATUS` could lead to incorrect results on some platforms. (Bug #53493)
- Builds of MySQL generated a large number of warnings. (Bug #53445)
- With `lower_case_table_names` set to a nonzero value, searches for table or database names in `INFORMATION_SCHEMA` tables could produce incorrect results. (Bug #53095)
- The ABI check for MySQL failed to compile with `gcc` 4.5. (Bug #52514)
- `mysql_secure_installation` sometimes failed to locate the `mysql` client. (Bug #52274)
- Reading a `ucs2` data file with `LOAD DATA INFILE` was subject to three problems. 1) Incorrect parsing of the file as `ucs2` data, resulting in incorrect length of the parsed string. This is fixed by truncating

the invalid trailing bytes (incomplete multibyte characters) when reading from the file. 2) Reads from a proper `ucs2` file did not recognize newline characters. This is fixed by first checking whether a byte is a newline (or any other special character) before reading it as a part of a multibyte character. 3) When using user variables to hold column data, the character set of the user variable was set incorrectly to the database charset. This is fixed by setting it to the character set specified in the `LOAD DATA INFILE` statement, if any. (Bug #51876)

- Searches in `INFORMATION_SCHEMA` tables for rows matching a nonexistent database produced an error instead of an empty query result. (Bug #49542)
- On FreeBSD, memory mapping for `MERGE` tables could fail if underlying tables were empty. (Bug #47139)
- The `my_like_range_xxx()` functions returned badly formed maximum strings for Asian character sets, which caused problems for storage engines. (Bug #45012)
- A debugging assertion could be raised after a write failure to a closed socket. (Bug #42496)
- An assertion failure occurred within yaSSL for very long keys. (Bug #29784)

References: See also: Bug #53463.

Changes in MySQL Enterprise 5.1.49sp1 [QSP] (2010-09-28)

This is a *Service Pack* release of the MySQL Enterprise Server 5.1.

Bugs Fixed

- Building MySQL on Solaris 8 x86 failed when using Sun Studio due to `gcc` inline assembly code. (Bug #55061)
- A client could supply data in chunks to a prepared statement parameter other than of type `TEXT` or `BLOB` using the `mysql_stmt_send_long_data()` C API function (or `COM_STMT_SEND_LONG_DATA` command). This led to a crash because other data types are not valid for long data. (Bug #54041)

Changes in MySQL 5.1.49 (2010-07-09)

- [InnoDB Plugin Notes](#)
- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

InnoDB Plugin Notes

- `InnoDB Plugin` has been upgraded to version 1.0.10. This version is considered of General Availability (GA) quality.

In this release, the `InnoDB Plugin` is included in source and binary distributions, except RHEL3, RHEL4, SuSE 9 (x86, x86_64, ia64), generic Linux RPM packages, and any builds produced with the `icc` compiler. It also does not work for FreeBSD 6 and HP-UX or for Linux on generic ia64.

Functionality Added or Changed

- **InnoDB:** The mechanism that checks if there is enough space for redo logs was improved, reducing the chance of encountering this message: `ERROR: the age of the last checkpoint is x, which exceeds the log group capacity y`. (Bug #39168)

Bugs Fixed

- **Security Fix; InnoDB:** After changing the values of the `innodb_file_format` or `innodb_file_per_table` configuration parameters, DDL statements could cause a server crash. (Bug #55039, CVE-2010-3676)
- **Security Fix:** Joins involving a table with a unique `SET` column could cause a server crash. (Bug #54575, CVE-2010-3677)
- **Security Fix:** Incorrect handling of `NULL` arguments could lead to a crash for `IN()` or `CASE` operations when `NULL` arguments were either passed explicitly as arguments (for `IN()`) or implicitly generated by the `WITH ROLLUP` modifier (for `IN()` and `CASE`). (Bug #54477, CVE-2010-3678)
- **Security Fix:** A malformed argument to the `BINLOG` statement could result in Valgrind warnings or a server crash. (Bug #54393, CVE-2010-3679)
- **Security Fix:** Use of `TEMPORARY InnoDB` tables with nullable columns could cause a server crash. (Bug #54044, CVE-2010-3680)
- **Security Fix:** The server could crash if there were alternate reads from two indexes on a table using the `HANDLER` interface. (Bug #54007, CVE-2010-3681)
- **Security Fix:** A security bug was fixed. (Bug #53933)
- **Security Fix:** Using `EXPLAIN` with queries of the form `SELECT ... UNION ... ORDER BY (SELECT ... WHERE ...)` could cause a server crash. (Bug #52711, CVE-2010-3682)
- **Security Fix:** `LOAD DATA INFILE` did not check for SQL errors and sent an OK packet even when errors were already reported. Also, an assert related to client/server protocol checking in debug servers sometimes was raised when it should not have been. (Bug #52512, CVE-2010-3683)
- **Security Fix:** A security bug was fixed. (Bug #52315)
- **InnoDB:** An `ALTER TABLE` statement could convert an `InnoDB` compressed table (with `row_format=compressed`) back to an uncompressed table (with `row_format=compact`). (Bug #54679)
- **InnoDB:** `InnoDB` could issue an incorrect message on startup, if tables were created under the setting `innodb_file_per_table=ON`. The message was of the form `InnoDB: Warning: allocated tablespace n, old maximum was 0`. If you encounter this message after upgrading, create an `InnoDB` table with `innodb_file_per_table = ON` and restart the server. The message should not be displayed any more. If you continue to encounter this message, or if you get it and haven't used a version without this fix, you might have corruption in your shared tablespace. If so, back up and reload your data. (Bug #54658)
- **InnoDB:** Fast index creation in the `InnoDB Plugin` could fail, leaving the new secondary index corrupted. (Bug #54330)
- **InnoDB:** Some combinations of `SELECT` and `SELECT FOR UPDATE` statements could fail with errors about locks, or incorrectly release a row lock during a `semi-consistent read` operation. (Bug #53674)
- **InnoDB:** Performing large numbers of `RENAME TABLE` statements caused excessive memory use. (Bug #47991)
- **Replication:** When using unique keys on `NULL` columns in row-based replication, the slave sometimes chose the wrong row when performing an update. This happened because a table having a unique key on such a column could have multiple rows containing `NULL` for the column used by the unique key, and the slave merely picked the first row containing `NULL` in that column. (Bug #53893)

- **Replication:** `FLUSH LOGS` could in some circumstances crash the server. This occurred because the I/O thread could concurrently access the relay log I/O cache while another thread was performing the `FLUSH LOGS`, which closes and reopens the relay log and, while doing so, initializes (or re-initializes) its I/O cache. This could cause problems if some other thread (in this case, the I/O thread) is accessing it at the same time.

Now the thread performing the `FLUSH LOGS` takes a lock on the relay log before actually flushing it. (Bug #53657)

References: See also: Bug #50364.

- **Replication:** Two related issues involving temporary tables and transactions were introduced by a fix made in MySQL 5.1.37:
 1. When a temporary table was created or dropped within a transaction, any failed statement that following the `CREATE TEMPORARY TABLE` or `DROP TEMPORARY TABLE` statement triggered a rollback, which caused the slave to diverge from the master.
 2. When a `CREATE TEMPORARY TABLE ... SELECT * FROM ...` statement was executed within a transaction in which only tables using transactional storage engines were used and the transaction was rolled back at the end, the changes—including the creation of the temporary table—were not written to the binary log.

The current fix restores the correct behavior in both of these cases. (Bug #53560)

References: This issue is a regression of: Bug #43929.

- **Replication:** When `CURRENT_USER()` or `CURRENT_USER` was used to supply the name and host of the affected user or of the definer in any of the statements `DROP USER`, `RENAME USER`, `GRANT`, `REVOKE`, and `ALTER EVENT`, the reference to `CURRENT_USER()` or `CURRENT_USER` was not expanded when written to the binary log. This resulted in `CURRENT_USER()` or `CURRENT_USER` being expanded to the user and host of the slave SQL thread on the slave, thus breaking replication. Now `CURRENT_USER()` and `CURRENT_USER` are expanded prior to being written to the binary log in such cases, so that the correct user and host are referenced on both the master and the slave. (Bug #48321)
- A signal-handler redefinition for `SIGUSR1` was removed. The redefinition could cause the server to encounter a kernel deadlock on Solaris when there are many active threads. Other POSIX platforms might also be affected. (Bug #54667)
- The `make_binary_distribution` target to `make` could fail on some platforms because the lines generated were too long for the shell. (Bug #54590)
- The server failed to disregard sort order for some zero-length tuples, leading to an assertion failure. (Bug #54459)
- The default value of `myisam_max_extra_sort_file_size` could be higher than the maximum accepted value, leading to warnings upon the server start. (Bug #54457)
- Inconsistent checking of the relationship between `SHOW` statements and `INFORMATION_SCHEMA` queries caused such queries to fail sometimes. (Bug #54422)
- If a session tried to drop a database containing a table opened with `HANDLER` in another session, any `DATABASE` statement (`CREATE`, `DROP`, `ALTER`) executed by that session produced a deadlock. (Bug #54360)
- Builds of the embedded `mysqld` failed due to a missing element of the `struct NET`. (Bug #53908, Bug #53912)

- The definition of the `MY_INIT` macro in `my_sys.h` included an extraneous semicolon, which could cause compilation failure. (Bug #53906)
- A client with automatic reconnection enabled saw the error message `Lost connection to MySQL server during query` if the connection was lost between the `mysql_stmt_prepare()` and `mysql_stmt_execute()` C API functions. However, `mysql_stmt_errno()` returned 0, not the corresponding error number 2013. (Bug #53899)
- Queries that used `MIN()` or `MAX()` on indexed columns could be optimized incorrectly. (Bug #53859)
- The `Lock_time` value in the slow query log was negative for stored routines. (Bug #53191)
- The results of some `ORDER BY ... DESC` queries were sorted incorrectly. (Bug #51431)
- `Index Merge` between three indexes could return incorrect results. (Bug #50389)
- The server could crash with an out of memory error when trying to parse a query that was too long to fit in memory. Now the parser rejects such queries with an `ER_OUT_OF_RESOURCES` error. (Bug #42064)
- `Sort-index_merge` for join tables other than the first table used excessive memory. (Bug #41660)
- Valgrind warnings in the `InnoDB compare_record()` function were corrected. (Bug #38999)
- `mysqld` could fail during execution when using SSL. (Bug #34236)
- The behavior of the RPM upgrade installation has changed. During an upgrade installation using the RPM packages, if the MySQL server is running when the upgrade occurs, the server is stopped, the upgrade occurs, and server is restarted. If the server is not already running when the RPM upgrade occurs, the server is not started at the end of the upgrade. The boot scripts for MySQL are installed in the appropriate directories in `/etc`, so the MySQL server will be restarted automatically at the next machine reboot. (Bug #27072)

Changes in MySQL 5.1.48 (2010-06-02)

- [InnoDB Plugin Notes](#)
- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

InnoDB Plugin Notes

- `InnoDB Plugin` has been upgraded to version 1.0.9. This version is considered of General Availability (GA) quality.

In this release, the `InnoDB Plugin` is included in source and binary distributions, except RHEL3, RHEL4, SuSE 9 (x86, x86_64, ia64), generic Linux RPM packages, and any builds produced with the `icc` compiler. It also does not work for FreeBSD 6 and HP-UX or for Linux on generic ia64.

Functionality Added or Changed

- The `Rows_examined` value in slow query log rows now is nonzero for `UPDATE` and `DELETE` statements that modify rows. (Bug #49756)

Bugs Fixed

- **Security Fix:** A security bug was fixed. (Bug #53907)
- **Security Fix:** A security bug was fixed. (Bug #52357)

- **Security Fix:** A security bug was fixed. (Bug #48157)
- **Important Change; Replication:** `MyISAM` transactions replicated to a transactional slave left the slave in an unstable condition. This was due to the fact that, when replicating from a nontransactional storage engine to a transactional engine with `autocommit` disabled, no `BEGIN` and `COMMIT` statements were written to the binary log; thus, on the slave, a never-ending transaction was started.

The fix for this issue includes enforcing `autocommit` mode on the slave by replicating all `autocommit=1` statements from the master. (Bug #29288)

- **InnoDB; Replication:** Reading from a table that used a self-logging storage engine and updating a table that used a transactional engine (such as `InnoDB`) generated changes that were written to the binary log using statement format which could make slaves diverge. However, when using mixed logging format, such changes should be written to the binary log using row format. (This issue did not occur when reading from tables using a self-logging engine and updating `MyISAM` tables, as this was already handled by checking for combinations of nontransactional and transactional engines.) Now such statements are classified as unsafe, and in mixed mode, cause a switch to row-based logging. (Bug #49019)
- **InnoDB:** The server could crash with a message `InnoDB: Assertion failure in thread nnnn`, typically during shutdown on a Windows system. (Bug #53947)
- **InnoDB:** Adding a unique key on multiple columns, where one of the columns is `NULL`, could mistakenly report duplicate key errors. (Bug #53290)
- **InnoDB:** Fixed a checksum error reported for compressed tables when the `--innodb_checksums` option is enabled. Although the message stated that the table was corrupted, the table is actually fine. (Bug #53248)
- **InnoDB:** Corrected the handling of the setting `innodb_change_buffering=default`. (The appropriate default value is different between MySQL 5.1 and 5.5.) (Bug #53165)
- **InnoDB:** Multi-statement execution could fail with an error about foreign key constraints. This problem could affect calls to `mysql_query()` and `mysql_real_query()`, and `CALL` statements that invoke stored procedures. (Bug #48024)
- **InnoDB:** If a crash occurs while creating an index using the InnoDB “Fast Index Creation” mechanism, the partially created index is dropped during the crash recovery processing when the database is restarted.
- **Partitioning:** `ALTER TABLE` statements that cause table partitions to be renamed or dropped (such as `ALTER TABLE ... ADD PARTITION`, `ALTER TABLE ... DROP PARTITION`, and `ALTER TABLE ... REORGANIZE PARTITION`) — when run concurrently with queries against the `INFORMATION_SCHEMA.PARTITIONS` table — could fail, cause the affected partitioned tables to become unusable, or both. This was due to the fact that the `INFORMATION_SCHEMA` database ignored the name lock imposed by the `ALTER TABLE` statement on the partitions affected. In particular, this led to problems with `InnoDB` tables, because `InnoDB` would accept the rename operation, but put it in a background queue, so that subsequent rename operations failed when `InnoDB` was unable to find the correct partition. Now, `INFORMATION_SCHEMA` honors name locks imposed by ongoing `ALTER TABLE` statements that cause partitions to be renamed or dropped. (Bug #50561)

References: See also: Bug #47343, Bug #45808.

- **Partitioning:** It was possible to execute a `CREATE TEMPORARY TABLE tmp LIKE pt` statement, where `pt` is a partitioned table, even though partitioned temporary tables are not permitted. This caused the server to crash. Now a check is performed to prevent such statements from being executed. (Bug #49477)

- **Partitioning:** When attempting to perform DDL on a partitioned table and the table's `.par` file could not be found, the server returned the inaccurate error message `Out of memory; restart server and try again (needed 2 bytes)`. Now in such cases, the server returns the error `Failed to initialize partitions from .par file`. (Bug #49161)
- **Replication:** In some cases, attempting to update a column with a value of an incompatible type resulted in a mismatch between master and slave because the column value was set to its implicit default value on the master (as expected), but the same column on the slave was set to `NULL`. (Bug #52868)
- **Replication:** When using a nontransactional table on the master with autocommit disabled, no `COMMIT` was recorded in the binary log following a statement affecting this table. If the slave's copy of the table used a transactional storage engine, the result on the slave was as though a transaction had been started, but never completed. (Bug #49522)

References: See also: Bug #29288.

- Valgrind warnings resulting from passing incomplete `DATETIME` values to the `TIMESTAMP()` function were corrected. (Bug #53942)
- `UPDATE` on an `InnoDB` table modifying the same index that was used to satisfy the `WHERE` condition could trigger a debug assertion under some circumstances. (Bug #53830)
- MySQL incorrectly processed `ALTER DATABASE `#mysql50#special` UPGRADE DATA DIRECTORY NAME` where `special` was `.`, `..`, or a sequence starting with `./` or `../`. It used the server data directory (which contains other regular databases) as the database directory. (Bug #53804, CVE-2010-2008)
- `InnoDB` crashed when replacing duplicates in a table after a fast `ALTER TABLE` added a unique index. (Bug #53592)
- For `InnoDB` tables, the error handler for a fast `CREATE INDEX` did not reset the error state of the transaction before attempting to undo a failed operation, resulting in a crash. (Bug #53591)
- For single-table `DELETE` statements that used quick select and index scan simultaneously caused a server crash or assertion failure. (Bug #53450)
- Incorrect results could be returned for `LEFT JOIN` of `InnoDB` tables with an impossible `WHERE` condition. (Bug #53334)
- Setting the `innodb_change_buffering` system variable to `DEFAULT` produced an incorrect result. (Bug #53165)
- `mysqldump` and `SELECT ... INTO OUTFILE` truncated long `BLOB` and `TEXT` values to 766 bytes. (Bug #53088)
- In the debug version of the server, the `FreeState()` function could in some circumstances be called twice, leading to an assertion failure. (Bug #52884)
- Aggregate functions could incorrectly return `NULL` in outer join queries. (Bug #52051)
- For outer joins, the optimizer could fail to properly calculate table dependencies. (Bug #52005)
- The Loose Index Scan optimization method assumed that it could depend on the partitioning engine to maintain interval endpoint information, as if it were a storage engine. (Bug #50939)
- Calculation of intervals for Event Scheduler events was not portable. (Bug #50087)

- Selecting from `INFORMATION_SCHEMA.ROUTINES` or `INFORMATION_SCHEMA.PARAMETERS` resulted in a memory leak. (Bug #48729)
- On Intel x86 machines, the optimizer could choose different execution plans for a query depending on the compiler version and optimization flags used to build the server binary. (Bug #48537)
- When the transaction isolation level was `REPEATABLE READ` and binary logging used statement or mixed format, `SELECT` statements with subqueries referencing `InnoDB` tables unnecessarily acquired shared locks on rows in these tables. (Bug #46947)
- Using an initial command with `mysql_options(..., MYSQL_INIT_COMMAND, ...)` that generated multiple result sets (such as a stored procedure or a multi-statement command) left the connection unusable. (Bug #42373)

Changes in MySQL 5.1.47 (2010-05-06)

- [InnoDB Plugin Notes](#)
- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

InnoDB Plugin Notes

- `InnoDB Plugin` has been upgraded to version 1.0.8. This version is considered of General Availability (GA) quality.

In this release, the `InnoDB Plugin` is included in source and binary distributions, except RHEL3, RHEL4, SuSE 9 (x86, x86_64, ia64), and generic Linux RPM packages. It also does not work for FreeBSD 6 and HP-UX or for Linux on generic ia64.

Functionality Added or Changed

- **InnoDB:** `InnoDB` stores redo log records in a hash table during recovery. On 64-bit systems, this hash table was 1/8 of the buffer pool size. To reduce memory usage, the dimension of the hash table was reduced to 1/64 of the buffer pool size (or 1/128 on 32-bit systems). (Bug #53122)

Bugs Fixed

- **Security Fix:** The server failed to check the table name argument of a `COM_FIELD_LIST` command packet for validity and compliance to acceptable table name standards. This could be exploited to bypass almost all forms of checks for privileges and table-level grants by providing a specially crafted table name argument to `COM_FIELD_LIST`.

In MySQL 5.0 and above, this permitted an authenticated user with `SELECT` privileges on one table to obtain the field definitions of any table in all other databases and potentially of other MySQL instances accessible from the server's file system.

Additionally, for MySQL version 5.1 and above, an authenticated user with `DELETE` or `SELECT` privileges on one table could delete or read content from any other table in all databases on this server, and potentially of other MySQL instances accessible from the server's file system. (Bug #53371, CVE-2010-1848)

- **Security Fix:** The server was susceptible to a buffer-overflow attack due to a failure to perform bounds checking on the table name argument of a `COM_FIELD_LIST` command packet. By sending long data for the table name, a buffer is overflowed, which could be exploited by an authenticated user to inject malicious code. (Bug #53237, CVE-2010-1850)

- **Security Fix:** The server could be tricked into reading packets indefinitely if it received a packet larger than the maximum size of one packet. (Bug #50974, CVE-2010-1849)
- **Performance; InnoDB:** Deadlock detection could be a bottleneck in InnoDB processing, if many transactions attempted to update the same row simultaneously. The algorithm has been improved to enhance performance and scalability, in the InnoDB Plugin for MySQL 5.1, and in InnoDB 1.1 for MySQL 5.5. (Bug #49047)
- **Important Change; Replication:** When invoked, `CHANGE MASTER TO` and `SET GLOBAL sql_slave_skip_counter` now cause information to be written to the error log about the slave's state prior to execution of the statement. For `CHANGE MASTER TO`, this information includes the previous values of `MASTER_HOST`, `MASTER_PORT`, `MASTER_LOG_FILE`, and `MASTER_LOG_POS`. For `SET GLOBAL sql_slave_skip_counter`, this information includes the previous values of `RELAY_LOG_FILE`, `RELAY_LOG_POS`, and `sql_slave_skip_counter`. (Bug #43406, Bug #43407)
- **InnoDB:** When reporting a foreign key constraint violation during `INSERT`, InnoDB could display uninitialized data for the `DB_TRX_ID` and `DB_ROLL_PTR` system columns. (Bug #53202)
- **InnoDB:** The values of `innodb_buffer_pool_pages_total` and `innodb_buffer_pool_pages_misc` in the `information_schema.global_status` table could be computed incorrectly. (Bug #52983)
- **InnoDB:** InnoDB page splitting could enter an infinite loop for compressed tables. (Bug #52964)
- **InnoDB:** An overly strict assertion could fail during the purge of delete-marked records in `DYNAMIC` or `COMPRESSED` InnoDB tables that contain column prefix indexes. (Bug #52746)
- **InnoDB:** InnoDB attempted to choose off-page storage without ensuring that there was an “off-page storage” flag in the record header. To correct this, in `DYNAMIC` and `COMPRESSED` formats, InnoDB stores locally any non-BLOB columns having a maximum length not exceeding 256 bytes. This is because there is no room for the “external storage” flag when the maximum length is 255 bytes or less. This restriction trivially holds in `REDUNDANT` and `COMPACT` formats, because there InnoDB always stores locally columns having a length up to `local_len` = 788 bytes. (Bug #52745)
- **InnoDB:** Connections waiting for an InnoDB row lock ignored `KILL` until the row lock wait ended. Now, `KILL` during lock wait results in “query interrupted” instead of “lock wait timeout exceeded”. The corresponding transaction is rolled back. (Bug #51920)
- **InnoDB:** InnoDB Plugin checks to see whether a row could possibly exceed the maximum size if all columns are fully used. This produced `Row size too large` errors for some tables that could be created with the built-in InnoDB. Now the check is only done when `innodb_strict_mode` is enabled or if the table is dynamic or compressed. (Bug #50495)
- **InnoDB:** A mismatch between index information maintained within the `.frm` files and the corresponding information in the InnoDB system tablespace could produce this error: `[ERROR] Index index of table has n columns unique inside InnoDB, but MySQL is asking statistics for m columns. Have you mixed up .frm files from different installations?` (Bug #44571)
- **Replication:** The failure of a `REVOKE` statement was logged with the wrong error code, causing replication slaves to stop even when the failure was expected on the master. (Bug #51987)
- Certain path names passed to `LOAD_FILE()` could cause a server crash. (Bug #53417)
- Semi-consistent read was implemented for InnoDB to address Bug #3300. Semi-consistent reads do not block when a nonmatching record is already locked by some other transaction. If the record is not locked, a lock is acquired, but is released if the record does not match the `WHERE` condition. However,

semi-consistent read was attempted even for `UPDATE` statements having a `WHERE` condition of the form `pk_coll=constant1, ..., pk_colN=constantN`. Some code failed that was designed with the assumption that semi-consistent read would be only attempted on table scans. (Bug #52663)

References: See also: Bug #3300.

- Setting `@@GLOBAL.debug` to an empty string failed to clear the current debug settings. (Bug #52629)
- A memory leak occurred due to missing deallocation of the `comparators` array (a member of the `Arg_comparator` class). (Bug #52124)
- For debug builds, creating a view containing a subquery that might require collation adjustment caused an assertion to be raised. For example, this could occur if some items had different collations but the result collation could be adjusted to the one of them. (Bug #52120)
- Locking involving the `LOCK_plugin`, `LOCK_global_system_variables`, and `LOCK_status` mutexes could deadlock. (Bug #51591)
- `InnoDB` fast index creation could incorrectly use a table copy in some cases. (Bug #50946)
- A syntactically invalid trigger could cause the server to crash when trying to list triggers. (Bug #50755)
- Setting `--secure-file-priv` to the empty string left the value unaffected. (Bug #50373)
- In MySQL 5.1, `READ COMMITTED` was changed to use less locking due to the availability of row-based binary logging (see the Note under `READ COMMITTED` at [SET TRANSACTION Syntax](#)). However, `READ UNCOMMITTED` did not have the same change, so it was using more locks than the higher isolation level, which is unexpected. This was changed so that `READ UNCOMMITTED` now also uses the lesser amount of locking and has the same restrictions for binary logging. (Bug #48607)
- `EXPLAIN` could cause a server crash for some queries with subqueries. (Bug #48419)
- On Windows, the server failed to find a description for Event ID 100. (Bug #48042)
- For updates to `InnoDB` tables, `TIMESTAMP` columns could be updated even when no values actually changed. (Bug #47453)
- `mysqld_safe` did not always pass `--open-files-limit` through to `mysqld`. `mysqld_safe` did not treat dashes and underscores as equivalent in option names. (Bug #47095)
- If the server is started with `--skip-grant-tables`, plugin loading and unloading should be prohibited, but the server failed to reject `INSTALL PLUGIN` and `UNINSTALL PLUGIN` statements. (Bug #46261)
- `InnoDB` could fail to create a unique index on `NULL` columns. (Bug #41904)
- Storage engine plugins on Windows could've been built with a definition of `time_t` which was different from the server expectations. The difference could cause affected plugins to crash. In addition, the use of the `time_t` type in the storage engine API layer has been enforced. (Bug #39802, Bug #40092)
- When using `UNINSTALL PLUGIN` to remove a loaded plugin, open tables and connections caused `mysqld` to hang until the open connections had been closed. (Bug #39053)
- 1) In rare cases, if a thread was interrupted during a `FLUSH PRIVILEGES` operation, a debug assertion occurred later due to improper diagnostics area setup. 2) A `KILL` operation could cause a console error message referring to a diagnostic area state without first ensuring that the state existed. (Bug #33982)

Changes in MySQL Enterprise 5.1.46sp1 [QSP] (2010-06-23)

This is a *Service Pack* release of the MySQL Enterprise Server 5.1.

**Important**

If you intend to use the plugin version of [InnoDB](#), we recommend that you use MySQL 5.1.48 or later instead of 5.1.46sp1. This is because 5.1.46sp1 contains the first production-ready version and the later version has fixes for some of the bugs found during more widespread production use.

Bugs Fixed

- **Security Fix:** The server failed to check the table name argument of a [COM_FIELD_LIST](#) command packet for validity and compliance to acceptable table name standards. This could be exploited to bypass almost all forms of checks for privileges and table-level grants by providing a specially crafted table name argument to [COM_FIELD_LIST](#).

In MySQL 5.0 and above, this permitted an authenticated user with [SELECT](#) privileges on one table to obtain the field definitions of any table in all other databases and potentially of other MySQL instances accessible from the server's file system.

Additionally, for MySQL version 5.1 and above, an authenticated user with [DELETE](#) or [SELECT](#) privileges on one table could delete or read content from any other table in all databases on this server, and potentially of other MySQL instances accessible from the server's file system. (Bug #53371, CVE-2010-1848)

- **Security Fix:** The server was susceptible to a buffer-overflow attack due to a failure to perform bounds checking on the table name argument of a [COM_FIELD_LIST](#) command packet. By sending long data for the table name, a buffer is overflowed, which could be exploited by an authenticated user to inject malicious code. (Bug #53237, CVE-2010-1850)
- **Security Fix:** The server could be tricked into reading packets indefinitely if it received a packet larger than the maximum size of one packet. (Bug #50974, CVE-2010-1849)
- **InnoDB:** [InnoDB](#) page splitting could enter an infinite loop for compressed tables. (Bug #52964)
- **InnoDB:** [InnoDB](#) attempted to choose off-page storage without ensuring that there was an “off-page storage” flag in the record header. To correct this, in [DYNAMIC](#) and [COMPRESSED](#) formats, [InnoDB](#) stores locally any non-[BLOB](#) columns having a maximum length not exceeding 256 bytes. This is because there is no room for the “external storage” flag when the maximum length is 255 bytes or less. This restriction trivially holds in [REDUNDANT](#) and [COMPACT](#) formats, because there [InnoDB](#) always stores locally columns having a length up to `local_len` = 788 bytes. (Bug #52745)
- MySQL incorrectly processed `ALTER DATABASE `#mysql50#special` UPGRADE DATA DIRECTORY NAME` where *special* was `.`, `...`, or a sequence starting with `./` or `../`. It used the server data directory (which contains other regular databases) as the database directory. (Bug #53804, CVE-2010-2008)
- A syntactically invalid trigger could cause the server to crash when trying to list triggers. (Bug #50755)
- Selecting from [INFORMATION_SCHEMA.ROUTINES](#) or [INFORMATION_SCHEMA.PARAMETERS](#) resulted in a memory leak. (Bug #48729)
- [EXPLAIN](#) could cause a server crash for some queries with subqueries. (Bug #48419)

Changes in MySQL 5.1.46 (2010-04-06)

**Important**

If you intend to use the plugin version of [InnoDB](#), we recommend that you use MySQL 5.1.48 or later instead of 5.1.46sp1. This is because 5.1.46 contains the

first production-ready version and the later version has fixes for some of the bugs found during more widespread production use.

- [InnoDB Plugin Notes](#)
- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

InnoDB Plugin Notes

- [InnoDB Plugin](#) has been upgraded to version 1.0.7. This version is considered of General Availability (GA) quality. [InnoDB Storage Engine Change History](#), may contain information in addition to those changes reported here.

In this release, the [InnoDB Plugin](#) is included in source and binary distributions, except RHEL3, RHEL4, SuSE 9 (x86, x86_64, ia64), and generic Linux RPM packages. It also does not work for FreeBSD 6 and HP-UX or for Linux on generic ia64.

Functionality Added or Changed

- A new system variable, [skip_name_resolve](#), is set from the value of the `--skip-name-resolve` server option. This provides a way to determine at runtime whether the server uses name resolution for client connections. (Bug #37168)

Bugs Fixed

- **Security Fix:** Privilege checking for [UNINSTALL PLUGIN](#) was incorrect. (Bug #51770, CVE-2010-1621)
- **Performance; InnoDB:** The redo scan during [InnoDB](#) recovery used excessive CPU. The efficiency of this scan was improved for [InnoDB Plugin](#), significantly speeding up crash recovery. (Bug #49535, Bug #29847)
- **Performance; InnoDB:** [InnoDB Plugin](#) page-freeing operations were made faster for compressed blocks, speeding up [ALTER TABLE](#), [DROP TABLE](#), and other operations on compressed tables that free compressed blocks. One symptom of the older behavior could be 100% CPU use during these operations. (Bug #35077)
- **Performance:** While looking for the shortest index for a covering index scan, the optimizer did not consider the full row length for a clustered primary key, as in [InnoDB](#). Secondary covering indexes are now preferred, making full table scans less likely. (Bug #39653)

References: See also: Bug #55656.

- **Important Change:** When using fast [ALTER TABLE](#), different internal ordering of indexes in the MySQL optimizer and the [InnoDB](#) storage engine could cause error messages about possibly mixed up [.frm](#) files and incorrect index use. (Bug #47622)
- **InnoDB; Replication:** [TRUNCATE TABLE](#) performed on a temporary table using the [InnoDB](#) storage engine was logged even when using row-based mode. (Bug #51251)
- **InnoDB; Replication:** Column length information generated by [InnoDB](#) did not match that generated by [MyISAM](#), which caused invalid metadata to be written to the binary log when trying to replicate [BIT](#) columns. (Bug #49618)
- **InnoDB:** For [InnoDB Plugin](#), bit fields were causing problems with concurrency on SMP systems because of word-packing issues. (Bug #52360)
- **InnoDB:** Fixed a performance issue on Windows systems that affected the [InnoDB Plugin](#), by turning off atomic instructions. (Bug #52102)

- **InnoDB:** The AIX implementation of `readdir_r()` caused InnoDB errors. (Bug #50691)
- **Partitioning:** Partition pruning on `RANGE` partitioned tables did not always work correctly; the last partition was not excluded if the range was beyond it (when not using `MAXVALUE`). Now the last partition is not included if the partitioning function value is not within the range. (Bug #51830)
- **Partitioning:** The `insert_id` server system variable was not reset following an insert that failed on a partitioned `MyISAM` table having an `AUTO_INCREMENT` column. (Bug #50392)
- **Partitioning:** Foreign keys are not supported on partitioned tables. However, it was possible using an `ALTER TABLE` statement to set a foreign key on a partitioned table; it was also possible to partition a table with a single foreign key. (Bug #50104)
- **Partitioning:** `GROUP BY` queries performed poorly for some partitioned tables. This was due to the block size not being set for partitioned tables, thus the keys per block was not correct, which could cause such queries to be optimized incorrectly. (Bug #48229)

References: See also: Bug #37252.

- **Partitioning:** `REPAIR TABLE` failed for partitioned `ARCHIVE` tables. (Bug #46565)
- **Replication:** When using temporary tables, the binary log needs to insert a pseudo-thread ID for threads that are using temporary tables, each time a switch happens between two threads, both of which are using temporary tables. However, if a thread issued a failing statement before exit, its ID was not recorded in the binary log, and this in turn caused the ID for the next thread that tried to do something with a temporary table not to be logged as well. Subsequent replays of the binary log failed with the error `Table ... doesn't exist`. (Bug #51226)

References: This issue is a regression of: Bug #35583.

- **Replication:** If the master was using `sql_mode='TRADITIONAL'`, duplicate key errors were not sent to the slave, which received 0 rather than the expected error code. This caused replication to fail even when such an error was expected. (Bug #51055)
- **Replication:** When run with the `--database` option, `mysqlbinlog` printed `ROLLBACK` statements but did not print any corresponding `SAVEPOINT` statements. (Bug #50407)
- **Replication:** When a `CREATE EVENT` statement was followed by an additional statement and the statements were executed together as a single statement, the `CREATE EVENT` statement was padded with “garbage” characters when written to the binary log. This led to a syntax error when the event was read from the log. (Bug #50095)
- **Replication:** The value of `Slave_IO_running` in the output of `SHOW SLAVE STATUS` did not distinguish between all 3 possible states of the slave I/O thread (not running; running but not connected; connected). Now the value `Connecting` (rather than `No`) is shown when the slave I/O thread is running but the slave is not connected to a replication master.

The server system variable `Slave_running` also reflects this change, and is now consistent with what is shown for `Slave_IO_running`. (Bug #30703, Bug #41613, Bug #51089)

- `EXPLAIN EXTENDED` crashed trying to resolve references to freed temporary table columns for `GROUP_CONCAT()` `ORDER BY` arguments. (Bug #52397)
- The optimizer could attempt to evaluate the `WHERE` clause before any rows had been read, resulting in a server crash. (Bug #52177)
- For LDML-defined collations, some data structures were not initialized properly to enable `UPPER()` and `LOWER()` to work correctly. (Bug #51976)

- On Windows, `LOAD_FILE()` could cause a crash for some pathnames. (Bug #51893)
- Invalid memory reads occurred for `HANDLER ... READ NEXT` after a failed `HANDLER ... READ FIRST`. (Bug #51877)
- After `TRUNCATE TABLE` of a `MyISAM` table, subsequent queries could crash the server if `myisam_use_mmap` was enabled. (Bug #51868)
- If `myisam_sort_buffer_size` was set to a small value, table repair for `MyISAM` tables with `FULLTEXT` indexes could crash the server. (Bug #51866)
- In `LOAD DATA INFILE`, using a `SET` clause to set a column equal to itself caused a server crash. (Bug #51850)
- A problem with equality propagation optimization for prepared statements and stored procedures caused a server crash upon re-execution of the prepared statement or stored procedure. (Bug #51650)

References: See also: Bug #8115, Bug #8849.

- The optimizer performed an incorrect join type when `COALESCE()` appeared within an `IN()` operation. (Bug #51598)
- The server crashed when the optimizer attempted to determine constant tables but a table storage engine did not support exact record count. (Bug #51494)
- A unique index on a column prefix could not be upgraded to a primary index even if there was no primary index already defined. (Bug #51378)
- The server could crash populating the `INFORMATION_SCHEMA.PROCESSLIST` table due to lack of mutex protection. (Bug #51377)
- Use of `HANDLER` statements with tables that had spatial indexes caused a server crash. (Bug #51357)
- With an XA transaction active, `SET autocommit = 1` could cause side effects such as memory corruption or a server crash. (Bug #51342)
- Following a bulk insert into a `MyISAM` table, if `MyISAM` failed to build indexes using repair by sort, data file corruption could occur. (Bug #51307)
- `CHECKSUM TABLE` could compute the checksum for `BIT` columns incorrectly. (Bug #51304)
- A `HAVING` clause on a joined table in some cases failed to eliminate rows which should have been excluded from the result set. (Bug #51242)
- The type inference used for view columns caused some columns in views to be handled as the wrong type, as compared to the same columns in base tables. `DATE` columns in base tables were treated as `TIME` columns in views, and base table `TIME` columns as view `DATETIME` columns. (Bug #50918)
- The `YEAR` values `2000` and `0000` could be treated as equal. (Bug #49910)
- Performing a single in-place `ALTER TABLE` containing `ADD INDEX` and `DROP INDEX` options that used the same index name could result in a corrupt table definition file. Now such `ALTER TABLE` statements are no longer performed in place. (Bug #49838)
- `mysql_upgrade` did not detect when `CSV` log tables incorrectly contained columns that could be `NULL`. Now these columns are altered to be `NOT NULL`. (Bug #49823)
- `InnoDB` returned an error when inserting a negative value into an auto-increment column. (Bug #49497)

- `InnoDB` did not reset table `AUTO_INCREMENT` values to the last used values after a server restart. (Bug #49032)
- If a stored function contained a `RETURN` statement with an `ENUM` value in the `ucs2` character set, `SHOW CREATE FUNCTION` and `SELECT DTD_IDENTIFIER FROM INFORMATION_SCHEMA.ROUTINES` returned incorrect values. (Bug #48766)
- A trigger could change the behavior of assigning `NULL` to a `NOT NULL` column. (Bug #48525)
- The server crashed when it could not determine the best execution plan for queries involving outer joins with nondeterministic `ON` clauses such as the ones containing the `RAND()` function, a user-defined function, or a `NOT DETERMINISTIC` stored function. (Bug #48483)
- The `MERGE` engine failed to open a child table from a different database if the child table or database name contained characters that were subject to table name to file name encoding.

Further, the `MERGE` engine did not properly open a child table from the same database if the child table name contained characters such as `'/'`, or `'#'`. (Bug #48265)

- A query that read from a derived table (of the form `SELECT ... FROM (SELECT ...)`) produced incorrect results when the following conditions were present:
 - The table subquery contained a derived query (`(SELECT ...) AS column`).
 - The derived query could potentially produce zero rows or a single `NULL` (that is, no rows matched, or the query used an aggregate function such as `SUM()` running over zero rows).
 - The table subquery joined at least two tables.
 - The join condition involved an index.

(Bug #47904)

- The optimization to read `MIN()` or `MAX()` values from an index did not properly handle comparisons with `NULL` values. This could produce incorrect results for `MIN()` or `MAX()` when the `WHERE` clause tested a `NOT NULL` column for `NULL`. (Bug #47762)
- Killing a query during the optimization phase of a subquery could cause a server crash. (Bug #47761)
- The query shown by `EXPLAIN EXTENDED` plus `SHOW WARNINGS` could produce results different from the original query. (Bug #47669)
- Renaming a column of an `InnoDB` table caused the server to go out of sync with the `InnoDB` data dictionary. To avoid this issue, renaming a column uses the older technique of copying all the table data rather than updating the table in-place. (Bug #47621)
- `MyISAM` could write uninitialized data to new index pages. Now zeros are written to unused bytes in the pages. (Bug #47598)
- Setting `myisam_repair_threads` larger than 1 could result in the cardinality for all indexes of a `MyISAM` table being set to 1 after parallel index repair. (Bug #47444)
- In debug builds, if the listed columns in the view definition of the table used in an `INSERT ... SELECT` statement mismatched, an assertion was raised in the query cache invalidation code following the failing statement. (Bug #46615)
- For a query that selected from a view and used an alias for the view, the metadata used the alias name rather than the view name in the `MYSQL_FIELD.table` member. (Bug #41788)

- `mysql_upgrade` did not create temporary files properly. (Bug #41057)
- It was possible for `DROP TABLE` of one `MyISAM` table to remove the data and index files of a different `MyISAM` table. (Bug #40980)
- If the arguments to a `CONCAT()` call included a local routine variable, selecting the return value into a user variable could produce an incorrect result. (Bug #40625)
- `SHOW CREATE VIEW` returned invalid SQL if the definition contained a `SELECT 'string'` statement where the `string` was longer than the maximum length of a column name, due to the fact that this text was also used as an alias (in the `AS` clause).

Because not all names retrieved from arbitrary `SELECT` statements can be used as view column names due to length and format restrictions, the server now checks the conformity of automatically generated column names and rewrites according to a predefined format any names that are not acceptable as view column names before storing the final view definition on disk.

In such cases, the name is now rewritten as `Name_exp_pos`, where `pos` is the position of the column. To avoid this conversion scheme, define explicit, valid names for view columns using the `column_list` clause of the `CREATE VIEW` statement.

As part of this fix, aliases are now generated only for top-level statements. (Bug #40277)

- `mysqlbinlog` option-processing code had a memory leak. (Bug #38468)
- The test for `readline` during configuration failed when trying to build MySQL in a directory other than the source tree root. (Bug #35250)
- A query on a `FEDERATED` table in which the data was ordered by a `TEXT` column returned incorrect results. For example, a query such as the following produced incorrect results if column `column1` was a `TEXT` column:

```
SELECT * FROM table1 ORDER BY column1;
```

(Bug #32426)

Changes in MySQL 5.1.45 (2010-03-01)

- [InnoDB Plugin Notes](#)
- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

InnoDB Plugin Notes

- This release includes `InnoDB Plugin` 1.0.6. This version is considered of Release Candidate (RC) quality.

In this release, the `InnoDB Plugin` is included in source and binary distributions, except RHEL3, RHEL4, SuSE 9 (x86, x86_64, ia64), and generic Linux RPM packages. It also does not work for FreeBSD 6 and HP-UX or for Linux on generic ia64.

Functionality Added or Changed

- `mysqltest` has a new `--max-connections` option to set a higher number of maximum permitted server connections than the default 128. This option can also be passed using `mysql-test-run.pl`. (Bug #51135)

- `mysql-test-run.pl` has a new `--portbase` option and a corresponding `MTR_PORT_BASE` environment variable for setting the port range, as an alternative to the existing `--build-thread` option. (Bug #50182)
- `mysql-test-run.pl` now has a `--gprof` option that runs the server through the `gprof` profiler, much the same way the currently supported `--gcov` option runs it through `gcov`. (Bug #49345)
- `mysqltest` now has a `lowercase_result` command that converts the output of the next statement to lowercase. This is useful for test cases where the lettercase may vary between platforms. (Bug #48863)
- `mysqltest` now has a `remove_files_wildcard` command that removes files matching a pattern from a directory. (Bug #39774)

Bugs Fixed

- **InnoDB:** `SHOW INNODB STATUS` could display incorrect information about deadlocks, when the deadlock detection routine stops early (to avoid excessive CPU usage). (Bug #49001)
- **Partitioning:** Attempting to drop a partitioned table from one connection while waiting for the completion of an `ALTER TABLE` that had been issued from a different connection, and that changed the storage engine used by the table, could cause the server to crash. (Bug #42438)
- **Replication:** Adding an index to a table on the master caused the slave to stop logging slow queries to the slow query log. (Bug #50620)
- **Replication:** Queries written to the slow query log on the master were not written to the slow query log on the slave. (Bug #23300)

References: See also: Bug #48632.

- `mysqld_multi` failed due to a syntax error in the script. (Bug #51468)
- Referring to a subquery result in a `HAVING` clause could produce incorrect results. (Bug #50995)
- The optimizer normally prefers use of `filesort` plus the join cache to a full index scan. But this combination was used even if the index is clustered, in which case, the clustered index scan can be faster. (Bug #50843)
- For debug builds, `SHOW BINARY LOGS` raised an assertion if binary logging was not enabled. (Bug #50780)
- The server did not recognize that the stored procedure cache became invalid if a view was created or modified within a procedure, resulting in a crash. (Bug #50624)
- Incorrect handling of `BIT` columns in temporary tables could lead to spurious duplicate-key errors. (Bug #50591)
- The second or subsequent invocation of a stored procedure containing `DROP TRIGGER` could cause a server crash. (Bug #50423)
- Full-text queries that used the truncation operator (`*`) could enter an infinite loop. (Bug #50351)
- For debug builds, an assertion was incorrectly raised in the optimizer when matching `ORDER BY` expressions. (Bug #50335)
- Queries optimized with `GROUP_MIN_MAX` did not clean up `KEYREAD` optimizations properly, causing subsequent queries to return incomplete rows. (Bug #49902)
- For dynamic format `MyISAM` tables containing `LONGTEXT` columns, a bulk `INSERT ... ON DUPLICATE KEY UPDATE` or bulk `REPLACE` could cause corruption. (Bug #49628)

- For debug builds, with `sql_safe_updates` enabled, a multiple-table `UPDATE` with the `IGNORE` modifier could raise an assertion. (Bug #49534)
- `EXPLAIN EXTENDED` crashed trying to print column names for a subquery in the `FROM` clause when the table had gone out of scope. (Bug #49487)
- For `InnoDB` tables, the test for using an index for `ORDER BY` sorting did not distinguish between primary keys and secondary indexes and expected primary key values to be concatenated to index values the way they are to secondary key values. (Bug #49324)
- `mysqltest` no longer permits you to execute an SQL statement on a connection after doing a `send` command, unless you do a `reap` first. This was previously accepted but could produce unpredictable results. (Bug #49269)
- For debug builds on Windows, warnings about incorrect use of debugging directives were written to the error log. The directives were rewritten to eliminate these messages. (Bug #49025)
- An `.ARZ` file missing from the database directory caused the server to crash. (Bug #48757)
- Running `SHOW CREATE TABLE` on a view `v1` that contained a function which accessed another view `v2` could trigger a infinite loop if the view referenced within the function (`v2`) caused a warning to be raised while being opened. (Bug #48449)
- Invalid memory reads could occur following a query that referenced a `MyISAM` table multiple times with a write lock. (Bug #48438)
- For debug builds, creating a view containing a row constructor raised an assertion. (Bug #48294)
- Slow `CALL` statements were not always logged to the slow query log because execution time for multiple-statement stored procedures was assessed incorrectly. (Bug #47905)
- For debug builds, killing a `SELECT` retrieving from a view that was processing a function raised an assertion. (Bug #47736)
- Failure to open a view with a nonexistent `DEFINER` was improperly handled and the server crashed later attempting to lock the view. (Bug #47734)
- If `EXPLAIN` encountered an error in the query, a memory leak occurred. (Bug #45989)
- Grouping by a subquery in a query with a `DISTINCT` aggregate function led to incorrect and unordered grouping values. (Bug #45640)
- Propagation of a large unsigned numeric constant in `WHERE` expressions could lead to incorrect results. This also affected `EXPLAIN EXTENDED`, which printed incorrect numeric constants in such transformed `WHERE` expressions. (Bug #45360)
- Valgrind warnings about uninitialized variables in optimizer code were corrected. (Bug #45195)
- `flush_cache_records()` did not correctly check for errors that should cause statement execution to stop, leading to a server crash. (Bug #39022)
- `InnoDB` logged an error repeatedly trying to load a page into the buffer pool, filling the error log and using excessive disk space. Now the number of attempts is limited to 100, after which the operation aborts with a message. (Bug #38901)
- When building MySQL when using a different target directory (for example using the `VPATH` environment variable), the build of the embedded `readline` component failed. (Bug #35250)

- [INSERT INTO ... VALUES\(DEFAULT\)](#) failed to insert the correct value for [ENUM](#) columns. For [MyISAM](#) tables, an empty value was inserted. For [CSV](#) tables, the table became corrupt. (Bug #33717)

Changes in MySQL 5.1.44 (2010-02-04)

- [InnoDB Plugin Notes](#)
- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

InnoDB Plugin Notes

- This release includes [InnoDB Plugin](#) 1.0.6. This version is considered of Release Candidate (RC) quality.

In this release, the [InnoDB Plugin](#) is included in source and binary distributions, except RHEL3, RHEL4, SuSE 9 (x86, x86_64, ia64), and generic Linux RPM packages. It also does not work for FreeBSD 6 and HP-UX or for Linux on generic ia64.

Functionality Added or Changed

- **Replication:** Introduced the [binlog_direct_non_transactional_updates](#) system variable. Enabling this variable causes updates using the statement-based logging format to tables using nontransactional engines to be written directly to the binary log, rather than to the transaction cache.

Before enabling this variable, be certain that you have no dependencies between transactional and nontransactional tables. A statement that both selects from an [InnoDB](#) table and inserts into a [MyISAM](#) table is an example of such a dependency. For more information, see [Binary Log Options and Variables](#). (Bug #46364)

References: See also: Bug #28976, Bug #40116.

Bugs Fixed

- **Performance:** The method for comparing [INFORMATION_SCHEMA](#) names and database names was nonoptimal and an improvement was made: When the database name length is already known, a length check is made first and content comparison skipped if the lengths are unequal. (Bug #49501)
- **Performance:** The [MD5\(\)](#) and [SHA1\(\)](#) functions had excessive overhead for short strings. (Bug #49491, Bug #11757443, Bug #60227, Bug #14134662)
- **InnoDB; Partitioning:** When an [ALTER TABLE ... REORGANIZE PARTITION](#) statement on an [InnoDB](#) table failed due to [innodb_lock_wait_timeout](#) expiring while waiting for a lock, [InnoDB](#) did not clean up any temporary files or tables which it had created. Attempting to reissue the [ALTER TABLE](#) statement following the timeout could lead to storage engine errors, or possibly a crash of the server. (Bug #47343)
- **InnoDB:** Creating or dropping a table with 1023 transactions active caused an assertion failure. (Bug #49238)
- **InnoDB:** If [innodb_force_recovery](#) was set to 4 or higher, the server could crash when opening an [InnoDB](#) table containing an auto-increment column. MySQL versions 5.1.31 and later were affected. (Bug #46193)
- **Replication:** In some cases, inserting into a table with many columns could cause the binary log to become corrupted. (Bug #50018)

References: See also: Bug #42749.

- **Replication:** When using row-based replication, setting a `BIT` or `CHAR` column of a `MyISAM` table to `NULL`, then trying to delete from the table, caused the slave to fail with the error `Can't find record in table`. (Bug #49481, Bug #49482)
- **Replication:** When logging in row-based mode, DDL statements are actually logged as statements; however, statements that affected temporary tables and followed DDL statements failed to reset the binary log format to `ROW`, with the result that these statements were logged using the statement-based format. Now the state of `binlog_format` is restored after a DDL statement has been written to the binary log. (Bug #49132)
- **Replication:** When using row-based logging, the statement `CREATE TABLE t IF NOT EXIST ... SELECT` was logged as `CREATE TEMPORARY TABLE t IF NOT EXIST ... SELECT` when `t` already existed as a temporary table. This was caused by the fact that the temporary table was opened and the results of the `SELECT` were inserted into it when a temporary table existed and had the same name.

Now, when this statement is executed, `t` is created as a base table, the results of the `SELECT` are inserted into it—even if there already exists a temporary table having the same name—and the statement is logged correctly. (Bug #47418)

References: See also: Bug #47442.

- **Replication:** Due to a change in the size of event representations in the binary log, when replicating from a MySQL 4.1 master to a slave running MySQL 5.0.60 or later, the `START SLAVE UNTIL` statement did not function correctly, stopping at the wrong position in the log. Now the slave detects that the master is using the older version of the binary log format, and corrects for the difference in event size, so that the slave stops in the correct position. (Bug #47142)
- The SSL certificates in the test suite were about to expire. They have been updated with expiration dates in the year 2015. (Bug #50642)
- The `printstack` function does not exist on Solaris 8 or earlier, which led to a compilation failure. (Bug #50409)
- A user could see tables in `INFORMATION_SCHEMA.TABLES` without appropriate privileges for them. (Bug #50276)
- Debug output for join structures was garbled. (Bug #50271)
- The `filesort` sorting method applied to a `CHAR(0)` column could lead to a server crash. (Bug #49897)
- `sql_buffer_result` had an effect on non-`SELECT` statements, contrary to the documentation. (Bug #49552)
- In some cases a subquery need not be evaluated because it returns only aggregate values that can be calculated from table metadata. This sometimes was not handled by the enclosing subquery, resulting in a server crash. (Bug #49512)
- Mixing full-text searches and row expressions caused a crash. (Bug #49445)
- `mysql-test-run.pl` now recognizes the `MTR_TESTCASE_TIMEOUT`, `MTR_SUITE_TIMEOUT`, `MTR_SHUTDOWN_TIMEOUT`, and `MTR_START_TIMEOUT` environment variables. If they are set, their values are used to set the `--testcase-timeout`, `--suite-timeout`, `--shutdown-timeout`, and `--start-timeout` options, respectively. (Bug #49210)

- The optimizer could continue to execute a query after a storage engine reported an error, leading to a server crash. (Bug #46175)

Changes in MySQL Enterprise 5.1.43sp1 [QSP] (2010-03-25)

This is a *Service Pack* release of the MySQL Enterprise Server 5.1.

Bugs Fixed

- **InnoDB; Partitioning:** When an `ALTER TABLE ... REORGANIZE PARTITION` statement on an `InnoDB` table failed due to `innodb_lock_wait_timeout` expiring while waiting for a lock, `InnoDB` did not clean up any temporary files or tables which it had created. Attempting to reissue the `ALTER TABLE` statement following the timeout could lead to storage engine errors, or possibly a crash of the server. (Bug #47343)
- **InnoDB:** If `innodb_force_recovery` was set to 4 or higher, the server could crash when opening an `InnoDB` table containing an auto-increment column. MySQL versions 5.1.31 and later were affected. (Bug #46193)
- Referring to a subquery result in a `HAVING` clause could produce incorrect results. (Bug #50995)
- The `filesort` sorting method applied to a `CHAR(0)` column could lead to a server crash. (Bug #49897)
- `sql_buffer_result` had an effect on non-`SELECT` statements, contrary to the documentation. (Bug #49552)
- In some cases a subquery need not be evaluated because it returns only aggregate values that can be calculated from table metadata. This sometimes was not handled by the enclosing subquery, resulting in a server crash. (Bug #49512)
- `flush_cache_records()` did not correctly check for errors that should cause statement execution to stop, leading to a server crash. (Bug #39022)

Changes in MySQL 5.1.43 (2010-01-15)

- [InnoDB Plugin Notes](#)
- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

InnoDB Plugin Notes

- This release includes `InnoDB Plugin` 1.0.6. This version is considered of Release Candidate (RC) quality.

In this release, the `InnoDB Plugin` is included in source and binary distributions, except RHEL3, RHEL4, SuSE 9 (x86, x86_64, ia64), and generic Linux RPM packages. It also does not work for FreeBSD 6 and HP-UX or for Linux on S/390, PowerPC, and generic ia64.

Functionality Added or Changed

- **Partitioning:** The `UNIX_TIMESTAMP()` function is now supported in partitioning expressions using `TIMESTAMP` columns. For example, it is now possible to create a partitioned table such as this one:

```
CREATE TABLE t (c TIMESTAMP)
PARTITION BY RANGE ( UNIX_TIMESTAMP(c) ) (
  PARTITION p0 VALUES LESS THAN (631148400),
```

```
PARTITION p1 VALUES LESS THAN (946681200),  
PARTITION p2 VALUES LESS THAN (MAXVALUE)  
);
```

All other expressions involving `TIMESTAMP` values are now rejected with an error for attempts to create a new partitioned table or to alter an existing partitioned table.

When accessing an existing partitioned table having a timezone-dependent partitioning function (where the table was using a previous version of MySQL), a warning rather than an error is issued. In such cases, you should fix the table. One way of doing this is to alter the table's partitioning expression so that it uses `UNIX_TIMESTAMP()`. (Bug #42849)

Bugs Fixed

- **Security Fix:** For servers built with yaSSL, a preauthorization buffer overflow could cause memory corruption or a server crash. We thank Evgeny Legerov from Intevydis for providing us with a proof-of-concept script that permitted us to reproduce this bug. (Bug #50227, CVE-2009-4484)
- **Performance; Partitioning:** When used on partitioned tables, the `records_in_range` handler call checked more partitions than necessary. The fix for this issue reduces the number of unpruned partitions checked for statistics in partition range checking, which has resulted in some partition operations being performed up to 2-10 times faster than before this change was made, when testing with tables having 1024 partitions. (Bug #48846)

References: See also: Bug #37252, Bug #47261.

- **Important Change; Replication:** The `RAND()` function is now marked as unsafe for statement-based replication. Using this function now generates a warning when `binlog_format=STATEMENT` and causes the format to switch to row-based logging when `binlog_format=MIXED`.

This change is being introduced because, when `RAND()` was logged in statement mode, the seed was also written to the binary log, so the replication slave generated the same sequence of random numbers as was generated on the master. While this could make replication work in some cases, the order of affected rows was still not guaranteed when this function was used in statements that could update multiple rows, such as `UPDATE` or `INSERT ... SELECT`; if the master and the slave retrieved rows in different order, they began to diverge. (Bug #49222)

- **InnoDB:** When compiling on Windows, an error in the `CMake` definitions for `InnoDB` caused the engine to be built incorrectly. (Bug #49502)
- **InnoDB:** The `InnoDB Monitor` could fail to print diagnostic information after a long lock wait. (Bug #47814)
- **InnoDB:** Crash recovery did not work for `InnoDB` temporary tables. (Bug #41609)
- **Partitioning:** A query that searched on a `ucs2` column failed if the table was partitioned. (Bug #48737)
- **Replication:** A `LOAD DATA INFILE` statement that loaded data into a table having a column name that had to be quoted (such as ``key` INT`) caused replication to fail when logging in mixed or statement mode. In such cases, the master wrote the `LOAD DATA` event into the binary log without quoting the column name. (Bug #49479)

References: See also: Bug #47927. This issue is a regression of: Bug #43746.

- **Replication:** Spatial data types caused row-based replication to crash. (Bug #48776)
- **Replication:** A flaw in the implementation of the purging of binary logs could result in orphaned files being left behind in the following circumstances:

- If the server failed or was killed while purging binary logs.

If the server failed or was killed after creating of a new binary log when the new log file was opened for the first time.

In addition, if the slave was not connected during the purge operation, it was possible for a log file that was in use to be removed; this could lead data loss and possible inconsistencies between the master and slave. (Bug #45292)

- **Replication:** When using the `STATEMENT` or `MIXED` logging format, the statements `LOAD DATA CONCURRENT LOCAL INFILE` and `LOAD DATA CONCURRENT INFILE` were logged as `LOAD DATA LOCAL INFILE` and `LOAD DATA LOCAL INFILE`, respectively (in other words, the `CONCURRENT` keyword was omitted). As a result, when using replication with either of these logging modes, queries on the slaves were blocked by the replication SQL thread while trying to execute the affected statements. (Bug #34628)
- **Replication:** Manually removing entries from the binary log index file on a replication master could cause the server to repeatedly send the same binary log file to slaves. (Bug #28421)
- **Cluster Replication:** When `expire_logs_days` was set, the thread performing the purge of the log files could deadlock, causing all binary log operations to stop. (Bug #49536)
- Within a stored routine, selecting the result of `CONCAT_WS()` with a routine parameter argument into a user variable could return incorrect results. (Bug #50096)
- The `IBMDB2I` storage engine was missing from the i5os 64-bit distribution of MySQL 5.1.42. It is now included again. (Bug #50059)
- `EXPLAIN EXTENDED UNION ... ORDER BY` caused a crash when the `ORDER BY` referred to a nonconstant or full-text function or a subquery. (Bug #49734)
- The `push_warning_printf()` function was being called with an invalid error level, `MYSQL_ERROR::WARN_LEVEL_ERROR`, causing an assertion failure. To fix the problem, `MYSQL_ERROR::WARN_LEVEL_ERROR` has been replaced by `MYSQL_ERROR::WARN_LEVEL_WARN`. (Bug #49638)
- Some prepared statements could raise an assertion when re-executed. (Bug #49570)
- A Valgrind error in `make_cond_for_table_from_pred()` was corrected. Thanks to Sergey Petrunya for the patch to fix this bug. (Bug #49506)
- Valgrind warnings for `CHECKSUM TABLE` were corrected. (Bug #49465)
- Specifying an index algorithm (such as `BTREE`) for `SPATIAL` or `FULLTEXT` indexes caused a server crash. These index types do not support algorithm specification, and it is not longer permitted to do so. (Bug #49250)
- The optimizer sometimes incorrectly handled conditions of the form `WHERE col_name='const1' AND col_name='const2'`. (Bug #49199)
- Execution of `DECODE()` and `ENCODE()` could be inefficient because multiple executions within a single statement reinitialized the random generator multiple times even with constant parameters. (Bug #49141)
- MySQL 5.1 does not support 2-byte collation numbers, but did not check the number and crashed for out-of-range values. (Bug #49134)

- With binary logging enabled, `REVOKE ... ON {PROCEDURE|FUNCTION} FROM ...` could cause a crash. (Bug #49119)
- The `LIKE` operator did not work correctly when using an index for a `ucs2` column. (Bug #49028)
- `check_key_in_view()` was missing a `DEBUG_RETURN` in one code branch, causing a crash in debug builds. (Bug #48995)
- Several `strmake()` calls had an incorrect length argument (too large by one). (Bug #48983)
- On Fedora 12, `strmov()` did not guarantee correct operation for overlapping source and destination buffer. Calls were fixed to use an overlap-safe version instead. (Bug #48866)
- Incomplete reset of internal `TABLE` structures could cause a crash with `eq_ref` table access in subqueries. (Bug #48709)
- Re-execution of a prepared statement could cause a server crash. (Bug #48508)
- The error message for `ER_UPDATE_INFO` was subject to buffer overflow or truncation. (Bug #48500)
- `SHOW BINLOG EVENTS` could fail with a error: `Wrong offset or I/O error`. (Bug #48357)
- Valgrind warnings related to binary logging of `LOAD DATA INFILE` statements were corrected. (Bug #48340)
- An aliasing violation in the C API could lead to a crash. (Bug #48284)
- With one thread waiting for a lock on a table, if another thread dropped the table and created a new table with the same name and structure, the first thread did not notice that the table had been re-created and tried to use cached metadata that belonged to the old table but had been freed. (Bug #48157)
- Queries containing `GROUP BY ... WITH ROLLUP` that did not use indexes could return incorrect results. (Bug #47650)
- If an invocation of a stored procedure failed in the table-open stage, subsequent invocations that did not fail in that stage could cause a crash. (Bug #47649)
- On Solaris, the server printed no stack trace to the error log after a crash. (Bug #47391)
- A crash occurred when a user variable that was assigned to a subquery result was used as a result field in a `SELECT` statement with aggregate functions. (Bug #47371)
- The first execution of `STOP SLAVE UNTIL` stopped too early. (Bug #47210)
- When the `mysql` client was invoked with the `--vertical` option, it ignored the `--skip-column-names` option. (Bug #47147)
- It was possible for `init_available_charsets()` not to initialize correctly. (Bug #45058)
- For a `VARCHAR(N)` column, `ORDER BY BINARY(col_name)` sorted using only the first `N` bytes of the column, even though column values could be longer than `N` bytes if they contained multibyte characters. (Bug #44131)
- Comparison with `NULL` values sometimes did not produce a correct result. (Bug #42760)
- The `mysql_upgrade` command added three columns to the `mysql.proc` table (`character_set_client`, `collation_connection`, and `db_collation`), but did not populate the columns with correct values. This led to error messages reported during stored procedure execution. (Bug #41569)

- When compressed [MyISAM](#) files were opened, they were always memory mapped, sometimes causing memory-swapping problems. To deal with this, a new system variable, [myisam_mmap_size](#), was added to permit limiting the amount of memory used for memory mapping of [MyISAM](#) files. (Bug #37408)
- A race condition on the privilege hash tables permitted one thread to try to delete elements that had already been deleted by another thread. A consequence was that [SET PASSWORD](#) or [FLUSH PRIVILEGES](#) could cause a crash. (Bug #35589, Bug #35591)
- [ALTER TABLE](#) with both [DROP COLUMN](#) and [ADD COLUMN](#) clauses could crash or lock up the server. (Bug #31145)

Changes in MySQL 5.1.42 (2009-12-15)

- [InnoDB Plugin Notes](#)
- [Platform-Specific Notes](#)
- [Bugs Fixed](#)

InnoDB Plugin Notes

- [InnoDB Plugin](#) has been upgraded to version 1.0.6. This version is considered of Release Candidate (RC) quality. [InnoDB Storage Engine Change History](#), may contain information in addition to those changes reported here.

In this release, the [InnoDB Plugin](#) is included in source and binary distributions, except RHEL3, RHEL4, SuSE 9 (x86, x86_64, ia64), and generic Linux RPM packages. It also does not work for FreeBSD 6 and HP-UX or for Linux on S/390, PowerPC, and generic ia64.

Platform-Specific Notes

- MySQL Server 5.1 is available on the following new platforms starting with the 5.1.42 release:
 - Mac OS X 10.6 x86/x64
 - HP-UX 11.31 IA64
 - SLES 11 x86/x64

Bugs Fixed

- **Performance:** When the query cache is fragmented, the size of the free block lists in the memory bins grows, which causes query cache invalidation to become slow. There is now a 50ms timeout for a [SELECT](#) statement waiting for the query cache lock. If the timeout expires, the statement executes without using the query cache. (Bug #39253)

References: See also: Bug #21074.

- **Important Change; Replication:** The following functions have been marked unsafe for statement-based replication:
 - [GET_LOCK\(\)](#)
 - [IS_FREE_LOCK\(\)](#)
 - [IS_USED_LOCK\(\)](#)
 - [MASTER_POS_WAIT\(\)](#)

- `RELEASE_LOCK()`
- `SLEEP()`
- `SYSDATE()`
- `VERSION()`

None of the functions just listed are guaranteed to replicate correctly when using the statement-based format because they can produce different results on the master and the slave. The use of any of these functions while `binlog_format` is set to `STATEMENT` is logged with the warning, `Statement is not safe to log in statement format`. When `binlog_format` is set to `MIXED`, the binary logging format is automatically switched to the row-based format whenever one of these functions is used. (Bug #47995)

- **Important Change:** After a binary upgrade to MySQL 5.1 from a MySQL 5.0 installation that contains `ARCHIVE` tables:

- Before MySQL 5.1.42, accessing those tables will cause the server to crash, even if you have run `mysql_upgrade` or `CHECK TABLE ... FOR UPGRADE`.
- As of MySQL 5.1.42, the server will not open 5.0 `ARCHIVE` tables at all.

In either case, the solution is to use `mysqldump` to dump all 5.0 `ARCHIVE` tables before upgrading, and reload them into MySQL 5.1 after upgrading. The same problem occurs for binary downgrades from MySQL 5.1 to 5.0. (Bug #47012)

- **Partitioning:** In some cases, it was not possible to add a new column to a table that had subpartitions. (Bug #48276)
- **Partitioning:** `SELECT COUNT(*)` from a partitioned table failed when using the `ONLY_FULL_GROUP_BY` SQL mode. (Bug #46923)

References: This issue is a regression of: Bug #45807.

- **Partitioning:** `SUBPARTITION BY KEY` failed with `DEFAULT CHARSET=utf8`. (Bug #45904)
- **Replication:** With row-based binary logging, the server crashed for statements of the form `CREATE TABLE IF NOT EXISTS existing_view LIKE temporary_table`. This occurred because the server handled the existing view as a table when logging the statement. (Bug #48506)
- **Replication:** When using row-based logging, `TRUNCATE TABLE` was written to the binary log even if the affected table was temporary, causing replication to fail. (Bug #48350)
- **Replication:** Replicating `TEXT` or `VARCHAR` columns declared as `NULL` on the master but `NOT NULL` on the slave caused the slave to crash. (Bug #43789)

References: See also: Bug #38850, Bug #43783, Bug #43785, Bug #47741, Bug #48091.

- **Replication:** When using row-based format, replication failed with the error `Could not execute Write_rows event on table ...; Field '...' doesn't have a default value` when an `INSERT` was made on the master without specifying a value for a column having no default, even if strict server SQL mode was not in use and the statement would otherwise have succeeded on the master. Now the SQL mode is checked, and the statement is replicated unless strict mode is in effect. For more information, see [Server SQL Modes](#). (Bug #38173)

References: See also: Bug #38262, Bug #43992.

- The result of comparison between nullable `BIGINT` and `INT` columns was inconsistent. (Bug #49517)
- Incorrect cache initialization prevented storage of converted constant values and could produce incorrect comparison results. (Bug #49489)
- Comparisons involving `YEAR` values could produce incorrect results. (Bug #49480)

References: See also: Bug #43668.

- If a query involving a table was terminated with `KILL`, a subsequent `SHOW CREATE TABLE` for that table caused a server crash. (Bug #48985)
- Privileges for stored routines were ignored for mixed-case routine names. (Bug #48872)

References: See also: Bug #41049.

- Building MySQL on Fedora Core 12 64-bit failed, due to errors in `comp_err`. (Bug #48864)
- Concurrent `ALTER TABLE` operations on an `InnoDB` table could raise an assertion. (Bug #48782)
- Certain `INTERVAL` expressions could cause a crash on 64-bit systems. (Bug #48739)
- During query execution, ranges could be merged incorrectly for `OR` operations and return an incorrect result. (Bug #48665)
- The `InnoDB` Table Monitor reported the `FLOAT` and `DOUBLE` data types incorrectly. (Bug #48526)
- `DISTINCT` was ignored for queries with `GROUP BY WITH ROLLUP` and only `const` tables. (Bug #48475)
- Loose index scan was inappropriately chosen for some `WHERE` conditions. (Bug #48472)
- The server could crash and corrupt the tablespace if the `InnoDB` tablespace was configured with too small a value, or if `innodb_file_per_table` was enabled and many `CREATE TEMPORARY TABLE` statements were executed and the temporary file directory filled up. (Bug #48469)
- Parts of the range optimizer could be initialized incorrectly, resulting in Valgrind errors. (Bug #48459)
- A bad typecast could cause query execution to allocate large amounts of memory. (Bug #48458)
- `GRANT` and `REVOKE` crashed if a user name was specified as `CURRENT_USER()`. (Bug #48319)
- On Windows, `InnoDB` could not be built as a statically linked library. (Bug #48317)
- `mysql_secure_installation` did not work on Solaris. (Bug #48086)
- When running `mysql_secure_installation`, the command failed if the `root` password contained multiple space, `'\'`, `'#'`, or quote characters. (Bug #48031)
- `MATCH IN BOOLEAN MODE` searches could return too many results inside a subquery. (Bug #47930)
- Using `REPLACE` to update a previously inserted negative value in an `AUTO_INCREMENT` column of an `InnoDB` table caused the table auto-increment value to be updated to 2147483647. (Bug #47720)
- If a session acquired a global read lock with `FLUSH TABLES WITH READ LOCK`, acquired a lock for one table with `LOCK TABLES`, and issued an `INSERT DELAYED` statement for another table, deadlock could occur. (Bug #47682)
- The `mysql` client `status` command displayed an incorrect value for the server character set. (Bug #47671)

- Connecting to a 4.1.x server from a 5.1.x or higher `mysql` client resulted in a memory-free error when disconnecting. (Bug #47655)
- Assignment of a system variable sharing the same base name as a declared stored program variable in the same context could lead to a crash. (Bug #47627)
- `mysqladmin debug` could crash on 64-bit systems. (Bug #47382)
- The `innodb_file_format_check` system variable could not be set at runtime to `DEFAULT` or to the value of a user-defined variable. (Bug #47167)
- The Mac OS X `MySQL Preference Pane` component was not built for 64-bit, which would trigger the System Preferences application to restart into 32-bit mode. (Bug #46935)
- The `IGNORE` clause on a `DELETE` statement masked an SQL statement error that occurred during trigger processing. (Bug #46425)
- On 64-bit systems, `--skip-innodb` did not skip `InnoDB` startup. (Bug #46043)
- Valgrind errors for `InnoDB Plugin` were corrected. (Bug #45992, Bug #46656)
- The return value was not checked for some `my_hash_insert()` calls. (Bug #45613)
- Truncation of `DECIMAL` values could lead to assertion failures; for example, when deducing the type of a table column from a literal `DECIMAL` value. (Bug #45261)

References: See also: Bug #48370.

- For `YEAR(2)` values, `MIN()`, `MAX()`, and comparisons could yield incorrect results. (Bug #43668)
- The server could crash when attempting to access a non-conformant `mysql.proc` system table. For example, the server could crash when invoking stored procedure-related statements after an upgrade from MySQL 5.0 to 5.1 without running `mysql_upgrade`. (Bug #41726)
- Multiple-statement execution could fail. (Bug #40877)
- Use of `InnoDB` monitoring (`SHOW ENGINE INNODB STATUS` or one of the `InnoDB` Monitor tables) could cause a server crash due to invalid access to a shared variable in a concurrent environment. This is a further fix for a regression introduced in MySQL 5.1.38 to the original fix in MySQL 5.1.31. (Bug #38883)
- On Windows, the `mysql_secure_installation` command failed to load the `Term::ReadKey` module, which was required for correct operation. (Bug #35106)
- If the `--log-bin` server option was set to a directory name with a trailing component separator character, the basename of the binary log files was empty, so that the created files were named `.000001` and `.index`. The same thing occurred with the `--log-bin-index`, `--relay-log`, and `--relay-log-index` options. Now the server reports and error and exits. (Bug #34739)
- If a comparison involved a constant value that required type conversion, the converted value might not be cached, resulting in repeated conversion and poorer performance. (Bug #34384)
- Using the `SHOW ENGINE INNODB STATUS` statement when using partitions in `InnoDB` tables caused `Invalid (old?) table or database name` errors to be logged. (Bug #32430)
- On some Windows systems, `InnoDB` could report `Operating system error number 995 in a file operation` due to transient driver or hardware problems. `InnoDB` now retries the operation and adds `Retry attempt is made` to the error message. (Bug #3139)

Changes in MySQL 5.1.41 (2009-11-05)

- [InnoDB Plugin Notes](#)
- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

InnoDB Plugin Notes

- [InnoDB Plugin](#) has been upgraded to version 1.0.5. This version is considered of Release Candidate (RC) quality. [InnoDB Storage Engine Change History](#), may contain information in addition to those changes reported here.

Functionality Added or Changed

- **Incompatible Change:** For [InnoDB Plugin](#), two status variables have been added: [Innodb_buffer_pool_read_ahead](#) and [Innodb_buffer_pool_read_ahead_evicted](#) indicate the number of pages read in by the [InnoDB](#) read-ahead background thread, and the number of such pages evicted without ever being accessed, respectively. Also, the status variables [Innodb_buffer_pool_read_ahead_rnd](#) and [Innodb_buffer_pool_read_ahead_seq](#) status variables have been removed.

The built-in version of [InnoDB](#) is not affected by these changes. (Bug #42885)

- The [InnoDB](#) buffer pool is divided into two sublists: A new sublist containing blocks that are heavily used by queries, and an old sublist containing less-used blocks and from which candidates for eviction are taken. In the default operation of the buffer pool, a block when read in is loaded at the midpoint and then moved immediately to the head of the new sublist as soon as an access occurs. In the case of a table scan (such as performed for a [mysqldump](#) operation), each block read by the scan ends up moving to the head of the new sublist because multiple rows are accessed from each block. This occurs even for a one-time scan, where the blocks are not otherwise used by other queries. Blocks may also be loaded by the read-ahead background thread and then moved to the head of the new sublist by a single access. These effects can be disadvantageous because they push blocks that are in heavy use by other queries out of the new sublist to the old sublist where they become subject to eviction.

[InnoDB](#) now provides two system variables that enable LRU algorithm tuning:

- [innodb_old_blocks_pct](#)

Specifies the approximate percentage of the buffer pool used for the old block sublist. The range of values is 5 to 95. The default value is 37 (that is, 3/8 of the pool).

- [innodb_old_blocks_time](#)

Specifies how long in milliseconds (ms) a block inserted into the old sublist must stay there after its first access before it can be moved to the new sublist. The default value is 0: A block inserted into the old sublist moves immediately to the new sublist the first time it is accessed, no matter how soon after insertion the access occurs. If the value is greater than 0, blocks remain in the old sublist until an access occurs at least that many ms after the first access. For example, a value of 1000 causes blocks to stay in the old sublist for 1 second after the first access before they become eligible to move to the new sublist.

See [The InnoDB Buffer Pool](#). (Bug #45015)

- The server now supports a Debug Sync facility for thread synchronization during testing and debugging. To compile in this facility, configure MySQL with the `--enable-debug-sync` option. The [debug_sync](#)

system variable provides the user interface Debug Sync. `mysqld` and `mysql-test-run.pl` support a `--debug-sync-timeout` option to enable the facility and set the default synchronization point timeout.

Bugs Fixed

- **Security Fix; Important Change:** Additional corrections were made for the symlink-related privilege problem originally addressed in MySQL 5.1.24. The original fix did not correctly handle the data directory path name if it contained symlinked directories in its path, and the check was made only at table-creation time, not at table-opening time later. (Bug #32167, CVE-2008-2079)

References: See also: Bug #39277.

- **Security Fix:** MySQL clients linked against OpenSSL could be tricked not to check server certificates. (Bug #47320, CVE-2009-4028)
- **InnoDB:** The `SHOW CREATE TABLE` statement could result in a serious error for some views, especially views involving comparisons between date/datetime/timestamp columns and strings returned by scalar subqueries. (Bug #11746321, Bug #25439)
- **InnoDB:** When a trigger inserts into a table containing an auto-increment column, an error `Error: Duplicate entry` could occur with the InnoDB Plugin if another insert was happening simultaneously. (Bug #26316)
- **Partitioning:** An `ALTER TABLE ... ADD PARTITION` statement that caused `open_files_limit` to be exceeded led to a MySQL server crash. (Bug #46922)

References: See also: Bug #47343.

- **Partitioning:** The cardinality of indexes on partitioned tables was calculated using the first partition in the table, which could result in suboptimal query execution plans being chosen. Now the partition having the most records is used instead, which should result in better use of indexes and thus improved performance of queries against partitioned tables in many if not most cases. (Bug #44059)
- **Replication:** This issue occurred in MySQL 5.1.40 only. (Bug #48297)
- **Replication:** When a session was closed on the master, temporary tables belonging to that session were logged with the wrong database names when either of the following conditions was true:
 1. The length of the name of the database to which the temporary table belonged was greater than the length of the current database name.
 2. The current database was not set.

(Bug #48216)

References: See also: Bug #46861, Bug #48297.

- **Replication:** When using row-based replication, changes to nontransactional tables that occurred early in a transaction were not immediately flushed upon committing a statement. This behavior could break consistency since changes made to nontransactional tables become immediately visible to other connections. (Bug #47678)

References: See also: Bug #47287, Bug #46864, Bug #43929, Bug #11752675, Bug #46129. This issue is a regression of: Bug #40116.

- **Replication:** When `mysqlbinlog --verbose` was used to read a binary log that had been written using row-based format, the output for events that updated some but not all columns of tables was not correct. (Bug #47323)

- **Replication:** When using the row-based format to replicate a transaction involving both transactional and nontransactional engines, which contained a DML statement affecting multiple rows, the statement failed. If this transaction was followed by a `COMMIT`, the master and the slave could diverge, because the statement was correctly rolled back on the master, but was applied on the slave. (Bug #47287)

References: See also: Bug #46864.

- **Replication:** A problem with the `BINLOG` statement in the output of `mysqlbinlog` could break replication; statements could be logged with the server ID stored within events by the `BINLOG` statement rather than the ID of the running server. With this fix, the server ID of the server executing the statements can no longer be overridden by the server ID stored in the binary log's `format description` statement. (Bug #46640)

References: This issue is a regression of: Bug #32407.

- **Replication:** When using statement-based replication and the transaction isolation level was set to `READ COMMITTED` or a less strict level, `InnoDB` returned an error even if the statement in question was filtered out according to the `--binlog-do-db` or `--binlog-ignore-db` rules in effect at the time. (Bug #42829)

- **Replication:** `FLUSH LOGS` did not close and reopen the binary log index file. (Bug #34582)

References: See also: Bug #48738.

- `SUM()` artificially increased the precision of a `DECIMAL` argument, which was truncated when a temporary table was created to hold the results. (Bug #48370)

References: See also: Bug #45261.

- If an outer query was invalid, a subquery might not be set up. `EXPLAIN EXTENDED` did not expect this and caused a crash by trying to dereference improperly set up information. (Bug #48295)
- A query containing a view using temporary tables and multiple tables in the `FROM` clause and `PROCEDURE ANALYSE()` caused a server crash.

As a result of this bug fix, `PROCEDURE ANALYSE()` is legal only in a top-level `SELECT`. (Bug #48293)

References: See also: Bug #46184.

- Error handling was missing for `SELECT` statements containing subqueries in the `WHERE` clause and that assigned a `SELECT` result to a user variable. The server could crash as a result. (Bug #48291)
- An assertion could fail if the optimizer used a `SPATIAL` index. (Bug #48258, Bug #47019)
- `InnoDB` mishandled memory-allocation failures in the `os_mem_alloc_large()` function. (Bug #48237)
- `WHERE` clauses with `outer_value_list NOT IN subquery` were handled incorrectly if the outer value list contained multiple items at least one of which could be `NULL`. (Bug #48177)
- A combination of `GROUP BY WITH ROLLUP`, `DISTINCT` and the `const` join type in a query caused a server crash when the optimizer used a temporary table to resolve `DISTINCT`. (Bug #48131)
- In some cases, using a null microsecond part in a `WHERE` condition (for example, `WHERE date_time_field <= 'YYYY-MM-DD HH:MM:SS.0000'`) could lead to incorrect results due to improper `DATETIME` comparison. (Bug #47963)
- A build configured using the `--without-server` option did not compile the `yaSSL` code, so if `--with-ssl` was also used, the build failed. (Bug #47957)

- When a query used a `DATE` or `DATETIME` value formatted using any separator characters other than hyphen ('-') and a `>=` condition matching only the greatest value in an indexed column, the result was empty if an index range scan was employed. (Bug #47925)
- `mysys/mf_keycache.c` requires threading, but no test was made for thread support. (Bug #47923)
- For debug builds, an assertion could fail during the next statement executed for a temporary table after a multiple-table `UPDATE` involving that table modified an `AUTO_INCREMENT` column with a user-supplied value. (Bug #47919)
- The `mysys/mf_strip.c` file, which defines the `strip_sp()` function, has been removed from the MySQL source. The function was no longer used within the main build, and the supplied function was causing symbol errors on Windows builds. (Bug #47857)
- The Windows build for MySQL would compile the `split.c` and `debug.c` files unnecessarily, causing additional symbols to be included in `mysqld`. (Bug #47850)
- When building storage engines on Windows it was not possible to specify additional libraries within the `CMake` file required for the build. An `$(engine)_LIBS` macro has been included in the files to support these additional storage-engine specific libraries. (Bug #47797)
- When building a pluggable storage engine on Windows, the engine name could be based on the directory name where the engine was located, rather than the configured storage engine name. (Bug #47795)
- During cleanup of a stored procedure's internal structures, the flag to ignore the errors for `INSERT IGNORE` or `UPDATE IGNORE` was not cleaned up, which could result in a server crash. (Bug #47788)
- If the first argument to `GeomFromWKB()` function was a geometry value, the function just returned its value. However, it failed to preserve the argument's `null_value` flag, which caused an unexpected `NULL` value to be returned to the caller, resulting in a server crash. (Bug #47780)
- `InnoDB` could crash when updating spatial values. (Bug #47777)
- On Windows, when an idle named pipe connection was forcibly closed with a `KILL` statement or because the server was being shut down, the thread that was closing the connection would hang infinitely.

As a result of the work done for this bug, the `net_read_timeout`, `net_write_timeout`, and `wait_timeout`, system variables now apply to connections over all transports, not just to TCP/IP. (Bug #47571, Bug #31621)

- A function call could end without throwing an error or setting the return value. For example, this could happen when an error occurred while calculating the return value. This is fixed by setting the value to `NULL` when an error occurs during evaluation of an expression. (Bug #47412)
- A simple `SELECT` with implicit grouping could return many rows rather than a single row if the query was ordered by the aggregated column in the select list. (Bug #47280)
- An assertion could be raised for `CREATE TABLE` if there was a pending `INSERT DELAYED` or `REPLACE DELAYED` for the same table. (Bug #47274)
- `InnoDB` raised errors in some cases in a manner not compatible with `SIGNAL` and `RESIGNAL`. (Bug #47233)
- If an `InnoDB` table was created with the `AUTO_INCREMENT` table option to specify an initial auto-increment value, and an index was added in a separate operation later, the auto-increment value was lost (subsequent inserts began at 1 rather than the specified value). (Bug #47125)

- Incorrect handling of predicates involving `NULL` by the range optimizer could lead to an infinite loop during query execution. (Bug #47123)
- Repair by sort or parallel repair of `MyISAM` tables might not fail over to repair with key cache. (Bug #47073)
- `InnoDB Plugin` did not compile on some Solaris systems. (Bug #47058)
- On Windows, when a failed I/O operation occurred with return code of `ERROR_WORKING_SET_QUOTA`, `InnoDB` intentionally crashed the server. Now `InnoDB` sleeps for 100ms and retries the failed operation. (Bug #47055)
- `InnoDB` now ignores negative values supplied by a user for an `AUTO_INCREMENT` column when calculating the next value to store in the data dictionary. Setting `AUTO_INCREMENT` columns to negative values is undefined behavior and this change should bring the behavior of `InnoDB` closer to what users expect. (Bug #46965)
- When MySQL crashed (or a snapshot was taken that simulates a crash), it was possible that internal XA transactions (used to synchronize the binary log and `InnoDB`) could be left in a `PREPARED` state, whereas they should be rolled back. This occurred when the `server_id` value changed before the restart, because that value was used to construct XID values.

Now the restriction is relaxed that the `server_id` value be consistent for XID values to be considered valid. The rollback phase should then be able to clean up all pending XA transactions. (Bug #46944)

- `InnoDB Plugin` did not compile using `gcc` 4.1 on PowerPC systems. (Bug #46718)
- If `InnoDB Plugin` reached its limit on the number of concurrent transactions (1023), it wrote a descriptive message to the error log but returned a misleading error message to the client, or an assertion failure occurred. (Bug #46672)

References: See also: Bug #18828.

- A Valgrind error during index creation by `InnoDB Plugin` was corrected. (Bug #46657)
- Concurrent `INSERT INTO ... SELECT` statements for an `InnoDB` table could cause an `AUTO_INCREMENT` assertion failure. (Bug #46650)
- If a transaction was rolled back inside `InnoDB` due to a deadlock or lock wait timeout, and a statement in the transaction had an `IGNORE` clause, the server could crash at the end of the statement or on shutdown. (Bug #46539)
- Trailing spaces were not ignored for user-defined collations that mapped spaces to a character other than 0x20. (Bug #46448)

References: See also: Bug #29468.

- The GPL and commercial license headers had different sizes, so that error log, backtrace, core dump, and cluster trace file line numbers could be off by one if they were not checked against the version of the source used for the build. (For example, checking a GPL build backtrace against commercial sources.) (Bug #46216)
- `InnoDB` did not disallow creation of an index with the name `GEN_CLUST_INDEX`, which is used internally. (Bug #46000)
- During the build of the Red Hat IA64 MySQL server RPM, the system library link order was incorrect. This made the resulting Red Hat IA64 RPM depend on "libc.so.6.1(GLIBC_PRIVATE)(64bit)", thus preventing installation of the package. (Bug #45706)

- The `caseinfo` member of the `CHARSET_INFO` structure was not initialized for user-defined Unicode collations, leading to a server crash. (Bug #45645)
- With `InnoDB Plugin`, renaming a table column and then creating an index on the renamed column caused a server crash due to the `.frm` file and the `InnoDB` data directory going out of sync. Now `InnoDB Plugin` 1.0.5 returns an error instead: `ERROR 1034 (HY000): Incorrect key file for table 'tbl_name'; try to repair it`. To work around the problem, create another table with the same structure and copy the original table to it. (Bug #44571)
- An `InnoDB` error message incorrectly referred to the nonexistent `innodb_max_files_open` variable rather than to `innodb_open_files`. (Bug #44338)
- For `ALTER TABLE`, renaming a `DATETIME` or `TIMESTAMP` column unnecessarily caused a table copy operation. (Bug #43508)
- The weekday names for the Romanian `lc_time_names` locale `'ro_RO'` were incorrect. Thanks to Andrei Boros for the patch to fix this bug. (Bug #43207)
- `XA START` could cause an assertion failure or server crash when it is called after a unilateral rollback issued by the Resource Manager (both in a regular transaction and after an XA transaction). (Bug #43171)
- The `FORCE INDEX FOR ORDER BY` index hint was ignored when join buffering was used. (Bug #43029)
- Incorrect handling of range predicates combined with `OR` operators could yield incorrect results. (Bug #42846)
- Failure to treat `BIT` values as unsigned could lead to unpredictable results. (Bug #42803)
- For the embedded server on Windows, `InnoDB` crashed when `innodb_file_per_table` was enabled and a table name was in full path format. (Bug #42383)
- Some queries with nested outer joins could lead to crashes or incorrect results because an internal data structure was handled improperly. (Bug #42116)
- In a replication scenario with `innodb_locks_unsafe_for_binlog` enabled on the slave, where rows were changed only on the slave (not through replication), in some rare cases, many messages of the following form were written to the slave error log: `InnoDB: Error: unlock row could not find a 4 mode lock on the record`. (Bug #41756)
- After renaming a user, granting that user privileges could result in the user having privileges additional to those granted. (Bug #41597)
- With a nonstandard `InnoDB` page size, some error messages became inaccurate.

**Note**

Changing the page size is not a supported operation and there is no guarantee that `InnoDB` will function normally with a page size other than 16KB. Problems compiling or running `InnoDB` may occur. In particular, `ROW_FORMAT=COMPRESSED` in the `InnoDB Plugin` assumes that the page size is at most 16KB and uses 14-bit pointers.

A version of `InnoDB` built for one page size cannot use data files or log files from a version built for a different page size.

(Bug #41490)

- In some cases, the server did not recognize lettercase differences between `GRANT` attributes such as table name or user name. For example, a user was able to perform operations on a table with privileges of another user with the same user name but in a different lettercase.

In consequence of this bug fix, the collation for the `Routine_name` column of the `mysql.proc` table is changed from `utf8_bin` to `utf8_general_ci`. (Bug #41049)

References: See also: Bug #48872.

- Simultaneous `ANALYZE TABLE` operations for an `InnoDB` tables could be subject to a race condition. (Bug #38996)
- Previously, `InnoDB` performed `REPLACE INTO T SELECT ... FROM S WHERE ...` by setting shared next-key locks on rows from `S`. Now `InnoDB` selects rows from `S` with shared locks or as a consistent read, as for `INSERT ... SELECT`. This reduces lock contention between sessions. (Bug #37232)
- When an `InnoDB` tablespace filled up, an error was logged to the client, but not to the error log. Also, the error message was misleading and did not indicate the real source of the problem. (Bug #31183)
- In `mysql`, using **Control+C** to kill the current query resulted in a `ERROR 1053 (08S01): Server shutdown in progress` message if the query was waiting for a lock. (Bug #28141)

Changes in MySQL Enterprise 5.1.40sp1 [QSP] (2009-11-25)

This is a *Service Pack* release of the MySQL Enterprise Server 5.1.

Bugs Fixed

- **Replication:** When using statement-based or mixed-format replication, the database name was not written to the binary log when executing a `LOAD DATA INFILE` statement. This caused problems when the table being loaded belonged to a database other than the current database; data could be loaded into the wrong table (if a table having the same name existed in the current database) or replication could fail (if no table having that name existed in the current database). Now a table referenced in a `LOAD DATA INFILE` statement is always logged using its fully qualified name when the database to which it belongs is not the current database. (Bug #48297)
- **Replication:** When a session was closed on the master, temporary tables belonging to that session were logged with the wrong database names when either of the following conditions was true:
 1. The length of the name of the database to which the temporary table belonged was greater than the length of the current database name.
 2. The current database was not set.

(Bug #48216)

References: See also: Bug #46861, Bug #48297.

- `SUM()` artificially increased the precision of a `DECIMAL` argument, which was truncated when a temporary table was created to hold the results. (Bug #48370)

References: See also: Bug #45261.

- If an outer query was invalid, a subquery might not be set up. `EXPLAIN EXTENDED` did not expect this and caused a crash by trying to dereference improperly set up information. (Bug #48295)

- A query containing a view using temporary tables and multiple tables in the `FROM` clause and `PROCEDURE ANALYSE()` caused a server crash.

As a result of this bug fix, `PROCEDURE ANALYSE()` is legal only in a top-level `SELECT`. (Bug #48293)

References: See also: Bug #46184.

- Error handling was missing for `SELECT` statements containing subqueries in the `WHERE` clause and that assigned a `SELECT` result to a user variable. The server could crash as a result. (Bug #48291)
- An assertion could fail if the optimizer used a `SPATIAL` index. (Bug #48258, Bug #47019)
- A combination of `GROUP BY WITH ROLLUP`, `DISTINCT` and the `const` join type in a query caused a server crash when the optimizer used a temporary table to resolve `DISTINCT`. (Bug #48131)
- In some cases, using a null microsecond part in a `WHERE` condition (for example, `WHERE date_time_field <= 'YYYY-MM-DD HH:MM:SS.0000'`) could lead to incorrect results due to improper `DATETIME` comparison. (Bug #47963)
- When a query used a `DATE` or `DATETIME` value formatted using any separator characters other than hyphen ('-') and a `>=` condition matching only the greatest value in an indexed column, the result was empty if an index range scan was employed. (Bug #47925)
- During cleanup of a stored procedure's internal structures, the flag to ignore the errors for `INSERT IGNORE` or `UPDATE IGNORE` was not cleaned up, which could result in a server crash. (Bug #47788)
- If the first argument to `GeomFromWKB()` function was a geometry value, the function just returned its value. However, it failed to preserve the argument's `null_value` flag, which caused an unexpected `NULL` value to be returned to the caller, resulting in a server crash. (Bug #47780)
- `InnoDB` could crash when updating spatial values. (Bug #47777)
- Incorrect handling of predicates involving `NULL` by the range optimizer could lead to an infinite loop during query execution. (Bug #47123)
- `InnoDB` now ignores negative values supplied by a user for an `AUTO_INCREMENT` column when calculating the next value to store in the data dictionary. Setting `AUTO_INCREMENT` columns to negative values is undefined behavior and this change should bring the behavior of `InnoDB` closer to what users expect. (Bug #46965)
- In a replication scenario with `innodb_locks_unsafe_for_binlog` enabled on the slave, where rows were changed only on the slave (not through replication), in some rare cases, many messages of the following form were written to the slave error log: `InnoDB: Error: unlock row could not find a 4 mode lock on the record.` (Bug #41756)

Changes in MySQL 5.1.40 (2009-10-06)

- [InnoDB Plugin Notes](#)
- [Bugs Fixed](#)

InnoDB Plugin Notes

- In this release, the `InnoDB Plugin` is included in source and binary distributions, except RHEL3, RHEL4, SuSE 9 (x86, x86_64, ia64), and generic Linux RPM packages. It also does not work for FreeBSD 6 and HP-UX or for Linux on S/390, PowerPC, and generic ia64.

Bugs Fixed

- **Incompatible Change; Replication:** Concurrent transactions that inserted rows into a table with an `AUTO_INCREMENT` column could break statement-based or mixed-format replication error 1062 `Duplicate entry '...' for key 'PRIMARY'` on the slave. This was especially likely to happen when one of the transactions activated a trigger that inserted rows into the table with the `AUTO_INCREMENT` column, although other conditions could also cause the issue to manifest.

As part of the fix for this issue, any statement that causes a trigger or function to update an `AUTO_INCREMENT` column is now considered unsafe for statement-based replication. For more information, see [Replication and AUTO_INCREMENT](#). (Bug #45677)

References: See also: Bug #42415, Bug #48608, Bug #50440, Bug #53079.

- **Incompatible Change:** In binary installations of MySQL, the supplied `binary-configure` script would start and configure MySQL, even when command help was requested with the `--help` command-line option. The `--help` option, if provided, no longer starts and installs the server. (Bug #30954)
- **Partitioning:** When reorganizing partitions, not all affected subpartitions were removed prior to renaming. One way in which the issue was visible was that attempting to reorganize two partitions into a single partition having the same name as one of the original partitions could lead to a crash of the server. (Bug #47029)

References: See also: Bug #45961, Bug #43729.

- **Partitioning:** An online or fast `ALTER TABLE` of a partitioned table could leave behind temporary files in the database directory.

This issue was observed in MySQL 5.1.31 and later only. (Bug #46483)

- **Partitioning:** When performing an `INSERT ... SELECT` into a partitioned table, `read_buffer_size` bytes of memory were allocated for every partition in the target table, resulting in consumption of large amounts of memory when the table had many partitions (more than 100).

This fix changes the method used to estimate the buffer size required for each partition and limits the total buffer size to a maximum of approximately 10 times `read_buffer_size`. (Bug #45840)

- **Partitioning:** Inserting negative values into an `AUTO_INCREMENT` column of a partitioned table could lead to apparently unrelated errors or a crash of the server.

This issue was observed in MySQL 5.1.31 and later only. (Bug #45823)

- **Partitioning:** Unnecessary calls were made in the server code for performing bulk inserts on partitions for which no inserts needed to be made. (Bug #35845)

References: See also: Bug #35843.

- **Replication:** Performing `ALTER TABLE ... DISABLE KEYS` on a slave table caused row-based replication to fail. (Bug #47312)
- **Replication:** `BEGIN` statements were not included in the output of `mysqlbinlog`. (Bug #46998)
- **Replication:** When using row-based replication, `DROP TEMPORARY TABLE IF EXISTS` was written to the binary log if the table named in the statement did not exist, even though a `DROP TEMPORARY TABLE` statement should never be logged in row-based logging mode, whether the table exists or not. (Bug #46572)
- **Replication:** When using row-based replication, importing a dump made with `mysqldump` and replicating a row with an `AUTO_INCREMENT` column set to 0, with `NO_AUTO_VALUE_ON_ZERO` active on the master, the row was inserted successfully on the master; however any setting for

`NO_AUTO_VALUE_ON_ZERO` was ignored on the slave. When the `AUTO_INCREMENT` column was incremented, this caused replication to fail on the slave due to a duplicate key error. In some cases it could also cause the slave to crash. (Bug #45999)

- **Replication:** By default, all statements executed by the `mysql_upgrade` program on the master are written to the binary log, then replicated to the slave. In some cases, this can result in problems; for example, it attempted to alter log tables on replicated databases (this failed due to logging being enabled).

As part of this fix, a `mysql_upgrade` option, `--write-binlog`, is added. Its inverse, `--skip-write-binlog`, can be used to disable binary logging while the upgrade is in progress. (Bug #43579)

- **Replication:** On the master, if a binary log event is larger than `max_allowed_packet`, the error message `ER_MASTER_FATAL_ERROR_READING_BINLOG` is sent to a slave when it requests a dump from the master, thus leading the I/O thread to stop. On a slave, the I/O thread stops when receiving a packet larger than `max_allowed_packet`.

In both cases, however, there was no `Last_IO_Error` reported, which made it difficult to determine why the slave had stopped in such cases. Now, `Last_IO_Error` is reported when `max_allowed_packet` is exceeded, and provides the reason for which the slave I/O thread stopped. (Bug #42914)

References: See also: Bug #14068, Bug #47200, Bug #47303.

- **API:** The fix for Bug #24507 could lead in some cases to client application failures due to a race condition. Now the server waits for the “dummy” thread to return before exiting, thus making sure that only one thread can initialize the POSIX threads library. (Bug #42850)

References: This issue is a regression of: Bug #24507.

- The `pthread_cond_wait()` implementations for Windows could deadlock in some rare circumstances. (Bug #47768)
- On Mac OS X or Windows, sending a `SIGHUP` signal to the server or an asynchronous flush (triggered by `flush_time`) caused the server to crash. (Bug #47525)
- Debug builds could not be compiled with the Sun Studio compiler. (Bug #47474)
- A multiple-table `UPDATE` involving a natural join and a mergeable view raised an assertion. (Bug #47150)
- Solaris binary packages now are compiled with `-g0` rather than `-g`. (Bug #47137)
- `EXPLAIN` caused a server crash for certain valid queries. (Bug #47106)
- The `configure` option `--without-server` did not work. (Bug #46980)
- The `ARCHIVE` storage engine lost records during a bulk insert. (Bug #46961)
- Failed multiple-table `DELETE` statements could raise an assertion. (Bug #46958)
- When creating a new instance on Windows using `mysqld-nt` and the `--install` parameter, the value of the service would be set incorrectly, resulting in a failure to start the configured service. (Bug #46917)
- `CONCAT_WS()` could return incorrect results due to an argument buffer also being used as a result buffer. (Bug #46815)
- The server crashed when re-using outer column references in correlated subqueries when the enclosing query used a temp table. (Bug #46791)

- For [InnoDB](#) tables, an unnecessary table rebuild for [ALTER TABLE](#) could sometimes occur for metadata-only changes. (Bug #46760)
- Assertion failure could result from repeated execution of a stored procedure containing an incorrect query with a subselect. (Bug #46629)
- The server ignored the setting of [sync_frm](#) for [CREATE TABLE ... LIKE](#). (Bug #46591)
- An attempt to create a table with the same name as an existing view could cause a server crash. (Bug #46384)
- A parser problem prevented properly stripping backquotes from an argument to a user-defined function (UDF). If the UDF was in an [ORDER BY](#) clause, its name would not be properly resolved against an alias with the same name in the select list. (Bug #46259)
- Dropping an [InnoDB](#) table that used an unknown collation (created on a different server, for example) caused a server crash. (Bug #46256)
- Certain [SELECT](#) statements containing [DISTINCT](#), [GROUP BY](#), and [HAVING](#) clauses could hang in an infinite loop. (Bug #46159)
- [InnoDB](#) did not disallow creation of an index with the name [GEN_CLUST_INDEX](#), which is used internally. (Bug #46000)
- [CREATE TEMPORARY TABLE](#) failed for [InnoDB](#) tables on systems with case-insensitive file systems when [lower_case_table_names = 2](#) and the pathname of the temporary file directory contained uppercase characters. (Bug #45638)
- Appending values to an [ENUM](#) or [SET](#) definition is a metadata change for which [ALTER TABLE](#) need not rebuild the table, but it was being rebuilt anyway. (Bug #45567)
- The [socket](#) system variable was unavailable on Windows. (Bug #45498)
- When re-installing MySQL on Windows on a server that has a data directory from a previous MySQL installation, the installer failed to identify the existence of the installation and the password configured for the [root](#) user. (Bug #45200)
- Client flags were incorrectly initialized for the embedded server, causing several tests in the [jp](#) test suite to fail. (Bug #45159)
- [InnoDB](#) did not always disallow creating tables containing columns with names that match the names of internal columns, such as [DB_ROW_ID](#), [DB_TRX_ID](#), [DB_ROLL_PTR](#), and [DB_MIX_ID](#). (Bug #44369)
- [SELECT ... WHERE ... IN \(NULL, ...\)](#) was executed using a full table scan, even if the same query without the [NULL](#) used an efficient range scan. (Bug #44139)

References: See also: Bug #18360.

- [InnoDB](#) use of [SELECT MAX\(autoinc_column\)](#) could cause a crash when MySQL data dictionaries went out of sync. (Bug #44030)
- [LOAD DATA INFILE](#) statements were written to the binary log in such a way that parsing problems could occur when re-executing the statement from the log. (Bug #43746)
- Selecting from the process list in the embedded server caused a crash. (Bug #43733)

References: See also: Bug #47304.

- Attempts to enable [large_pages](#) with a shared memory segment larger than 4GB caused a server crash. (Bug #43606)

- A test for stack growth failed on some platforms, leading to server crashes. (Bug #42213)
References: See also: Bug #62856.
- The server used the wrong lock type (always `TL_READ` instead of `TL_READ_NO_INSERT` when appropriate) for tables used in subqueries of `UPDATE` statements. This led in some cases to replication failure because statements were written in the wrong order to the binary log. (Bug #42108)
- The `mysql-stress-test.pl` test script was missing from the `noinstall` packages on Windows. (Bug #41546)
- Privileges for `SHOW CREATE VIEW` were not being checked correctly. (Bug #35996)
- Different invocations of `CHECKSUM TABLE` could return different results for a table containing columns with spatial data types. (Bug #35570)
- Concurrent execution of `FLUSH TABLES` along with `SHOW FUNCTION STATUS` or `SHOW PROCEDURE STATUS` could cause a server crash. (Bug #34895)
- `myisamchk` performed parameter value casting at startup that generated unnecessary warning messages. (Bug #33785)
- When using the `ARCHIVE` storage engine, `SHOW TABLE STATUS` displayed incorrect information for `Max_data_length`, `Data_length` and `Avg_row_length`. (Bug #29203)
- When building MySQL on Windows from source, the `WITH_BERKELEY_STORAGE_ENGINE` option would fail to configure `BDB` support correctly. (Bug #27693)

Changes in MySQL 5.1.39 (2009-09-04)

Bugs Fixed

- **Performance:** For `MyISAM` tables with `bulk_insert_buffer_size` values larger than 256KB, the performance of bulk insert operations such as multiple-row `INSERT` and `INSERT ... SELECT` operations has been improved greatly when up to a hundred rows are inserted at the same time. (Bug #44723)
- **Partitioning:** An `INSERT ... SELECT` statement on an empty partition of a partitioned table failed with `ERROR 1030 (HY000): Got error 124 from storage engine`. This issue also caused queries run against a partitioned table while a `LOAD DATA CONCURRENT INFILE` statement was in progress to fail with the same error. (Bug #46639)
References: See also: Bug #35845, Bug #44657, Bug #45840.
- **Partitioning:** A partitioned table having a `TIMESTAMP` column with a default value of `CURRENT_TIMESTAMP` and this column was not defined using an `ON UPDATE` option, an `ALTER TABLE ... REORGANIZE PARTITION` statement on the table caused the `TIMESTAMP` column value to be set to `CURRENT_TIMESTAMP` regardless. (Bug #46478)
- **Partitioning:** Partition pruning did not always work correctly when the table's partitioning key used the `TO_DAYS()` function. (Bug #46362)
- **Partitioning:** Attempting to access a partitioned table when partitioning support was disabled in a MySQL server binary that had been compiled with partitioning support caused the server to crash. (Bug #39893)
- **Partitioning:** The use of `TO_DAYS()` in the partitioning expression led to selection failures when the column having the date value contained invalid dates. This occurred because the function returns `NULL` in such cases, and the partition containing `NULL` values was pruned away. For example, this

problem occurred if '2001-02-00' was inserted into a `DATE` column of such a table, and a subsequent query on this table used `WHERE date_col < '2001-02-00'`—while '2001-01-01' is less than '2001-02-00', `TO_DAYS('2001-02-00')` evaluates as `NULL`, and so the row containing '2001-01-01' was not returned. Now, for tables using `RANGE` or `LIST` partitioning and having `TO_DAYS()` in the partitioning expression, the `NULL` partition is also scanned instead of being ignored.

The fix for this issue also corrects misbehavior such that a query of the form `SELECT * FROM table WHERE date_col < date_val` on a table partitioned by `RANGE` or `LIST` was handled as though the server SQL mode included `ALLOW_INVALID_DATES` even if this was not actually part of the server SQL mode at the time the query was issued. (Bug #20577)

References: See also: Bug #18198, Bug #32021, Bug #46362.

- **Replication:** Performing a multi-row update of the `AUTO_INCREMENT` column of a transactional table could result in an inconsistency between master and slave when there was a trigger on the transactional table that updated a nontransactional table. When such an update failed on the master, no rows were updated on the master, but some rows could (erroneously) be updated on the slave. (Bug #46864)
- **Replication:** When using the `--replicate-rewrite-db` option and the database referenced by this option on the master was the current database when the connection to the slave was closed, any temporary tables existing in this database were not properly dropped. (Bug #46861)
- **Replication:** When a statement that changed both transactional and nontransactional tables failed, the transactional changes were automatically rolled back on the master but the slave ignored the error and did not roll them back, thus leading to inconsistencies between master and slave.

This issue is fixed by automatically rolling back a statement that fails on the slave; however, the transaction is not rolled back unless a corresponding `ROLLBACK` statement is found in the relay log file. (Bug #46130)

References: See also: Bug #33864.

- **Replication:** When `slave_transaction_retries` is set, a statement that replicates, but is then rolled back due to a deadlock on the slave, should be retried. However, in certain cases, replication was stopped with error 1213 (`Deadlock found when trying to get lock; try restarting transaction`) instead, even when this variable was set. (Bug #45694)
- **Replication:** The binary logging behavior (and thus, the replication behavior) of `CREATE DATABASE IF NOT EXISTS`, `CREATE TABLE IF NOT EXISTS`, and `CREATE EVENT IF NOT EXISTS` was not consistent among these statements, nor with that of `DROP DATABASE IF EXISTS`, `DROP TABLE IF EXISTS`, and `DROP EVENT IF EXISTS`: A `DROP ... IF EXISTS` statement is always logged even if the database object named in the statement does not exist. However, of the `CREATE ... IF NOT EXISTS` statements, only the `CREATE EVENT IF NOT EXISTS` statement was logged when the database object named in the statement already existed.

Now, every `CREATE ... IF NOT EXISTS` statement is written to the binary log (and thus replicated), whether the database object named in the statement exists or not. For more information, see [Replication of CREATE ... IF NOT EXISTS Statements](#).

Exception. Replication and logging of `CREATE TABLE IF NOT EXISTS ... SELECT` continues to be handled according to existing rules. See [Replication of CREATE TABLE ... SELECT Statements](#), for more information.

(Bug #45574)

- **Replication:** When using statement-based replication, database-level character sets were not always honored by the replication SQL thread. This could cause data inserted on the master using `LOAD DATA` to be replicated using the wrong character set.

**Note**

This was not an issue when using row-based replication.

(Bug #45516)

- **Replication:** In some cases, a `STOP SLAVE` statement could cause the replication slave to crash. This issue was specific to MySQL on Windows or Macintosh platforms. (Bug #45238, Bug #45242, Bug #45243, Bug #46013, Bug #46014, Bug #46030)

References: See also: Bug #40796.

- **Replication:** Creating a scheduled event whose `DEFINER` clause was either set to `CURRENT_USER` or not set explicitly caused the master and the slave to become inconsistent. This issue stems from the fact that, in both cases, the `DEFINER` is set to the `CURRENT_USER` of the current thread. (On the master, the `CURRENT_USER` is the `mysqld` user; on the slave, the `CURRENT_USER` is empty.)

This behavior has been modified as follows:

- If `CURRENT_USER` is used as the `DEFINER`, it is replaced with the *value* of `CURRENT_USER` before the `CREATE EVENT` statement is written to the binary log.
- If the definer is not set explicitly, a `DEFINER` clause using the value of `CURRENT_USER` is added to the `CREATE EVENT` statement before it is written to the binary log.

(Bug #44331)

References: See also: Bug #42217.

- **Replication:** When using the statement-based logging format, the only possible safe combination of transactional and nontransactional statements within the same transaction is to perform any updates on nontransactional tables (such as `MyISAM` tables) first, before updating any transactional tables (such as those using the `InnoDB` storage engine). This is due to the fact that, although a modification made to a nontransactional table is immediately visible to other connections, the update is not immediately written to the binary log, which can lead to inconsistencies between master and slave. (Other combinations may hide a causal dependency, thus making it impossible to write statements updating nontransactional tables to the binary log in the correct order.)

However, in some cases, this situation was not handled properly, and the determination whether a given statement was safe or not under these conditions was not always correct. In particular, a multi-table update that affected both transactional and nontransactional tables or a statement modifying data in a nontransactional table having a trigger that operated on a transactional table (or the reverse) was not determined to be unsafe when it should have been.

With this fix, the following determinations regarding replication safety are made when combining updates to transactional and nontransactional tables within the same transaction in statement-based logging mode:

1. Any statement modifying data in a nontransactional table within a given transaction is considered safe if it is issued prior to any data modification statement accessing a transactional table within the same transaction.

2. A statement that updates transactional tables only is always considered safe.
3. A statement affecting both transactional and nontransactional tables within a transaction is always considered unsafe. It is not necessary that both tables be modified for this to be true; for example, a statement such as `INSERT INTO innodb_table SELECT * FROM myisam_table` is also considered unsafe.

**Note**

The current fix is valid only when using statement-based logging mode; we plan to address similar issues occurring when using the `MIXED` or `ROW` format in a future MySQL release.

(Bug #28976)

- Stack overflow checking did not account for the size of the structure stored in the heap. (Bug #46807)
- The server could crash for queries with the following elements: 1. An “impossible where” in the outermost `SELECT`; 2. An aggregate in the outermost `SELECT`; 3. A correlated subquery with a `WHERE` clause that includes an outer field reference as a top-level `WHERE` sargable predicate; (Bug #46749)
- `CREATE TABLE ... SELECT` could cause assertion failure if a table already existed with the same name and contained an `AUTO_INCREMENT` column. (Bug #46616)
- `SHOW CREATE TRIGGER` for a `MERGE` table trigger caused an assertion failure. (Bug #46614)
- In queries for which the loose index scan access method was chosen, using a condition of the form `col_name` rather than the equivalent `col_name <> 0` caused an assertion failure. (Bug #46607)
- `TRUNCATE TABLE` for a table that was opened with `HANDLER` did not close the handler and left it in an inconsistent state that could lead to a server crash. Now `TRUNCATE TABLE` for a table closes all open handlers for the table. (Bug #46456)
- A query containing a subquery in the `FROM` clause and `PROCEDURE ANALYSE()` caused a server crash. (Bug #46184)

References: See also: Bug #48293.

- Killing a query that was performing a sort could result in a memory leak. (Bug #45962)
- Truncation of `DECIMAL` values could lead to assertion failures; for example, when deducing the type of a table column from a literal `DECIMAL` value. (Bug #45261)

References: See also: Bug #48370.

- A buffer overflow could occur during handling of `IS NULL` ranges. (Bug #37044)
- `mysqladmin --wait ping` crashed on Windows systems. (Bug #35132)
- Installation of MySQL on Windows failed to set the correct location for the character set files, which could lead to `mysqld` and `mysql` failing to initialize properly. (Bug #17270)

Changes in MySQL 5.1.38 (2009-09-01)

- [InnoDB Plugin Notes](#)
- [Functionality Added or Changed](#)

- [Bugs Fixed](#)

InnoDB Plugin Notes

- As of MySQL 5.1.38, the [InnoDB Plugin](#) is included in MySQL 5.1 releases, in addition to the built-in version of [InnoDB](#) that has been included in previous releases. The version of the [InnoDB Plugin](#) in this release is 1.0.4 and is considered of Beta quality.

The [InnoDB Plugin](#) offers new features, improved performance and scalability, enhanced reliability and new capabilities for flexibility and ease of use. Among the features of the [InnoDB Plugin](#) are “Fast index creation,” table and index compression, file format management, new [INFORMATION_SCHEMA](#) tables, capacity tuning, multiple background I/O threads, and group commit.

The [InnoDB Plugin](#) is included in source and binary distributions, except RHEL3, RHEL4, SuSE 9 (x86, x86_64, ia64), and generic Linux RPM packages.

For instructions on replacing the built-in version of [InnoDB](#) with [InnoDB Plugin](#), see [Using InnoDB Plugin Instead of the Built-In InnoDB](#).

Functionality Added or Changed

- **Replication:** With statement-based logging (SBL), repeatedly calling statements that are unsafe for SBL caused a warning message to be written to the error log for each statement, and there was no way to disable this behavior. Now the server logs messages about statements that are unsafe for statement-based logging only if the [log_warnings](#) variable is greater than 0. (Bug #46265)
- The undocumented [TRANSACTIONAL](#) and [PAGE_CHECKSUM](#) keywords were removed from the grammar. (Bug #45829)
- Previously, [mysqldump](#) would not dump the [INFORMATION_SCHEMA](#) database and ignored it if it was named on the command line. Now, [mysqldump](#) will dump [INFORMATION_SCHEMA](#) if it is named on the command line. Currently, this requires that the [--skip-lock-tables](#) (or [--skip-opt](#)) option be given. (Bug #33762)
- Previously, [SELECT ... INTO OUTFILE](#) dumped column values without character set conversion, which could produce data files that cannot be imported without error if different columns used different character sets. A consequence of this is that [mysqldump](#) ignored the [--default-character-set](#) option if the [--tab](#) option was given (which causes [SELECT ... INTO OUTFILE](#) to be used to dump data.)

[INTO OUTFILE](#) now can be followed by a [CHARACTER SET](#) clause indicating the character set to which dumped values should be converted. Also, [mysqldump](#) adds a [CHARACTER SET](#) clause to the [SELECT ... INTO OUTFILE](#) statement used to dump data, so that [--default-character-set](#) is no longer ignored if [--tab](#) is given.

Other changes are that [SELECT ... INTO OUTFILE](#) enforces that [ENCLOSED BY](#) and [ESCAPED BY](#) arguments must be a single character, and [SELECT ... INTO OUTFILE](#) and [LOAD DATA INFILE](#) produce warnings if non-ASCII field or line separators are specified. (Bug #30946)

- Pluggable storage engines now can be built for Windows.
- The MySQL [euckr](#) character set now can store extended codes [81...FE][41..5A,61..7A,81..FE], which makes [euckr](#) compatible with the Microsoft [cp949](#) character set.

Bugs Fixed

- **Performance:** The table cache lock ([LOCK_open](#)) is now an adaptive mutex, which should improve performance in workloads where this lock is heavily contended. (Bug #43435)

- **Partitioning:** Attempting to create a table using an invalid or inconsistent subpartition definition caused the server to crash. An example of such a statement is shown here:

```
CREATE TABLE t2 (s1 INT, s2 INT)
PARTITION BY LIST (s1) SUBPARTITION BY HASH (s2) SUBPARTITIONS 1
(
    PARTITION p1 VALUES IN (1),
    PARTITION p2 VALUES IN (2) (SUBPARTITION p3)
);
```

(Bug #46354)

- **Partitioning:** When using a debug build of MySQL, if a query against a partitioned table having an index on one or more `DOUBLE` columns used that index, the server failed with an assertion. (Bug #45816)
- **Partitioning:** A failed `RENAME TABLE` operation on a table with user-defined partitioning left the table in an unusable state, due to only some of the table files having been renamed. (Bug #30102)
- **Replication:** When a statement that changes a nontransactional table failed, the transactional cache was flushed, causing a mismatch between the execution and logging histories. Now we avoid flushing the transactional cache unless a `COMMIT` or `ROLLBACK` is issued. (Bug #46129)

References: This issue is a regression of: Bug #43929, Bug #11752675.

- **Replication:** The internal function `get_master_version_and_clock()` (defined in `sql/slave.cc`) ignored errors and passed directly when queries failed, or when queries succeeded but the result retrieved was empty. Now this function tries to reconnect the master if a query fails due to transient network problems, and to fail otherwise. The I/O thread now prints a warning if the same system variables do not exist on master (in the event the master is a very old version of MySQL, compared to the slave.) (Bug #45214)
- **Replication:** When using the `MIXED` logging format, after creating a temporary table and performing an update that switched the logging format to `ROW`, the format switch persisted following the update. This prevented any subsequent DDL statements on temporary tables from being written to the binary log until the temporary table was dropped. (Bug #43046)

References: See also: Bug #40013. This issue is a regression of: Bug #20499.

- **Replication:** If the `--log-bin-trust-function-creators` option is not enabled, `CREATE FUNCTION` requires one of the modifiers `DETERMINISTIC`, `NO SQL`, or `READS SQL DATA`. When using statement-based mode, the execution of a stored function should follow the same rules; however, only functions defined with `DETERMINISTIC` could actually be executed. In addition, the wrong error was generated (`ER_BINLOG_ROW_RBR_TO_SBR` instead of `ER_BINLOG_UNSAFE_ROUTINE`).

Now execution of stored functions is compatible with creation in this regard; when a stored function without one of the modifiers above is executed in `STATEMENT` mode, the correct error is raised, and functions defined using `NO SQL`, `READS SQL DATA`, or both (that is, without using `DETERMINISTIC`) can be executed. (Bug #41166)

- The test suite was missing from RPM packages. (Bug #46834)
- Incorrect index optimization could lead to incorrect results or server crashes. (Bug #46454)
- The server printed warnings at startup about adjusting the value of the `max_join_size` system variable. (These were harmless, but might be seen by users as significant.) (Bug #46385)

- `mysql` did not handle backspace properly for multibyte characters. This has been fixed now if `mysql` is linked with the `readline` library. It is not fixed if `mysql` is linked with `libedit`, which does not contain the necessary support for multibyte character sets. (Bug #46310)
- After an error such as a table-full condition, `INSERT IGNORE` could cause an assertion failure for debug builds. (Bug #46075)
- An optimization that moved an item from a subquery to an outer query could cause a server crash. (Bug #46051)
- Several Valgrind warnings were corrected. (Bug #46003, Bug #46034, Bug #46042)
- `CREATE TABLE ... SELECT` could cause a server crash if no default database was selected. (Bug #45998)
- The MySQL Server crashed when performing a `REPLACE` into a `MERGE` table if there was a duplicate. (Bug #45800)
- An infinite hang and 100% CPU usage occurred after a handler tried to open a merge table.

If the command `mysqladmin shutdown` was executed during the hang, the debug server generated the following assert:

```
mysqld: table.cc:407: void free_table_share(TABLE_SHARE*): Assertion `share->ref_count == 0' failed.  
090610 14:54:04 - mysqld got signal 6 ;
```

(Bug #45781)

- For problems reading SSL files during SSL initialization, the server wrote error messages to `stderr` rather than to the error log. (Bug #45770)
- The vendor name change from MySQL AB to Sun Microsystems, Inc. in RPM packages was not handled gracefully when upgrading MySQL using an RPM package. (Bug #45534)
- A Windows Installation using the GUI installer failed with:

```
MySQL Server 5.1 Setup Wizard ended prematurely  
  
The wizard was interrupted before MySQL Server 5.1. could be completely installed.  
  
Your system has not been modified. To complete installation at another time, please run  
setup again.  
  
Click Finish to exit the wizard
```

This was due to a step in the MSI installer that could fail to execute correctly on some environments. (Bug #45418)

- Invalid memory reads could occur using the compressed client/server protocol. (Bug #45031)
- The `mysql_real_connect()` C API function only attempted to connect to the first IP address returned for a hostname. This could be a problem if a hostname mapped to multiple IP address and the server was not bound to the first one returned. Now `mysql_real_connect()` attempts to connect to all IPv4 or IPv6 addresses that a domain name maps to. (Bug #45017)

References: See also: Bug #47757.

- Invalid input could cause invalid memory reads by the parser. (Bug #45010)

- Some files in an AIX `tar` file distribution unpacked with incorrect permissions. (Bug #44647)
- For debug builds, executing a stored procedure as a prepared statement could sometimes cause an assertion failure. (Bug #44521)
- Using `mysql_stmt_execute()` to call a stored procedure could cause a server crash. (Bug #44495)
- Creating a new instance after previously removing an instance failed to complete the installation properly because the security settings could not be applied correctly. (Bug #44428)
- `mysqslslap` ignored the `--csv` option if it was given without an argument. (Bug #44412)
- Enabling the event scheduler from within the file specified by `--init-file` caused a server crash. (Bug #43587)
- The server did not always check the return value of calls to the `hash_init()` function. (Bug #43572)
- `mysqladmin --count=X --sleep=Y` incorrectly delayed `Y` seconds after the last iteration before exiting. (Bug #42639)
- A test for stack growth failed on some platforms, leading to server crashes. (Bug #42213)

References: See also: Bug #62856.

- `mysqladmin` did not have enough space allocated for tracking all variables when using `--vertical` or `--relative` with `extended-status`. (Bug #40395)
- Partitioning a log table caused a server crash. (Bug #40281)
- When using quick access methods to search for rows in `UPDATE` and `DELETE` statements, there was no check whether a fatal error had already been sent to the client while evaluating the quick condition. Consequently, a false OK (following the error) was sent to the client, causing the error to be incorrectly transformed into a warning. (Bug #40113)
- `SHOW PROCESSLIST` could access freed memory of a stored procedure run in a concurrent session. (Bug #38816)
- During installation on Windows, the MySQL Instance Configuration Wizard window could be opened at a size too small to be usable. (Bug #38723)
- `make_binary_distribution` did not always generate correct distribution names. (Bug #37808)
- The server crashed when executing a prepared statement containing a duplicated `MATCH()` function call in the select list and `ORDER BY` clause; for example, `SELECT MATCH(a) AGAINST('test') FROM t1 ORDER BY MATCH(a) AGAINST('test')`. (Bug #37740)
- The output of `mysqldump --tab` for views included a `DROP TABLE` statement without the `IF EXISTS` qualifier. (Bug #37377)
- `mysql_upgrade` silently ignored the `--basedir` and `--datadir` options, which it accepts for backward compatibility. Now it prints a warning. (Bug #36558)
- `mysqlimport` was not always compiled correctly to enable thread support, which is required for the `--use-threads` option. (Bug #32991)
- `mysqlcheck` failed to fix table names when the `--fix-table-names` and `--all-in-1` options were both specified. (Bug #31821)
- If the MySQL server was killed without the PID file being removed, attempts to stop the server with `mysql.server stop` waited 900 seconds before giving up. (Bug #31785)

- When performing an installation on Windows using the GUI installer, the installer failed to wait long enough during installation for the MySQL service to be installed, which would cause the installation to fail and may cause security settings, such as the `root` password to not be applied correctly. (Bug #30525)
- `mysql` included extra spaces at the end of some result set lines. (Bug #29622)
- The `mysql` client inconsistently handled NUL bytes in column data in various output formats. (Bug #28203)
- `mysqlimport` did not correctly quote and escape table identifiers and file names. (Bug #28071)
- When installing the Windows service, using quotation marks around command-line configuration parameters could cause the quotation marks to be incorrectly placed around the entire command-line option, and not just the value. (Bug #27535)
- If the `mysql` client was built with the `readline` library and the `.inputrc` file mapped `Space` to the `magic-space` function, it became impossible to enter spaces. (Bug #27439)
- If `InnoDB` reached its limit on the number of concurrent transactions (1023), it wrote a descriptive message to the error log but returned a misleading error message to the client, or an assertion failure occurred. (Bug #18828)

References: See also: Bug #46672.

Changes in MySQL Enterprise 5.1.37sp1 [QSP] (2009-10-10)

This is a *Service Pack* release of the MySQL Enterprise Server 5.1.

Bugs Fixed

- The test suite was missing from RPM packages. (Bug #46834)
- The server could crash for queries with the following elements: 1. An “impossible where” in the outermost `SELECT`; 2. An aggregate in the outermost `SELECT`; 3. A correlated subquery with a `WHERE` clause that includes an outer field reference as a top-level `WHERE` sargable predicate; (Bug #46749)
- `SHOW CREATE TRIGGER` for a `MERGE` table trigger caused an assertion failure. (Bug #46614)
- Incorrect index optimization could lead to incorrect results or server crashes. (Bug #46454)
- A query containing a subquery in the `FROM` clause and `PROCEDURE ANALYSE()` caused a server crash. (Bug #46184)

References: See also: Bug #48293.

- `CREATE TABLE ... SELECT` could cause a server crash if no default database was selected. (Bug #45998)
- A Windows Installation using the GUI installer failed with:

```
MySQL Server 5.1 Setup Wizard ended prematurely
```

```
The wizard was interrupted before MySQL Server 5.1. could be completely installed.
```

```
Your system has not been modified. To complete installation at another time, please run setup again.
```

```
Click Finish to exit the wizard
```


This was due to a step in the MSI installer that could fail to execute correctly on some environments. (Bug #45418)

- For debug builds, executing a stored procedure as a prepared statement could sometimes cause an assertion failure. (Bug #44521)
- Using `mysql_stmt_execute()` to call a stored procedure could cause a server crash. (Bug #44495)

Changes in MySQL 5.1.37 (2009-07-13)

- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

Functionality Added or Changed

- **Important Change; Replication:** `RESET MASTER` and `RESET SLAVE` now reset the values shown for `Last_IO_Error`, `Last_IO_Errno`, `Last_SQL_Error`, and `Last_SQL_Errno` in the output of `SHOW SLAVE STATUS`. (Bug #44270)

References: See also: Bug #34654.

Bugs Fixed

- **Security Fix; Partitioning:** Accessing a table having user-defined partitioning when the server SQL mode included `ONLY_FULL_GROUP_BY` caused the MySQL server to crash. For example, the following sequence of statements crashed the server:

```
DROP TABLE IF EXISTS t1;

SET SESSION sql_mode='ONLY_FULL_GROUP_BY';

CREATE TABLE t1 (id INT, KEY(id))
  PARTITION BY HASH(id) PARTITIONS 2;
```

(Bug #45807)

- **Security Fix:** The `strxnmov()` library function could write a null byte after the end of the destination buffer. (Bug #44834)
- **Performance:** With `InnoDB` tables, MySQL used a less-selective secondary index to avoid a filesort even if a prefix of the primary key was much more selective.

The fix for this problem might cause other queries to run more slowly. (Bug #45828)

- **Important Change; Replication:** When using `STATEMENT` or `MIXED` binary logging format, a statement that changes both nontransactional and transactional tables must be written to the binary log whenever there are changes to nontransactional tables. This means that the statement goes into the binary log even when the changes to the transactional tables fail. In particular, in the event of a failure such statement is annotated with the error number and wrapped inside a pair of `BEGIN` and `ROLLBACK` statements.

On the slave, while applying the statement, it is expected that the same failure and the rollback prevent the transactional changes from persisting. However, statements that fail due to concurrency issues such as deadlocks and timeouts are logged in the same way, causing the slave to stop since the statements are applied sequentially by the SQL thread.

To address this issue, we ignore concurrency failures on the slave. Specifically, the following failures are now ignored: `ER_LOCK_WAIT_TIMEOUT`, `ER_LOCK_DEADLOCK`, and `ER_XA_RBDEADLOCK`. (Bug #44581)

- **Partitioning:** Truncating a partitioned `MyISAM` table did not reset the `AUTO_INCREMENT` value. (Bug #35111)
- **Replication:** The `SHOW SLAVE STATUS` connection thread competed with the slave SQL thread for use of the error message buffer. As a result, the connection thread sometimes received incomplete messages. This issue was uncovered with `valgrind` when message strings were passed without `NULL` terminators, causing the error `Conditional jump or move depends on uninitialized value(s)`. (Bug #45511)

References: See also: Bug #43076.

- **Replication:** For replication of a stored procedure that uses the `gbk` character set, the result on the master and slave differed. (Bug #45485)
- **Replication:** The internal function `purge_relay_logs()` did not propagate an error occurring in another internal function `count_relay_log_space()`. (Bug #44115)
- **Replication:** Large transactions and statements could corrupt the binary log if the size of the cache (as set by `max_binlog_cache_size`) was not large enough to store the changes.

Now, for transactions that do not fit into the cache, the statement is not logged, and the statement generates an error instead.

For nontransactional changes that do not fit into the cache, the statement is also not logged—an incident event is logged after committing or rolling back any pending transaction, and the statement then raises an error.



Note

If a failure occurs before the incident event is written the binary log, the slave does not stop, and the master does not report any errors.

(Bug #43929, Bug #11752675)

References: See also: Bug #37148, Bug #11748696, Bug #46166, Bug #11754544.

- **Replication:** The `--database` option for `mysqlbinlog` was ignored when using the row-based logging format. (Bug #42941)
- **Replication:** Statements using `LIMIT` generated spurious `Statement is not safe to log in statement format` warnings in the error log, causing the log to grow rapidly in size. (Bug #42851)

References: See also: Bug #46265, Bug #42415. This issue is a regression of: Bug #34768.

- **Replication:** Shutting down the slave while executing `FLUSH LOGS`, `CHANGE MASTER TO`, or `STOP SLAVE` could sometimes cause it to crash. (Bug #38240)
- **Replication:** When reading a binary log that was in use by a master or that had not been properly closed (possibly due to a crash), the following message was printed: `Warning: this binlog was not closed properly. Most probably mysqld crashed writing it.` This message did not take into account the possibility that the file was merely in use by the master, which caused some users concern who were not aware that this could happen.

To make this clear, the original message has been replaced with `Warning: this binlog is either is use or was not closed properly`. (Bug #34687)

- The server crashed if evaluation of `GROUP_CONCAT(... ORDER BY)` required allocation of a sort buffer but allocation failed. (Bug #46080)
- When creating tables using the `IBMDB2I` storage engine with the `ibmdb2i_create_index_option` option set to 1, creating an `IBMDB2I` table with a primary key should produce an additional index that uses EBCDIC hexadecimal sorting, but this index was not created. (Bug #45983)
- The server crashed for attempts to use `REPLACE` or `INSERT ... ON DUPLICATE KEY UPDATE` with a view defined using a join. (Bug #45806)
- Some collations were causing `IBMDB2I` to report inaccurate key range estimations to the optimizer for `LIKE` clauses that select substrings. This can be seen by running `EXPLAIN`. This problem primarily affects multibyte and unicode character sets. (Bug #45803)
- Invalid memory reads and writes were generated when altering merge and base tables. This could lead to a crash or Valgrind errors:

```
==28038== Invalid write of size 1
at: memset (mc_replace_strmem.c:479)
by: myrg_attach_children (myrg_open.c:433)
by: ha_myisammrg::attach_children() (ha_myisammrg.cc:546)
by: ha_myisammrg::extra(ha_extra_function) (ha_myisammrg.cc:944)
by: attach_merge_children(TABLE_LIST*) (sql_base.cc:4147)
by: open_tables(THD*, TABLE_LIST**, unsigned*, unsigned) (sql_base.cc:4709)
by: open_and_lock_tables_derived(THD*, TABLE_LIST*, bool) (sql_base.cc:4977)
by: open_n_lock_single_table (mysql_priv.h:1550)
by: mysql_alter_table(sql_table.cc:6428)
by: mysql_execute_command(THD*) (sql_parse.cc:2860)
by: mysql_parse(THD*, char const*, unsigned, char const**) (sql_parse.cc:5933)
by: dispatch_command (sql_parse.cc:1213)
```

(Bug #45796)

- Inserting data into a table using the `macce` character set with the `IBMDB2I` storage engine failed. (Bug #45793)
- There was a race condition when changing `innodb_commit_concurrency` at runtime to the value `DEFAULT`. (Bug #45749)

References: See also: Bug #42101.

- Performing an empty XA transaction caused the server to crash for the next XA transaction. (Bug #45548)
- `SHOW CREATE TRIGGER` requires the `TRIGGER` privilege but was not checking privileges. (Bug #45412)
- An assertion failure could occur if `InnoDB` tried to unlock a record when the clustered index record was unknown. (Bug #45357)
- `--enable-plugin_name` options (for example, `--enable-innodb`) did not work correctly. (Bug #45336)

References: See also: Bug #19027.

- If `autocommit` was enabled, `InnoDB` did not roll back `DELETE` or `UPDATE` statements if the statement was killed. (Bug #45309)

- The optimizer mishandled “impossible range” conditions and returned empty results due to an uninitialized variable. (Bug #45266)
- Use of `DECIMAL` constants with more than 65 digits in `CREATE TABLE ... SELECT` statements led to spurious errors or assertion failures. (Bug #45262)
- The `mysql` client could misinterpret some character sequences as commands under some circumstances. (Bug #45236)
- Use of `CONVERT()` with an empty `SET` value could cause an assertion failure. (Bug #45168)
- `InnoDB` recovery could hang due to redo logging of doublewrite buffer pages. (Bug #45097)
- When reading binary data, the concatenation function for geometry data collections did not rigorously check for available data, leading to invalid reads and server crashes. (Bug #44684)
- If an error occurred during the creation of a table (for example, the table already existed) having an `AUTO_INCREMENT` column and a `BEFORE` trigger that used the `INSERT ... SELECT` construct, an internal flag was not reset properly. This led to a crash the next time the table was opened again. (Bug #44653)
- `configure.in` contained references to literal instances of `nm` and `libc`, rather than to variables parameterized for the proper values on the current platform. (Bug #42721)
- `configure.in` did not properly check for the `pthread_setschedprio()` function. (Bug #42599)
- `SHOW ERRORS` returned an empty result set after an attempt to drop a nonexistent table. (Bug #42364)
- A workaround for a Sun Studio bug was instituted. (Bug #41710)
- For queries with a sufficient number of subqueries in the `FROM` clause of this form:

```
SELECT * FROM (SELECT 1) AS t1,  
              (SELECT 2) AS t2,  
              (SELECT 3) AS t3, ...
```

The query failed with a `Too high level of nesting for select` error, as though the query had this form:

```
SELECT * FROM (SELECT 1 FROM (SELECT 2 FROM (SELECT 3 FROM ...
```

(Bug #41156)

- Some `UPDATE` statements that affected no rows returned a rows-affected count of one. (Bug #40565)
- Valgrind warnings that occurred for `SHOW TABLE STATUS` with `InnoDB` tables were silenced. (Bug #38479)
- In the `mysql` client, if the server connection was lost during repeated `status` commands, the client failed to detect this and command output would be inconsistent. (Bug #37274)
- A Valgrind error during subquery execution was corrected. (Bug #36995)
- When invoked to start multiple server instances, `mysqld_multi` sometimes failed to start them all due to not changing location into the base directory for each instance. (Bug #36654)
- Rows written to the slow query log could have an indeterminate `Rows_examined` value due to improper initialization. (Bug #34002)

- Renaming a column that appeared in a foreign key definition did not update the foreign key definition with the new column name. (Bug #21704)

Changes in MySQL 5.1.36 (2009-06-16)

- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

Functionality Added or Changed

- **Important Change; Replication:** Previously, incident log events were represented as comments in the output from `mysqlbinlog`, making them effectively silent when playing back the binary log.

(An incident log event represents an incident that could cause the contents of the database to change without that event being recorded in the binary log.)

This meant that, if the SQL were applied to a server, it could potentially lead to the master and the slave having different data. To make it possible to handle incident log events without breaking applications that expect the previous behavior, the nonsense statement `RELOAD DATABASE` is added to the SQL output for that incident log event, which causes an error.

To use this functionality currently requires hand editing of the dump file and handling of each case on an individual basis by a database administrator before applying the output to a server. (Bug #44442)

- `mysql_upgrade` now displays a message indicating the connection parameters it uses when invoking `mysqlcheck`. (Bug #44638)
- The time zone tables available at <http://dev.mysql.com/downloads/timezones.html> have been updated. These tables can be used on systems such as Windows or HP-UX that do not include zoneinfo files. (Bug #39923)
- The `mysqltest` program now has a `move_file from_file to_file` command for renaming files. This should be used in test cases rather than invoking an external command that might be platform specific. (Bug #39542)
- The maximum value for `max_binlog_cache_size` has been increased from $2^{32} - 1$ to $2^{64} - 1$ (even on 32-bit platforms), which enables transactions 4GB and larger to be performed when binary logging is enabled. (Bug #10206)

Bugs Fixed

- **Security Fix:** The server crashed if an account with the `CREATE ROUTINE` privilege but not the `EXECUTE` privilege attempted to create a stored procedure. (Bug #44798)
- **Security Fix:** The server crashed if an account without the proper privileges attempted to create a stored procedure. (Bug #44658)
- **Security Fix:** Four potential format string vulnerabilities were fixed (discovered by the Veracode code analysis). (Bug #44166)
- **Performance:** The `InnoDB` adaptive hash latch is released (if held) for several potentially long-running operations. This improves throughput for other queries if the current query is removing a temporary table, changing a temporary table from memory to disk, using `CREATE TABLE ... SELECT`, or performing a `MyISAM` repair on a table used within a transaction. (Bug #32149)

- **Incompatible Change:** The server can load plugins under the control of startup options. For example, many storage engines can be built in pluggable form and loaded when the server starts. In the following descriptions, `plugin_name` stands for a plugin name such as `innodb`.

Previously, plugin options were handled like other boolean options (see [Program Option Modifiers](#)). That is, any of these options enabled the plugin:

```
--plugin_name
--plugin_name=1
--enable-plugin_name
```

And these options disabled the plugin:

```
--plugin_name=0
--disable-plugin_name
--skip-plugin_name
```

However, use of a boolean option for plugin loading did not provide control over what to do if the plugin failed to start properly: Should the server exit, or start with the plugin disabled? The actual behavior has been that the server starts with the plugin disabled, which can be problematic. For example, if `InnoDB` fails to start, existing `InnoDB` tables become inaccessible, and attempts to create new `InnoDB` tables result in tables that use the default storage engine unless the `NO_ENGINE_SUBSTITUTION` SQL mode has been enabled to cause an error to occur instead.

Now, there is a change in the options used to control plugin loading, such that they have a tristate format:

- `--plugin_name=OFF`

Do not enable the plugin.

- `--plugin_name [=ON]`

Enable the plugin. If plugin initialization fails, start the server anyway, but with the plugin disabled. Specifying the option as `--plugin_name` without a value also enables the plugin.

- `--plugin_name=FORCE`

Enable the plugin. If plugin initialization fails, do not start the server. In other words, force the server to run with the plugin or not at all.

The values `OFF`, `ON`, and `FORCE` are not case sensitive.

Suppose that `CSV` and `InnoDB` have been built as pluggable storage engines and that you want the server to load them at startup, subject to these conditions: The server is permitted to run if `CSV` initialization fails, but must require that `InnoDB` initialization succeed. To accomplish that, use these lines in an option file:

```
[mysqld]
csv=ON
innodb=FORCE
```

This change is incompatible with the previous implementation if you used options of the form `--plugin_name=0` or `--plugin_name=1`, which should be changed to `--plugin_name=OFF` or `--plugin_name=ON`, respectively.

`--enable-plugin_name` is still supported and is the same as `--plugin_name=ON`. `--disable-plugin_name` and `--skip-plugin_name` are still supported and are the same as `--plugin_name=OFF`. (Bug #19027)

References: See also: Bug #45336.

- **Important Change; Replication:** `BEGIN`, `COMMIT`, and `ROLLBACK` statements are no longer affected by `--replicate-do-db` or `--replicate-ignore-db` rules. (Bug #43263)
- **Partitioning:** Queries using `DISTINCT` on multiple columns or `GROUP BY` on multiple columns did not return correct results with partitioned tables. (Bug #44821)

References: See also: Bug #41136.

- **Replication:** When using row-based logging, the length of an event for which the field metadata exceeded 255 bytes in size was incorrectly calculated. This could lead to corruption of the binary log, or cause the server to hang. (Bug #42749)

References: See also: Bug #44548, Bug #44672, Bug #44752.

- **Replication:** The warning `Statement is not safe to log in statement format`, issued in situations when it cannot be determined that a statement or other database event can be written reliably to the binary log using the statement-based format, has been changed to `Statement may not be safe to log in statement format`. (Bug #42415)
- **Replication:** The `Query_log_event` used by replication to transfer a query to the slave has been refactored. `Query_log_event` also stores and sends the error code resulting from the execution since it, in some cases, is necessary to execute the statement on the slave as well, which should result in the same error code. The `Query_log_event` constructor previously worked out for itself the error code using a complex routine, the result of which was often set aside within the constructor itself. This was also involved with at least 2 known bugs relating to invalid errors, and taken as a clear sign that the constructor was not well-designed and needed to be re-written. (Bug #41948)

References: See also: Bug #37145.

- **Replication:** When stopping and restarting the slave while it was replicating temporary tables, the slave server could crash or raise an assertion failure. This was due to the fact that, although temporary tables were saved between slave thread restarts, the reference to the thread being used (`table->in_use`) was not being properly updated when restarting, continuing to reference the old thread instead of the new one. This issue affected statement-based replication only. (Bug #41725)
- A separator was added between the time tag and the thread ID in the general query log file. (Bug #45387)
- The combination of `MIN()` or `MAX()` in the select list with `WHERE` and `GROUP BY` clauses could lead to incorrect results. (Bug #45386)
- Linker failures with `libmysqld` on VC++ 2008 were fixed. (Bug #45326)
- Compiler warnings on Mac OS X were fixed. (Bug #45286)
- Running a `SELECT` query over an `IBMDB2I` table using the `cp1250` character set would produce an error

```
ibmdb2i error 2027: Error converting single-byte sort sequence to UCS-2
```

(Bug #45197)

- Use of `ROUND()` on a `LONGTEXT` or `LONGBLOB` column of a derived table could cause a server crash. (Bug #45152)
- `DROP USER` could fail to drop all privileges for an account if the `PAD_CHAR_TO_FULL_LENGTH` SQL mode was enabled. (Bug #45100)
- `GROUP BY` on a `constant` (single-row) `InnoDB` table joined to other tables caused a server crash. (Bug #44886)
- `ALTER TABLE` on a view crashed the server. (Bug #44860)
- When using partitioning with the `IBMDB2I` storage engine, the engine could report that a valid character set was not supported. (Bug #44856)
- Running queries on tables with the `IBMDB2I` storage engine using the `utf8` character would fail when using the 64-bit version of MySQL. (Bug #44811)
- Index Merge followed by a filesort could result in a server crash if `sort_buffer_size` was not large enough for all sort keys. (Bug #44810)

References: See also: Bug #40974.

- `UNCOMPRESSED_LENGTH()` returned a garbage result when passed a string shorter than 5 bytes. Now `UNCOMPRESSED_LENGTH()` returns `NULL` and generates a warning. (Bug #44796)
- Several Valgrind warnings were silenced. (Bug #44774, Bug #44792)
- Selecting `RAND(N)` function where `N` is a column of a `constant` table (table with a single row) failed with a `SIGFPE` signal. (Bug #44768)
- The `PASSWORD()` and `OLD_PASSWORD()` functions could read memory outside of an internal buffer when used with `BLOB` arguments. (Bug #44767)
- Conversion of a string to a different character set could use the same buffer for input and output, leading to incorrect results or warnings. (Bug #44743, Bug #44766)
- `mysqld_safe` could fail to find the `logger` program. (Bug #44736)
- Code that optimized a read-only XA transaction failed to reset the XID once the transaction was no longer active. (Bug #44672)
- A Valgrind warning related to transaction processing was silenced. (Bug #44664)
- Some Perl scripts in AIX packages contained an incorrect path to the `perl` executable. (Bug #44643)
- When creating tables using the `IBMDB2I` storage engine, the `RCDFMT` (record format) that would be applied to the corresponding files within the IBM i would be set according to the table name. During whole table operations, the name could get modified to a value inconsistent with the table name. In addition, the record format would be inconsistent compared to the file content. The `IBMDB2I` storage engine now adds an explicit `RCDFMT` clause to the `CREATE TABLE` statement passed down to the DB2 storage engine layer. (Bug #44610)
- `innochecksum` could incorrectly determine the input file name from the arguments. (Bug #44484)
- Incorrect time was reported at the end of `mysqldump` output. (Bug #44424)
- Caching of `GROUP BY` expressions could lead to mismatches between compile-time and runtime calculations and cause a server crash. (Bug #44399)

- Lettercase conversion in multibyte `cp932` or `sjis` character sequences could produce incorrect results. (Bug #44352)
- `InnoDB` was missing `DB_ROLL_PTR` information in Table Monitor `COLUMNS` output. (Bug #44320)
- Assertion failure could occur for duplicate-key errors in `INSERT INTO ... SELECT` statements. (Bug #44306)
- Trying to use an unsupported character set on an `IBMDB2I` table would produce DB2 error 2501 or 2511. The error has been updated to produce Error 2504 (Character set is unsupported). (Bug #44232)
- On 64-bit Windows systems, `myisamchk` did not handle `key_buffer_size` values larger than 4GB. (Bug #43940)
- For user-defined `utf8` collations, attempts to store values too long for a column could cause a server crash. (Bug #43827)
- Invalidation of query cache entries due to table modifications could cause threads to hang inside the query cache with state “freeing items”. (Bug #43758)
- `EXPLAIN EXTENDED` could crash for `UNION` queries in which the last `SELECT` was not parenthesized and included an `ORDER BY` clause. (Bug #43612)
- Multiple-table updates for `InnoDB` tables could produce unexpected results. (Bug #43580)
- If the client lost the connection to the MySQL server after `mysql_stmt_prepare()`, the first call to `mysql_stmt_execute()` returned an error (as expected) but consecutive calls to `mysql_stmt_execute()` or `mysql_stmt_close()` crashed the client. (Bug #43560)
- For `DELETE` statements with `ORDER BY var`, where `var` was a global system variable with a `NULL` value, the server could crash. (Bug #42778)
- Builds linked against OpenSSL had a memory leak in association with use of X509 certificates. (Bug #42158)
- There was a race condition when changing `innodb_commit_concurrency` at runtime from zero to nonzero or from nonzero to zero. Now this variable cannot be changed at runtime from zero to nonzero or vice versa. The value can still be changed from one nonzero value to another. (Bug #42101)

References: See also: Bug #45749.

- `SELECT ... INTO @var` could produce values different from `SELECT ...` without the `INTO` clause. (Bug #42009)
- `mysql_zap` did not work on Mac OS X. (Bug #41883)
- A crash occurred due to a race condition between the merge table and `table_cache` evictions.

```
00000001403C452F  mysqlld.exe!memcpy() [memcpy.asm:151]
00000001402A275F  mysqlld.exe!ha_myisammrg::info() [ha_myisammrg.cc:854]
00000001402A2471  mysqlld.exe!ha_myisammrg::attach_children() [ha_myisammrg.cc:488]
00000001402A2788  mysqlld.exe!ha_myisammrg::extra() [ha_myisammrg.cc:863]
000000014015FC5D  mysqlld.exe!attach_merge_children() [sql_base.cc:4135]
000000014016A4C1  mysqlld.exe!open_tables() [sql_base.cc:4697]
000000014016A898  mysqlld.exe!open_and_lock_tables_derived() [sql_base.cc:4956]
000000014018BB54  mysqlld.exe!mysql_insert() [sql_insert.cc:613]
000000014019EDD3  mysqlld.exe!mysql_execute_command() [sql_parse.cc:3066]
00000001401A2F06  mysqlld.exe!mysql_parse() [sql_parse.cc:5791]
00000001401A3C1A  mysqlld.exe!dispatch_command() [sql_parse.cc:1202]
```

```
00000001401A4CD7  mysql!do_command()[sql_parse.cc:857]
0000000140246327  mysql!handle_one_connection()[sql_connect.cc:1115]
00000001402B82C5  mysql!pthread_start()[my_winthread.c:85]
00000001403CAC37  mysql!_callthreadstart()[thread.c:295]
00000001403CAD05  mysql!_threadstart()[thread.c:275]
0000000077D6B69A  kernel32.dll!BaseThreadStart()
Trying to get some variables.
Some pointers may be invalid and cause the dump to abort...
```

(Bug #41212)

- Shared-memory connections did not work in Vista if `mysqld` was started from the command line. (Bug #41190)
- For views created with a column list clause, column aliases were not substituted when selecting through the view using a `HAVING` clause. (Bug #40825)
- A multiple-table `DELETE` involving a table self-join could cause a server crash. (Bug #39918)
- Creating an `InnoDB` table with a comment containing a '#' character caused foreign key constraints to be omitted. (Bug #39793)
- `ALTER TABLE` neglected to preserve `ROW_FORMAT` information from the original table, which could cause subsequent `ALTER TABLE` and `OPTIMIZE TABLE` statements to lose the row format for `InnoDB` tables. (Bug #39200)
- The `mysql` option `--ignore-spaces` was nonfunctional. (Bug #39101)
- If a query was such as to produce the error `1054 Unknown column '...' in 'field list'`, using `EXPLAIN EXTENDED` with the query could cause a server crash. (Bug #37362)
- In the `mysql` client, using a default character set of `binary` caused internal commands such as `DELIMITER` to become case sensitive. (Bug #37268)
- `mysqldump --tab` dumped triggers to `stdout` rather than to the `.sql` file for the corresponding table. (Bug #34861)
- If the `MYSQL_HISTFILE` environment variable was set to `/dev/null`, the `mysql` client overwrote the `/dev/null` device file as a normal file. (Bug #34224)
- `mysqld_safe` mishandled certain parameters if they contained spaces. (Bug #33685)
- `mysqladmin kill` did not work for thread IDs larger than 32 bits. (Bug #32457)
- Several client programs failed to interpret `--skip-password` as “send no password.” (Bug #28479)
- Output from `mysql --html` did not encode the '<', '>', or '&' characters. (Bug #27884)
- `mysql_convert_table_format` did not prevent conversion of tables to `MEMORY` or `BLACKHOLE` tables, which could result in data loss. (Bug #27149)

Changes in MySQL 5.1.35 (2009-05-13)

- [Windows Notes](#)
- [Bugs Fixed](#)

Windows Notes

- This MySQL release has two known outstanding issues for Windows:

- The `.msi` installer does not detect an existing `root` password on the initial configuration attempt. To work around this, install and configure MySQL as normal, but skip any changes to security. (There is a check box that enables this on the security screen of the configuration wizard.) Then check your settings:
 - If the old `root` password and security settings are okay, you are done and can proceed to use MySQL.
 - Otherwise, reconfigure with the wizard and make any changes on the second configuration attempt. The wizard will properly prompt for the existing `root` password and permit changes to be made.

This issue has been filed as Bug #45200 for correction in a future release.

- The Windows configuration wizard permits changes to `InnoDB` settings during a reconfiguration operation. For an upgrade, this may cause difficulties. To work around this, use one of the following alternatives:
 - Do not change `InnoDB` settings.
 - Copy files from the old `InnoDB` location to the new one.

This issue has been filed as Bug #45201 for correction in a future release.

Bugs Fixed

- **Performance:** `InnoDB` uses random numbers to generate dives into indexes for calculating index cardinality. However, under certain conditions, the algorithm did not generate random numbers, so `ANALYZE TABLE` did not update cardinality estimates properly. A new algorithm has been introduced with better randomization properties, together with a system variable, `innodb_use_legacy_cardinality_algorithm`, that controls which algorithm to use. The default value of the variable is 1 (`ON`), to use the original algorithm for compatibility with existing applications. The variable can be set to 0 (`OFF`) to use the new algorithm with improved randomness. (Bug #43660)
- **Performance:** If the character set for a column being compared was neither the default server character set nor `latin1`, `InnoDB` was slower than necessary due to excessive contention for a character set mutex.

As a workaround for earlier versions, set the default server character set to the character set other than `latin1` that is most often used in indexed columns. (Bug #42649)

- **Important Change; Replication:** The transactional behavior of `STOP SLAVE` has changed. Formerly, it took effect immediately, even inside a transaction; now, it waits until the current replication event group (if any) has finished executing, or until the user issues a `KILL QUERY` or `KILL CONNECTION` statement.

This was done to solve the problem encountered when replication was stopped while a nontransactional slave was replicating a transaction on the master. (It was impossible to roll back a mixed-engines transaction when one of the engines was nontransactional, which meant that the slave could not safely re-apply any transaction that had been interrupted by `STOP SLAVE`.) (Bug #319, Bug #38205)

References: See also: Bug #43217.

- **Partitioning:** When a value was equal to a `PARTITION ... VALUES LESS THAN (value)` value other than `MAXVALUE`, the corresponding partition was not pruned. (Bug #42944)
- **Replication:** Unrelated errors occurring during the execution of `RESET SLAVE` could cause the slave to crash. (Bug #44179)

- **Replication:** The `--slave-skip-errors` option had no effect when using row-based logging format. (Bug #39393)
- **Replication:** The following errors were not correctly reported:
 - Failures during slave thread initialization
 - Failures while initializing the relay log position (immediately following the starting of the slave thread)
 - Failures while processing queries passed through the `--init_slave` option.

Information about these types of failures can now be found in the output of `SHOW SLAVE STATUS`. (Bug #38197)

- **Replication:** Killing the thread executing a DDL statement, after it had finished its execution but before it had written the binary log event, caused the error code in the binary log event to be set (incorrectly) to `ER_SERVER_SHUTDOWN` or `ER_QUERY_INTERRUPTED`, which caused replication to fail. (Bug #37145)

References: See also: Bug #27571, Bug #22725.

- **Replication:** Column aliases used inside subqueries were ignored in the binary log. (Bug #35515)
- Valgrind warnings for the `DECODE()`, `ENCRYPT()`, and `FIND_IN_SET()` functions were corrected. (Bug #44358, Bug #44365, Bug #44367)
- On Windows, entries for `build-vs9.bat` and `build-vs9_x64.bat` were missing in `win/Makefile.am`. (Bug #44353)
- Incomplete cleanup of `JOIN_TAB::select` during the filesort of rows for a `GROUP BY` clause inside a subquery caused a server crash. (Bug #44290)
- Not all lock types had proper descriptive strings, resulting in garbage output from `mysqladmin debug`. (Bug #44164)
- Use of `HANDLER` statements with `INFORMATION_SCHEMA` tables caused a server crash. Now `HANDLER` is prohibited with such tables. (Bug #44151)
- MySQL Server permitted the creation of a merge table based on views but crashed when attempts were made to read from that table. The following example demonstrates this:

```
#Create a test table
CREATE TABLE tmp (id int, c char(2));

#Create two VIEWS upon it
CREATE VIEW v1 AS SELECT * FROM tmp;
CREATE VIEW v2 AS SELECT * FROM tmp;

#Finally create a MERGE table upon the VIEWS
CREATE TABLE merge (id int, c char(2))
ENGINE=MERGE UNION(v1, v2);

#Reading from the merge table lead to a crash
SELECT * FROM merge;
```

The final statement generated the crash. (Bug #44040)

- Some schema names longer than 8 characters were not supported by `IBMDB2I`. The engine has been updated to permit digits and underscore characters to be used in names longer than 8 characters. (Bug #44025)

- In some circumstances, when a table is created with the `IBMDB2I` engine, the `CREATE TABLE` statement will return successfully but the table will not exist. (Bug #44022)
 - The `ucs2_swedish_ci` and `utf8_swedish_ci` collations did not work with indexes using the `IBMDB2I` storage engine. Support is now provided for MySQL when running on IBM i 6.1 or higher. (Bug #44020)
 - Invoking `SHOW TABLE STATUS` from within a stored procedure could cause a `Packets out of order` error. (Bug #43962)
 - `myisamchk` could display a negative `Max keyfile length` value. (Bug #43950)
 - On 64-bit systems, a `key_buffer_size` value larger than 4GB could cause `MyISAM` index corruption. (Bug #43932)
 - `mysqld_multi` incorrectly passed `--no-defaults` to `mysqld_safe`. (Bug #43876)
 - `SHOW VARIABLES` did not properly display the value of `slave_skip_errors`. (Bug #43835)
 - On Windows, a server crash occurred for attempts to insert a floating-point value into a `CHAR` column with a maximum length less than the converted floating-point value length. (Bug #43833)
 - Incorrect initialization of `MyISAM` table indexes could cause incorrect query results. (Bug #43737)
 - `libmysqld` crashed when it was reinitialized. (Bug #43706, Bug #44091)
 - `UNION` of floating-point numbers did unnecessary rounding. (Bug #43432)
 - `ALTER DATABASE ... UPGRADE DATA DIRECTORY NAME` failed when the database contained views. (Bug #43385)
 - Certain statements might open a table and then wait for an impending global read lock without noticing whether they hold a table being waiting for by the global read lock, causing a hang. Affected statements are `SELECT ... FOR UPDATE`, `LOCK TABLES ... WRITE`, `TRUNCATE TABLE`, and `LOAD DATA INFILE`. (Bug #43230)
 - Using an XML function such as `ExtractValue()` more than once in a single query could produce erroneous results. (Bug #43183)
- References: See also: Bug #43937.
- Full-text prefix searches could hang the connection and cause 100% CPU consumption. (Bug #42907)
 - Incorrect elevation of warning messages to error messages for unsafe statements caused a server crash. (Bug #42640)
 - `CHECK TABLE` suggested use of `REPAIR TABLE` for corrupt tables for storage engines not supported by `REPAIR TABLE`. Now `CHECK TABLE` suggests that the user dump and reload the table. (Bug #42563)
 - Compressing a table with the `myisampack` utility caused the server to produce Valgrind warnings when it opened the table. (Bug #41541)
 - For a `MyISAM` table with `DELAY_KEY_WRITE` enabled, the index file could be corrupted without the table being marked as crashed if the server was killed. (Bug #41330)
 - For some queries, an equality propagation problem could cause `a = b` and `b = a` to be handled differently. (Bug #40925)
 - Killing an `INSERT ... SELECT` statement for a `MyISAM` table could cause table corruption if the table had indexes. (Bug #40827)

- A multiple-table `DELETE IGNORE` statement involving a foreign key constraint caused an assertion failure. (Bug #40127)
- Multiple-table `UPDATE` statements did not properly activate triggers. (Bug #39953)
- The `mysql_setpermission` operation for removing database privileges removed global privileges instead. (Bug #39852)
- A stored routine contain a C-style comment could not be dumped and reloaded. (Bug #39559)
- In an `UPDATE` or `DELETE` through a secondary index, `InnoDB` did not store the cursor position. This made `InnoDB` crash in semi-consistent read while attempting to unlock a nonmatching record. (Bug #39320)
- The functions listed in [MySQL-Specific Functions That Create Geometry Values](#), previously accepted WKB arguments and returned WKB values. They now accept WKB or geometry arguments and return geometry values.

The functions listed in [Functions That Create Geometry Values from WKB Values](#), previously accepted WKB arguments and returned geometry values. They now accept WKB or geometry arguments and return geometry values. (Bug #38990)

- On Windows, running the server with `myisam_use_mmap` enabled caused `MyISAM` table corruption. (Bug #38848)
- `CHECK TABLE` did not properly check whether `MyISAM` tables created by servers from MySQL 4.0 or older needed to be upgraded. This could cause problems upgrading to MySQL 5.1 or higher. (Bug #37631)
- An `UPDATE` statement that updated a column using the same `DES_ENCRYPT()` value for each row actually updated different rows with different values. (Bug #35087)
- For shared-memory connections, the read and write methods did not properly handle asynchronous close events, which could lead to the client locking up waiting for a server response. For example, a call to `mysql_real_query()` would block forever on the client side if the executed statement was aborted on the server side. Thanks to Armin Schöffmann for the bug report and patch. (Bug #33899)
- `CHECKSUM TABLE` was not killable with `KILL QUERY`. (Bug #33146)
- `myisamchk` and `myisampack` were not being linked with the library that enabled support for `*` file name pattern expansion. (Bug #29248)
- For `InnoDB` tables that have their own `.ibd` tablespace file, a superfluous `ibuf cursor restoration fails!` message could be written to the error log. This warning has been suppressed. (Bug #27276)
- `COMMIT` did not delete savepoints if there were no changes in the transaction. (Bug #26288)
- Several memory allocation functions were not being checked for out-of-memory return values. (Bug #25058)

Changes in MySQL Enterprise 5.1.34sp1 [QSP] (2009-06-25)

This is a *Service Pack* release of the MySQL Enterprise Server 5.1.

This section documents all changes and bugfixes that have been applied since the last MySQL Enterprise Server release (5.1.34).

**Note**

The fix for Bug #40974 in MySQL 5.1.31 caused the regression problem reported in Bug #44810. Users for whom stability is of utmost priority should note that 5.1.34sp1 is affected by this problem because Bug #44810 is not fixed until MySQL 5.1.36.

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 Enterprise* (a commercial MySQL offering). For more details please see <http://www.mysql.com/products/enterprise/advisors.html>.

Bugs Fixed

- Incomplete cleanup of `JOIN_TAB::select` during the filesort of rows for a `GROUP BY` clause inside a subquery caused a server crash. (Bug #44290)
- Use of `HANDLER` statements with `INFORMATION_SCHEMA` tables caused a server crash. Now `HANDLER` is prohibited with such tables. (Bug #44151)
- On 64-bit systems, a `key_buffer_size` value larger than 4GB could cause `MyISAM` index corruption. (Bug #43932)
- On Windows, a server crash occurred for attempts to insert a floating-point value into a `CHAR` column with a maximum length less than the converted floating-point value length. (Bug #43833)
- `libmysqld` crashed when it was reinitialized. (Bug #43706, Bug #44091)
- Certain statements might open a table and then wait for an impending global read lock without noticing whether they hold a table being waiting for by the global read lock, causing a hang. Affected statements are `SELECT ... FOR UPDATE`, `LOCK TABLES ... WRITE`, `TRUNCATE TABLE`, and `LOAD DATA INFILE`. (Bug #43230)
- Using an XML function such as `ExtractValue()` more than once in a single query could produce erroneous results. (Bug #43183)

References: See also: Bug #43937.

- Incorrect elevation of warning messages to error messages for unsafe statements caused a server crash. (Bug #42640)
- In an `UPDATE` or `DELETE` through a secondary index, `InnoDB` did not store the cursor position. This made `InnoDB` crash in semi-consistent read while attempting to unlock a nonmatching record. (Bug #39320)
- The functions listed in [MySQL-Specific Functions That Create Geometry Values](#), previously accepted WKB arguments and returned WKB values. They now accept WKB or geometry arguments and return geometry values.

The functions listed in [Functions That Create Geometry Values from WKB Values](#), previously accepted WKB arguments and returned geometry values. They now accept WKB or geometry arguments and return geometry values. (Bug #38990)

Changes in MySQL 5.1.34 (2009-04-02)

- [AIX Notes](#)
- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

AIX Notes

- Support Ending for AIX 5.2: Per the <http://www.mysql.com/about/legal/lifecycle/> regarding ending support for OS versions that have reached vendor end of life, we plan to discontinue building or supporting MySQL binaries for AIX 5.2 as of April 30, 2009. This release of MySQL 5.1 (5.1.34) is the last MySQL 5.1 release with support for AIX 5.2. For more information, see the March 24, 2009 note at [MySQL Product Support EOL Announcements](#).

Functionality Added or Changed

- The `optimizer_switch` system variable is now available to control optimizations that can be switched on and off. See [Controlling Switchable Optimizations](#).

Bugs Fixed

- **Important Note; Replication:** Binary logging with `--binlog-format=ROW` failed when a change to be logged included more than 251 columns. This issue was not known to occur with mixed-format or statement-based logging. (Bug #42977)

References: See also: Bug #42914.
- **Replication:** Assigning an invalid directory for the `--slave-load-tmpdir` caused the replication slave to crash. (Bug #42861)
- **Replication:** The `mysql.procs_priv` system table was not replicated. (Bug #42217)
- **Replication:** An `INSERT DELAYED` into a `TIMESTAMP` column issued concurrently with an insert on the same column not using `DELAYED`, but applied after the other insert, was logged using the same timestamp as generated by the other (non-`DELAYED`) insert. (Bug #41719)
- **Replication:** The `MIXED` binary logging format did not switch to row-based mode for statements containing the `LOAD_FILE()` function. (Bug #39701)
- **Replication:** When the server SQL mode included `IGNORE_SPACE`, statement-based replication of `LOAD DATA INFILE ... INTO tbl_name` failed because the statement was read incorrectly from the binary log; a trailing space was omitted, causing the statement to fail with a syntax error when run on the slave. (Bug #22504)

References: See also: Bug #43746.

- An attempt by a user who did not have the `SUPER` privilege to kill a system thread could cause a server crash. (Bug #43748)
- On Windows, incorrectly specified link dependencies in `CMakeLists.txt` resulted in link errors for `mysql_embedded`, `mysqltest_embedded`, and `mysql_client_test_embedded`. (Bug #43715)
- `mysql` crashed if a request for the current database name returned an empty result, such as after the client has executed a preceding `SET sql_select_limit=0` statement. (Bug #43254)
- If the value of the `version_comment` system variable was too long, the `mysql` client displayed a truncated startup message. (Bug #43153)
- Queries of the following form returned an empty result:

```
SELECT ... WHERE ... (col=col AND col=col) OR ... (false expression)
```

(Bug #42957)

- The `strings/CHARSET_INFO.txt` file was not included in source distributions. (Bug #42937)

- A dangling pointer in `mysys/my_error.c` could lead to client crashes. (Bug #42675)
- Passing an unknown time zone specification to `CONVERT_TZ()` resulted in a memory leak. (Bug #42502)
- The MySQL Instance Configuration Wizard failed to start correctly on Windows Vista. (Bug #42386)
- With more than two arguments, `LEAST()`, `GREATEST()`, and `CASE` could unnecessarily return `Illegal mix of collations` errors. (Bug #41627)
- The `mysql` client could misinterpret its input if a line was longer than an internal buffer. (Bug #41486)
- In the `help` command output displayed by `mysql`, the description for the `\c` (`clear`) command was misleading. (Bug #41268)
- The `load_defaults()`, `my_search_option_files()` and `my_print_default_files()` functions in the C client library were subject to a race condition in multi-threaded operation. (Bug #40552)
- If `--basedir` was specified, `mysqld_safe` did not use it when attempting to locate `my_print_defaults`. (Bug #39326)
- When running the MySQL Instance Configuration Wizard in command-line only mode, the service name would be ignored (effectively creating all instances with the default `MySQL` service name), irrespective of the name specified on the command line. However, the wizard would attempt to start the service with the specified name, and failed. (Bug #38379)
- When MySQL was configured with the `--with-max-indexes=128` option, `mysqld` crashed. (Bug #36751)
- Setting the `join_buffer_size` variable to its minimum value produced spurious warnings. (Bug #36446)
- The use of `NAME_CONST()` can result in a problem for `CREATE TABLE ... SELECT` statements when the source column expressions refer to local variables. Converting these references to `NAME_CONST()` expressions can result in column names that are different on the master and slave servers, or names that are too long to be legal column identifiers. A workaround is to supply aliases for columns that refer to local variables.

Now a warning is issued in such cases that indicate possible problems. (Bug #35383)
- An attempt to check or repair an `ARCHIVE` table that had been subjected to a server crash returned a 144 internal error. The data appeared to be irrecoverable. (Bug #32880)
- The `Time` column for `SHOW PROCESSLIST` output and the value of the `TIME` column of the `INFORMATION_SCHEMA.PROCESSLIST` table now can have negative values. Previously, the column was unsigned and negative values were displayed incorrectly as large positive values. Negative values can occur if a thread alters the time into the future with `SET TIMESTAMP = value` or the thread is executing on a slave and processing events from a master that has its clock set ahead of the slave. (Bug #22047)
- Restoring a `mysqldump` dump file containing `FEDERATED` tables failed because the file contained the data for the table. Now only the table definition is dumped (because the data is located elsewhere). (Bug #21360)

Changes in MySQL 5.1.33 (2009-03-13)

- [AIX Notes](#)

- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

AIX Notes

- Support Ending for AIX 5.2: Per the <http://www.mysql.com/about/legal/lifecycle/> regarding ending support for OS versions that have reached vendor end of life, we plan to discontinue building or supporting MySQL binaries for AIX 5.2 as of April 30, 2009. The next release of MySQL 5.1 (5.1.34) will be the last MySQL 5.1 release with support for AIX 5.2. For more information, see the March 24, 2009 note at [MySQL Product Support EOL Announcements](#).

Functionality Added or Changed

- **Performance:** The query cache now checks whether a `SELECT` statement begins with `SQL_NO_CACHE` to determine whether it can skip checking for the query result in the query cache. This is not supported when `SQL_NO_CACHE` occurs within a comment. (Bug #37416)
- `mysql-test-run.pl` now supports an `--experimental=file_name` option. It enables you to specify a file that contains a list of test cases that should be displayed with the `[exp-fail]` code rather than `[fail]` if they fail. (Bug #42888)
- The MD5 algorithm now uses the Xfree implementation. (Bug #42434)

Bugs Fixed

- **Partitioning:** A duplicate key error raised when inserting into a partitioned table using a different error code from that returned by such an error raised when inserting into a table that was not partitioned. (Bug #38719)

References: See also: Bug #28842.

- **Partitioning:** Several error messages relating to partitioned tables were incorrect or missing. (Bug #36001)
- **Replication:** When `binlog_format` was set to `STATEMENT`, a statement unsafe for statement-based logging caused an error or warning to be issued even if `sql_log_bin` was set to 0. (Bug #41980)
- **Replication:** When using `MIXED` replication format and temporary tables were created in statement-based mode, but a later operation in the same session caused a switch to row-based mode, the temporary tables were not dropped on the slave at the end of the session. (Bug #40013)

References: See also: Bug #43046. This issue is a regression of: Bug #20499.

- **Replication:** When using the `MIXED` replication format, `UPDATE` and `DELETE` statements that searched for rows where part of the key had nullable `BIT` columns failed. This occurred because operations that inserted the data were replicated as statements, but `UPDATE` and `DELETE` statements affecting the same data were replicated using row-based format.

This issue did not occur when using statement-based replication (only) or row-based replication (only). (Bug #39753)

References: See also: Bug #39648.

- **Replication:** The server SQL mode in effect when a stored procedure was created was not retained in the binary log. This could cause a `CREATE PROCEDURE` statement that succeeded on the master to fail on the slave.

This issue was first noticed when a stored procedure was created when `ANSI_QUOTES` was in effect on the master, but could possibly cause failed `CREATE PROCEDURE` statements and other problems on the slave when using other server SQL modes as well. (Bug #39526)

- **Replication:** If `--secure-file-priv` was set on the slave, it was unable to execute `LOAD DATA INFILE` statements sent from the master when using mixed-format or statement-based replication.

As a result of this fix, this security restriction is now ignored on the slave in such cases; instead the slave checks whether the files were created and should be read by the slave in its `--slave-load-tmpdir`. (Bug #38174)

- **Replication:** Server IDs greater than $2^{32} - 1$ were represented by negative numbers in the binary log. (Bug #37313)
- **Replication:** When its disk becomes full, a replication slave may wait while writing the binary log, relay log or `MyISAM` tables, continuing after space has been made available. The error message provided in such cases was not clear about the frequency with which checking for free space is done (once every 60 seconds), and how long the server waits after space has been freed before continuing (also 60 seconds); this caused users to think that the server had hung.

These issues have been addressed by making the error message clearer, and dividing it into two separate messages:

1. The error message `Disk is full writing 'filename' (Errcode: error_code). Waiting for someone to free space... (Expect up to 60 secs delay for server to continue after freeing disk space)` is printed only once.
2. The warning `Retry in 60 secs, Message reprinted in 600 secs` is printed once every for every 10 times that the check for free space is made; that is, the check is performed once each 60 seconds, but the reminder that space needs to be freed is printed only once every 10 minutes (600 seconds).

(Bug #22082)

- **Replication:** The statements `DROP PROCEDURE IF EXISTS` and `DROP FUNCTION IF EXISTS` were not written to the binary log if the procedure or function to be dropped did not exist. (Bug #13684)

References: See also: Bug #25705.

- The IBM DB2i storage engine has been added to this release for the IBM i Series platform. For more information, see [The IBMDB2I Storage Engine](#). (Bug #44217)
- On 64-bit debug builds, code in `safemalloc` resulted in errors due to use of a 32-bit value for 64-bit allocations. (Bug #43885)
- `make distcheck` failed to properly handle subdirectories of `storage/ndb`. (Bug #43614)
- Use of `USE INDEX` hints could cause `EXPLAIN EXTENDED` to crash. (Bug #43354)
- For `InnoDB` tables, overflow in an `AUTO_INCREMENT` column could cause a server crash. (Bug #43203)
- On 32-bit Windows, `mysqld` could not use large buffers due to a 2GB user mode address limit. (Bug #43082)
- `stderr` should be unbuffered, but when the server redirected `stderr` to a file, it became buffered. (Bug #42790)

- The `DATA_TYPE` column of the `INFORMATION_SCHEMA.COLUMNS` table displayed the `UNSIGNED` attribute for floating-point data types. (The column should contain only the data type name.) (Bug #42758)
- For `InnoDB` tables, spurious duplicate-key errors could occur when inserting into an `AUTO_INCREMENT` column. (Bug #42714)
- `mysqldump` included views that were excluded with the `--ignore-table` option. (Bug #42635)
- An earlier bug fix resulted in the problem that the `InnoDB` plugin could not be used with a server that was compiled with the built-in `InnoDB`. To handle this two changes were made:
 - The server now supports an `--ignore-builtin-innodb` option that causes the server to behave as if the built-in `InnoDB` is not present. This option causes other `InnoDB` options not to be recognized.
 - For the `INSTALL PLUGIN` statement, the server reads option (`my.cnf`) files just as during server startup. This enables the plugin to pick up any relevant options from those files. Consequently, a plugin no longer is started with each option set to its default value.

Because of this change, it is possible to add plugin options to an option file even before loading a plugin (if the `loose` prefix is used). It is also possible to uninstall a plugin, edit `my.cnf`, and install the plugin again. Restarting the plugin this way enables it to the new option values without a server restart.

**Note**

`InnoDB` Plugin versions 1.0.4 and higher will take advantage of this bug fix. Although the `InnoDB` Plugin is source code compatible with multiple MySQL releases, a given binary `InnoDB` Plugin can be used only with a specific MySQL release. When `InnoDB` Plugin 1.0.4 is released, it is expected to be compiled for MySQL 5.1.34. For 5.1.33, you can use `InnoDB` Plugin 1.0.3, but you must build from source.

(Bug #42610)

References: This issue is a regression of: Bug #29263.

- With the `ONLY_FULL_GROUP_BY` SQL mode enabled, some legal queries failed. (Bug #42567)
- Tables could enter open table cache for a thread without being properly cleaned up, leading to a server crash. (Bug #42419)
- For `InnoDB` tables, inserting into floating-point `AUTO_INCREMENT` columns failed. (Bug #42400)
- The `InnoDB btr_search_drop_page_hash_when_freed()` function had a race condition. (Bug #42279)
- For `InnoDB` tables, there was a race condition for `ALTER TABLE`, `OPTIMIZE TABLE`, `CREATE INDEX`, and `DROP INDEX` operations when periodically checking whether table copying can be committed. (Bug #42152)
- Parsing of the optional microsecond component of `DATETIME` values did not fail gracefully when that component width was larger than the permitted six places. (Bug #42146)
- In `InnoDB` recovery after a server crash, table lookup could fail and corrupt the data dictionary cache. (Bug #42075)
- `mysqldumpslow` parsed the `--debug` and `--verbose` options incorrectly. (Bug #42027)
- Queries that used the loose index scan access method could return no rows. (Bug #41610)

- In `InnoDB` recovery after a server crash, rollback of a transaction that updated a column from `NULL` to `NULL` could cause another crash. (Bug #41571)
- The error message for a too-long column comment was `Unknown error` rather than a more appropriate message. (Bug #41465)
- Use of `SELECT *` permitted users with rights to only some columns of a view to access all columns. (Bug #41354)
- If the tables underlying a `MERGE` table had a primary key but the `MERGE` table itself did not, inserting a duplicate row into the `MERGE` table caused a server crash. (Bug #41305)
- The server did not robustly handle problems hang if a table opened with `HANDLER` needed to be re-opened because it had been altered to use a different storage engine that does not support `HANDLER`. The server also failed to set an error if the re-open attempt failed. These problems could cause the server to crash or hang. (Bug #41110, Bug #41112)
- `SELECT` statements executed concurrently with `INSERT` statements for a `MyISAM` table could cause incorrect results to be returned from the query cache. (Bug #41098)
- For prepared statements, multibyte character sets were not taking into account when calculating `max_length` for string values and `mysql_stmt_fetch()` could return truncated strings. (Bug #41078)
- Deprecation warnings that referred to MySQL 5.2 were changed to refer to MySQL 6.0. (Bug #41077)
- For user-defined variables in a query result, incorrect length values were returned in the result metadata. (Bug #41030)
- On Windows, starting the server with an invalid value for `innodb_flush_method` caused a crash. (Bug #40757)
- MySQL 5.1 crashed with index merge algorithm and merge tables.
A query in the `MyISAM` merge table caused a crash if the index merge algorithm was being used. (Bug #40675)
- With strict SQL mode enabled, setting a system variable to an out-of-bounds value caused an assertion failure. (Bug #40657)
- Table temporary scans were slower than necessary due to use of `mmap` rather than caching, even with the `myisam_use_mmap` system variable disabled. (Bug #40634)
- For a view that references a table in another database, `mysqldump` wrote the view name qualified with the current database name. This makes it impossible to reload the dump file into a different database. (Bug #40345)
- On platforms where long and pointer variables have different sizes, `MyISAM` could copy key statistics incorrectly, resulting in a server crash or incorrect cardinality values. (Bug #40321)
- `DELETE` tried to acquire write (not read) locks for tables accessed within a subquery of the `WHERE` clause. (Bug #39843)
- `perror` did not produce correct output for error codes 153 to 163. (Bug #39370)
- Several functions in `libmysqld` called `exit()` when an error occurred rather than returning an error to the caller. (Bug #39289)
- The `innodb_log_arch_dir` system variable is no longer available but was present in some of the sample option files included with MySQL distributions (such as `my-huge.cnf`). The line was present as

- a comment but uncommenting it would cause server startup failure so the line has been removed. (Bug #38249)
- Setting a savepoint with the same name as an existing savepoint incorrectly deleted any other savepoints that had been set in the meantime. For example, setting savepoints named `a`, `b`, `c`, `b` resulted in savepoints `a`, `b`, rather than the correct savepoints `a`, `c`, `b`. (Bug #38187)
 - `--help` output for `myisamchk` did not list the `--HELP` option. (Bug #38103)
 - Comparisons between row constructors, such as `(a, b) = (c, d)` resulted in unnecessary `Illegal mix of collations` errors for string columns. (Bug #37601)
 - If a user created a view that referenced tables for which the user had disjoint privileges, an assertion failure occurred. (Bug #37191)
 - An argument to the `MATCH()` function that was an alias for an expression other than a column name caused a server crash. (Bug #36737)
 - The `event`, `general_log`, and `slow_log` tables in the `mysql` database store `server_id` values, but did not use an `UNSIGNED` column and thus were not able to store the full range of ID values. (Bug #36540)
 - On Windows, the `_PC` macro in `my_global.h` was causing problems for modern compilers. It has been removed because it is no longer used. (Bug #34309)
 - For `DROP FUNCTION` with names that were qualified with a database name, the database name was handled in case-sensitive fashion even with `lower_case_table_names` set to 1. (Bug #33813)
 - `mysqldump --compatible=mysql40` emitted statements referring to the `character_set_client` system variable, which is unknown before MySQL 4.1. Now the statements are enclosed in version-specific comments. (Bug #33550)
 - Detection by `configure` of several functions such as `setsockopt()`, `bind()`, `sched_yield()`, and `getty()` could fail. (Bug #31506)
 - Use of MBR spatial functions such as `MBRTouches()` with columns of `InnoDB` tables caused a server crash rather than an error. (Bug #31435)
 - The `mysql` client mishandled input parsing if a `delimiter` command was not first on the line. (Bug #31060)
 - `SHOW PRIVILEGES` listed the `CREATE ROUTINE` privilege as having a context of `Functions,Procedures`, but it is a database-level privilege. (Bug #30305)
 - `mysqld --help` did not work as `root`. (Bug #30261)
 - `CHECK TABLE`, `REPAIR TABLE`, `ANALYZE TABLE`, and `OPTIMIZE TABLE` erroneously reported a table to be corrupt if the table did not exist or the statement was terminated with `KILL`. (Bug #29458)
 - `SHOW TABLE STATUS` could fail to produce output for tables with non-ASCII characters in their name. (Bug #25830)
 - Allocation of stack space for error messages could be too small on HP-UX, leading to stack overflow crashes. (Bug #21476)
 - Floating-point numbers could be handled with different numbers of digits depending on whether the text or prepared-statement protocol was used. (Bug #21205)
 - Incorrect length metadata could be returned for `LONG TEXT` columns when a multibyte server character set was used. (Bug #19829)

- [ROUND\(\)](#) sometimes returned different results on different platforms. (Bug #15936)

Changes in MySQL 5.1.32 (2009-02-14)

- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

Functionality Added or Changed

- The [libedit](#) library was upgraded to version 2.11. (Bug #42433)

Bugs Fixed

- **Security Fix:** Using an XPath expression employing a scalar expression as a [FilterExpr](#) with [ExtractValue\(\)](#) or [UpdateXML\(\)](#) caused the server to crash. Such expressions now cause an error instead. (Bug #42495)
- **Incompatible Change:** The fix for Bug #33699 introduced a change to the [UPDATE](#) statement such that assigning [NULL](#) to a [NOT NULL](#) column caused an error even when strict SQL mode was not enabled. The original behavior before was that such assignments caused an error only in strict SQL mode, and otherwise set the column to the implicit default value for the column data type and generated a warning. (For information about implicit default values, see [Data Type Default Values](#).)

The change caused compatibility problems for applications that relied on the original behavior. It also caused replication problems between servers that had the original behavior and those that did not, for applications that assigned [NULL](#) to [NOT NULL](#) columns in [UPDATE](#) statements without strict SQL mode enabled. This change has been reverted so that [UPDATE](#) again had the original behavior. Problems can still occur if you replicate between servers that have the modified [UPDATE](#) behavior and those that do not. (Bug #39265)

References: See also: Bug #33699.

- **Important Change:** When using the MySQL Instance Configuration Wizard with a configuration where you already have an existing installation with a custom [datadir](#), the wizard could reset the data to the default data directory. When performing an upgrade installation in this situation, you must re-specify your custom settings, including the [datadir](#), to ensure that your configuration file is not reset to the default values. (Bug #37534)
- **Important Change:** Uninstalling MySQL using the MySQL installer on Windows would delete the [my.ini](#) file. The file is no longer deleted. In addition, when a new installation is conducted, any existing configuration file will be renamed to [myDATETIME.ini.bak](#) during configuration. (Bug #36493)
- **Important Change:** When installing MySQL on Windows, it was possible to install multiple editions (Complete, and Essential, for example) of the same version of MySQL, leading to two separate entries in the installed packages which were impossible to isolate. This could lead to problems with installation and uninstallation. The MySQL installer on Windows no longer permits multiple installations of the same version of MySQL on a single machine. (Bug #4217)
- **Replication:** [START SLAVE UNTIL](#) did not work correctly with [--replicate-same-server-id](#) enabled; when started with this option, the slave did not perform events recorded in the relay log and that originated from a different master.

Log rotation events are automatically generated and written when rotating the binary log or relay log. Such events for relay logs are usually ignored by the slave SQL thread because they have the same server ID as that of the slave. However, when [--replicate-same-server-id](#) was enabled, the rotation event for the relay log was treated as if it originated on the master, because the log's name and

position were incorrectly updated. This caused the `MASTER_POS_WAIT()` function always to return `NULL` and thus to fail. (Bug #38734, Bug #38934)

- **Replication:** `TRUNCATE TABLE` statements failed to replicate when statement-based binary logging mode was not available. The issue was observed when using `InnoDB` with the transaction isolation level set to `READ UNCOMMITTED` (thus forcing `InnoDB` not to permit statement-based logging). However, the same behavior could be reproduced using any transactional storage engine supporting only row-based logging, regardless of the isolation level. This was due to two separate problems:
 1. An error was printed by `InnoDB` for `TRUNCATE TABLE` when using statement-based logging mode where the transaction isolation level was set to `READ COMMITTED` or `READ UNCOMMITTED`, because `InnoDB` permits statement-based replication for DML statements. However, `TRUNCATE TABLE` is not transactional; since it is the equivalent of `DROP TABLE` followed by `CREATE TABLE`, it is actually DDL, and should therefore be permitted to be replicated as a statement.
 2. `TRUNCATE TABLE` was not logged in mixed mode because of the error just described; however, this error was not reported to the client.

As a result of this fix, `TRUNCATE TABLE` is now treated as DDL for purposes of binary logging and replication; that is, it is always logged as a statement and so no longer causes an error when replicated using a transactional storage engine such as `InnoDB`. (Bug #36763)

References: See also: Bug #42643.

- **Replication:** `mysqlbinlog` replay of `CREATE TEMPORARY TABLE ... LIKE` statements and of `TRUNCATE TABLE` statements used on temporary tables failed with Error 1146 (`Table ... doesn't exist`). (Bug #35583)
- **Replication:** In statement mode, `mysqlbinlog` failed to issue a `SET @@autocommit` statement when the autocommit mode was changed. (Bug #34541)
- **Replication:** `LOAD DATA INFILE` statements did not replicate correctly from a master running MySQL 4.1 to a slave running MySQL 5.1 or later. (Bug #31240)
- The use by `libedit` of the `__weak_reference()` macro caused compilation failure on FreeBSD. (Bug #42817)
- A `'%'` character in SQL statements could cause the server to crash. (Bug #42634)
- An optimization introduced for Bug #37553 required an explicit cast to be added for some uses of `TIMEDIFF()` because automatic casting could produce incorrect results. (It was necessary to use `TIME(TIMEDIFF(...))`.) (Bug #42525)

References: See also: Bug #37553.

- On the IBM i5 platform, the MySQL configuration process caused the system version of `pthread_setschedprio()` to be used. This function returns `SIGILL` on i5 because it is not supported, causing the server to crash. Now the `my_pthread_setprio()` function in the `mysys` library is used instead. (Bug #42524)
- The SSL certificates included with MySQL distributions were regenerated because the previous ones had expired. (Bug #42366)
- User variables within triggers could cause a crash if the `mysql_change_user()` C API function was invoked. (Bug #42188)
- Dependent subqueries such as the following caused a memory leak proportional to the number of outer rows:


```
SELECT COUNT(*) FROM t1, t2 WHERE t2.b
  IN (SELECT DISTINCT t2.b FROM t2 WHERE t2.b = t1.a);
```

(Bug #42037)

- Some queries using `NAME_CONST(.. COLLATE ...)` led to a server crash due to a failed type cast. (Bug #42014)
- On Mac OS X, some of the universal client libraries were not actually universal and were missing code for one or more architectures. (Bug #41940)
- String reallocation could cause memory overruns. (Bug #41868)
- `mysql_install_db` did not pass some relevant options to `mysqld`. (Bug #41828)
- Setting `innodb_locks_unsafe_for_binlog` should be equivalent to setting the transaction isolation level to `READ COMMITTED`. However, if both of those things were done, nonmatching semi-consistently read rows were not unlocked when they should have been. (Bug #41671)
- `REPAIR TABLE` crashed for compressed `MyISAM` tables. (Bug #41574)
- For a `TIMESTAMP NOT NULL DEFAULT ...` column, storing `NULL` as the return value from some functions caused a “cannot be NULL” error. `NULL` returns now correctly cause the column default value to be stored. (Bug #41370)
- The server cannot execute `INSERT DELAYED` statements when statement-based binary logging is enabled, but the error message displayed only the table name, not the entire statement. (Bug #41121)
- `FULLTEXT` indexes did not work for Unicode columns that used a custom UCA collation. (Bug #41084)
- The Windows installer displayed incorrect product names in some images. (Bug #40845)
- Changing `innodb_thread_concurrency` at runtime could cause errors. (Bug #40760)
- `SELECT` statements could be blocked by `INSERT DELAYED` statements that were waiting for a lock, even with `low_priority_updates` enabled. (Bug #40536)
- For `InnoDB` tables that used `ROW_FORMAT=REDUNDANT`, storage size of `NULL` columns could be determined incorrectly. (Bug #40369)
- The query cache stored only partial query results if a statement failed while the results were being sent to the client. This could cause other clients to hang when trying to read the cached result. Now if a statement fails, the result is not cached. (Bug #40264)
- When a `MEMORY` table became full, the error generated was returned to the client but was not written to the error log. (Bug #39886)
- With row-based binary logging, replication of `InnoDB` tables containing `NULL`-valued `BIT` columns could fail. (Bug #39648)
- The expression `ROW(...) IN (SELECT ... FROM DUAL)` always returned `TRUE`. (Bug #39069)
- The greedy optimizer could cause a server crash due to improper handling of nested outer joins. (Bug #38795)
- Use of `COUNT(DISTINCT)` prevented `NULL` testing in the `HAVING` clause. (Bug #38637)
- The `innodb_stats_on_metadata` system variable was not displayed by `SHOW VARIABLES` and was not settable at runtime. (Bug #38189)

- Enabling the `sync_frm` system variable had no effect on the handling of `.frm` files for views. (Bug #38145)
- The embedded server truncated some error messages. (Bug #37995)
- For comparison of `NULL` to a subquery result inside `IS NULL`, the comparison could evaluate to `NULL` rather than to `TRUE` or `FALSE`. This occurred for expressions such as:

```
SELECT ... WHERE NULL IN (SELECT ...) IS NULL
```

(Bug #37822)

- Setting `myisam_repair_threads` greater than 1 caused a server crash for table repair or alteration operations for `MyISAM` tables with multiple `FULLTEXT` indexes. (Bug #37756)
- When using the MySQL MSI Installer on Windows and selecting **Back** after a choosing Repair, you would be returned to the Fresh Install section of the installer. You are now correctly returned to the Install, Repair, Modify screen. (Bug #37294)
- The `mysql` client sometimes improperly interpreted string escape sequences in nonstring contexts. (Bug #36391)
- The query cache stored packets containing the server status of the time when the cached statement was run. This might lead to an incorrect transaction status on the client side if a statement was cached during a transaction and later served outside a transaction context (or vice versa). (Bug #36326)
- If the system time was adjusted backward during query execution, the apparent execution time could be negative. But in some cases these queries would be written to the slow query log, with the negative execution time written as a large unsigned number. Now statements with apparent negative execution time are not written to the slow query log. (Bug #35396)
- `libmysqld` was not built with all character sets. (Bug #32831)
- For `mysqld_multi`, using the `--mysqld=mysqld_safe` option caused the `--defaults-file` and `--defaults-extra-file` options to behave the same way. (Bug #32136)
- Attempts to open a valid `MERGE` table sometimes resulted in a `ER_WRONG_MRG_TABLE` error. This happened after failure to open an invalid `MERGE` table had also generated an `ER_WRONG_MRG_TABLE` error. (Bug #32047)
- For Solaris package installation using `pkgadd`, the postinstall script failed, causing the system tables in the `mysql` database not to be created. (Bug #31164)
- If the default database was dropped, the value of `character_set_database` was not reset to `character_set_server` as it should have been. (Bug #27208)

References: See also: Bug #63524.

Changes in MySQL Enterprise 5.1.31sp1 [QSP] (2009-03-19)

This is a *Service Pack* release of the MySQL Enterprise Server 5.1.

This section documents all changes and bugfixes that have been applied since the last MySQL Enterprise Server release (5.1.31).

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 Enterprise* (a commercial MySQL offering). For more details please see <http://www.mysql.com/products/enterprise/advisors.html>.

- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

Functionality Added or Changed

- The `libedit` library was upgraded to version 2.11. (Bug #42433)

Bugs Fixed

- **Security Fix:** Using an XPath expression employing a scalar expression as a [FilterExpr](#) with [ExtractValue\(\)](#) or [UpdateXML\(\)](#) caused the server to crash. Such expressions now cause an error instead. (Bug #42495)
- On the IBM i5 platform, the MySQL configuration process caused the system version of [pthread_setschedprio\(\)](#) to be used. This function returns `SIGILL` on i5 because it is not supported, causing the server to crash. Now the [my_pthread_setprio\(\)](#) function in the `mysys` library is used instead. (Bug #42524)
- The SSL certificates included with MySQL distributions were regenerated because the previous ones had expired. (Bug #42366)
- User variables within triggers could cause a crash if the [mysql_change_user\(\)](#) C API function was invoked. (Bug #42188)
- Some queries using [NAME_CONST\(... COLLATE ...\)](#) led to a server crash due to a failed type cast. (Bug #42014)

Changes in MySQL 5.1.31 (2009-01-19)

- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

Functionality Added or Changed

- [MySQL-shared-compat-advanced-gpl-5.1.31-0.*.rpm](#) and [MySQL-shared-compat-advanced-5.1.31-0.*.rpm](#) packages are now available. These client library compatibility packages are like the [MySQL-shared-compat](#) package, but are for the “MySQL Enterprise Server — Advanced Edition” products. Install these packages rather than the normal [MySQL-shared-compat](#) package if you want to include shared client libraries for older MySQL versions. (Bug #41838)
- A new status variable, [Queries](#), indicates the number of statements executed by the server. This includes statements executed within stored programs, unlike the [Questions](#) variable which includes only statements sent to the server by clients. (Bug #41131)
- Performance of [SELECT *](#) retrievals from [INFORMATION_SCHEMA.COLUMNS](#) was improved slightly. (Bug #38918)
- Previously, index hints did not work for [FULLTEXT](#) searches. Now they work as follows:

For natural language mode searches, index hints are silently ignored. For example, [IGNORE INDEX\(i\)](#) is ignored with no warning and the index is still used.

For boolean mode searches, index hints with [FOR ORDER BY](#) or [FOR GROUP BY](#) are silently ignored. Index hints with [FOR JOIN](#) or no [FOR](#) modifier are honored. In contrast to how hints apply for

non-[FULLTEXT](#) searches, the hint is used for all phases of query execution (finding rows and retrieval, grouping, and ordering). This is true even if the hint is given for a non-[FULLTEXT](#) index. (Bug #38842)

Bugs Fixed

- **Performance:** For an [InnoDB](#) table, [DROP TABLE](#) or [ALTER TABLE ... DISCARD TABLESPACE](#) could take a long time or cause a server crash. (Bug #39939)
- **Important Change; Replication:** If a trigger was defined on an [InnoDB](#) table and this trigger updated a nontransactional table, changes performed on the [InnoDB](#) table were replicated and were visible on the slave before they were committed on the master, and were not rolled back on the slave after a successful rollback of those changes on the master.

As a result of the fix for this issue, the semantics of mixing nontransactional and transactional tables in a transaction have changed. Previously, if the initial statements in a transaction contained nontransactional changes, those statements were written directly to the binary log. Now, any statement appearing after a [BEGIN](#) (or immediately following a [COMMIT](#) if [autocommit](#) = 0) is always considered part of the transaction and cached. This means that nontransactional changes do not propagate to the slave until the transaction is committed and thus written to the binary log.

See [Replication and Transactions](#), for more information about this change in behavior. (Bug #40116)

- **Important Change:** The MSI installer packages for Windows are now digitally signed with a certificate, enabling installation on Windows where only certified packages are permitted by group policy or configuration.

As part of this change, and to comply with the certified installer requirements, the [Setup.exe](#) versions of the MySQL installer have been discontinued. You must have Windows Installer support in your Windows installation to use the MSI install package. This is a standard component on Windows XP SP2 and higher. For earlier versions, you can download the Microsoft Installer support from Microsoft.com. (Bug #36409)

- **Partitioning; Replication:** Changing the transaction isolation level while replicating partitioned [InnoDB](#) tables could cause statement-based logging to fail. (Bug #39084)
- **Partitioning:** A comparison with an invalid [DATE](#) value in a query against a partitioned table could lead to a crash of the MySQL server.



Note

Invalid [DATE](#) and [DATETIME](#) values referenced in the [WHERE](#) clause of a query on a partitioned table are treated as [NULL](#). See [Partition Pruning](#), for more information.

(Bug #40972)

- **Partitioning:** A query on a user-partitioned table caused MySQL to crash, where the query had the following characteristics:
 - The query's [WHERE](#) clause referenced an indexed column that was also in the partitioning key.
 - The query's [WHERE](#) clause included a value found in the partition.
 - The query's [WHERE](#) clause used the [<](#) or [<>](#) operators to compare with the indexed column's value with a constant.
 - The query used an [ORDER BY](#) clause, and the same indexed column was used in the [ORDER BY](#) clause.

- The `ORDER BY` clause used an explicit or implicit `ASC` sort priority.

Two examples of such a query are given here, where `a` represents an indexed column used in the table's partitioning key:

1.

```
SELECT * FROM table WHERE a < constant ORDER BY a;
```
2.

```
SELECT * FROM table WHERE a <> constant ORDER BY a;
```

This bug was introduced in MySQL 5.1.29. (Bug #40954)

References: This issue is a regression of: Bug #30573, Bug #33257, Bug #33555.

- **Partitioning:** With `READ COMMITTED` transaction isolation level, `InnoDB` uses a semi-consistent read that releases nonmatching rows after MySQL has evaluated the `WHERE` clause. However, this was not happening if the table used partitions. (Bug #40595)
- **Partitioning:** A query that timed out when run against a partitioned table failed silently, without providing any warnings or errors, rather than returning `Lock wait timeout exceeded`. (Bug #40515)
- **Partitioning:** `ALTER TABLE ... REORGANIZE PARTITION` could crash the server when the number of partitions was not changed. (Bug #40389)

References: See also: Bug #41945.

- **Partitioning:** For a partitioned table having an `AUTO_INCREMENT` column: If the first statement following a start of the server or a `FLUSH TABLES` statement was an `UPDATE` statement, the `AUTO_INCREMENT` column was not incremented correctly. (Bug #40176)
- **Partitioning:** The server attempted to execute the statements `ALTER TABLE ... ANALYZE PARTITION`, `ALTER TABLE ... CHECK PARTITION`, `ALTER TABLE ... OPTIMIZE PARTITION`, and `ALTER TABLE ... REORGANIZE PARTITION` on tables that were not partitioned. (Bug #39434)

References: See also: Bug #20129.

- **Partitioning:** The value of the `CREATE_COLUMNS` column in `INFORMATION_SCHEMA.TABLES` was not `partitioned` for partitioned tables. (Bug #38909)
- **Partitioning:** When executing an `ORDER BY` query on a partitioned `InnoDB` table using an index that was not in the partition expression, the results were sorted on a per-partition basis rather than for the table as a whole. (Bug #37721)
- **Partitioning:** Dropping or creating an index on a partitioned table managed by the `InnoDB` Plugin locked the table. (Bug #37453)
- **Partitioning:** Partitioned table checking sometimes returned a warning with an error code of 0, making proper response to errors impossible. The fix also renders the error message subject to translation in non-English deployments. (Bug #36768)
- **Partitioning:** `SHOW TABLE STATUS` could show a nonzero value for the `Mean record length` of a partitioned `InnoDB` table, even if the table contained no rows. (Bug #36312)
- **Partitioning:** When `SHOW CREATE TABLE` was used on a partitioned table, all of the table's `PARTITION` and `SUBPARTITION` clauses were output on a single line, making it difficult to read or parse. (Bug #14326)

- **Replication:** Per-table `AUTO_INCREMENT` option values were not replicated correctly for `InnoDB` tables. (Bug #41986)
- **Replication:** Some `log_event` types did not skip the post-header when reading. (Bug #41961)
- **Replication:** Attempting to read a binary log containing an `Incident_log_event` having an invalid incident number could cause the debug server to crash. (Bug #40482)
- **Replication:** When using row-based replication, an update of a primary key that was rolled back on the master due to a duplicate key error was not rolled back on the slave. (Bug #40221)
- **Replication:** When rotating relay log files, the slave deletes relay log files and then edits the relay log index file. Formerly, if the slave shut down unexpectedly between these two events, the relay log index file could then reference relay logs that no longer existed. Depending on the circumstances, this could when restarting the slave cause either a race condition or the failure of replication. (Bug #38826, Bug #39325)
- **Replication:** With row-based replication, `UPDATE` and `DELETE` statements using `LIMIT` and a table's primary key could produce different results on the master and slave. (Bug #38230)
- `resolve_stack_dump` was unable to resolve the stack trace format produced by `mysqld` in MySQL 5.1 and up (see [Using a Stack Trace](#)). (Bug #41612)
- In example option files provided in MySQL distributions, the `thread_stack` value was increased from 64K to 128K. (Bug #41577)
- The optimizer could ignore an error and rollback request during a filesort, causing an assertion failure. (Bug #41543)
- `DATE_FORMAT()` could cause a server crash for year-zero dates. (Bug #41470)
- `SET PASSWORD` caused a server crash if the account name was given as `CURRENT_USER()`. (Bug #41456)
- When a repair operation was carried out on a `CSV` table, the debug server crashed. (Bug #41441)
- When substituting system constant functions with a constant result, the server was not expecting `NULL` function return values and could crash. (Bug #41437)
- Queries such as `SELECT ... CASE AVG(...) WHEN ...` that used aggregate functions in a `CASE` expression crashed the server. (Bug #41363)
- `INSERT INTO ... SELECT ... FROM` and `CREATE TABLE ... SELECT ... FROM` a `TEMPORARY` table could inadvertently change the locking type of the temporary table from a write lock to a read lock, causing statement failure. (Bug #41348)
- The `INFORMATION_SCHEMA.SCHEMA_PRIVILEGES` table was limited to 7680 rows. (Bug #41079)
- In debug builds, obsolete debug code could be used to crash the server. (Bug #41041)
- Some queries that used a “range checked for each record” scan could return incorrect results. (Bug #40974)

References: See also: Bug #44810.

- Certain `SELECT` queries could fail with a `Duplicate entry` error. (Bug #40953)
- For debug servers, `OPTIMIZE TABLE` on a compressed table caused a server crash. (Bug #40949)
- Accessing user variables within triggers could cause a server crash. (Bug #40770)

- `IF(..., CAST(longtext_val AS UNSIGNED), signed_val)` as an argument to an aggregate function could cause an assertion failure. (Bug #40761)
- For single-table `UPDATE` statements, an assertion failure resulted from a runtime error in a stored function (such as a recursive function call or an attempt to update the same table as in the `UPDATE` statement). (Bug #40745)
- When executing concurrent `CREATE TABLE ... SELECT` statements on a Maria table, the error `Error: Memory allocated at trnman.c:129 was underrun, discovered at ma_close.c:65` error would be logged in the error file, and the server would eventually crash. (Bug #40416)
- `TRUNCATE TABLE` for an InnoDB table did not flush cached queries for the table. (Bug #40386)
- Prepared statements permitted invalid dates to be inserted when the `ALLOW_INVALID_DATES` SQL mode was not enabled. (Bug #40365)
- `mc.exe` is no longer needed to compile MySQL on Windows. This makes it possible to build MySQL from source using Visual Studio Express 2008. (Bug #40280)
- The `'` character was incorrectly not permitted in table names. (Bug #40104)
- Support for the `revision` field in `.frm` files has been removed. This addresses the downgrading problem introduced by the fix for Bug #17823. (Bug #40021)

References: See also: Bug #17823.

- Retrieval speed from the following `INFORMATION_SCHEMA` tables was improved by shortening the `VARIABLE_VALUE` column to 1024 characters: `GLOBAL_VARIABLES`, `SESSION_VARIABLES`, `GLOBAL_STATUS`, and `SESSION_STATUS`.

As a result of this change, any variable value longer than 1024 characters will be truncated with a warning. This affects only the `init_connect` system variable. (Bug #39955)

- If the operating system is configured to return leap seconds from OS time calls or if the MySQL server uses a time zone definition that has leap seconds, functions such as `NOW()` could return a value having a time part that ends with `:59:60` or `:59:61`. If such values are inserted into a table, they would be dumped as is by `mysqldump` but considered invalid when reloaded, leading to backup/restore problems.

Now leap second values are returned with a time part that ends with `:59:59`. This means that a function such as `NOW()` can return the same value for two or three consecutive seconds during the leap second. It remains true that literal temporal values having a time part that ends with `:59:60` or `:59:61` are considered invalid.

For additional details about leap-second handling, see [Time Zone Leap Second Support](#). (Bug #39920)

- The server could crash during a sort-order optimization of a dependent subquery. (Bug #39844)
- For a server started with the `--temp-pool` option on Windows, temporary file creation could fail. This option now is ignored except on Linux systems, which was its original intended scope. (Bug #39750)
- `ALTER TABLE` on a table with `FULLTEXT` index that used a pluggable `FULLTEXT` parser could cause debug servers to crash. (Bug #39746)
- Performing an `INSERT` on a Maria table with a `UNIQUE` column, MySQL could deadlock. (Bug #39697)
- With the `ONLY_FULL_GROUP_BY` SQL mode enabled, the check for nonaggregated columns in queries with aggregate functions, but without a `GROUP BY` clause was treating all the parts of the query as if they

- were in the select list. This is fixed by ignoring the nonaggregated columns in the `WHERE` clause. (Bug #39656)
- The server crashed if an integer field in a CSV file did not have delimiting quotation marks. (Bug #39616)
 - Creating a table with a comment of 62 characters or longer caused a server crash. (Bug #39591)
 - The `do_abi_check` program run during the build process depends on `mysql_version.h` but that file was not created first, resulting in build failure. (Bug #39571)
 - `CHECK TABLE` failed for `MyISAM INFORMATION_SCHEMA` tables. (Bug #39541)
 - On 64-bit Windows systems, the server accepted `key_buffer_size` values larger than 4GB, but allocated less. (For example, specifying a value of 5GB resulted in 1GB being allocated.) (Bug #39494)
 - `InnoDB` could hang trying to open an adaptive hash index. (Bug #39483)
 - Following `ALTER TABLE ... DISCARD TABLESPACE` for an `InnoDB` table, an attempt to determine the free space for the table before the `ALTER TABLE` operation had completely finished could cause a server crash. (Bug #39438)
 - Use of the `PACK_KEYS` or `MAX_ROWS` table option in `ALTER TABLE` should have triggered table reconstruction but did not. (Bug #39372)
 - Multiple concurrent inserts to a `Maria` table could lead to a deadlock situation. (Bug #39363)
 - The server returned a column type of `VARBINARY` rather than `DATE` as the result from the `COALESCE()`, `IFNULL()`, `IF()`, `GREATEST()`, or `LEAST()` functions or `CASE` expression if the result was obtained using `filesort` in an anonymous temporary table during the query execution. (Bug #39283)
 - `Maria` could fail to find data in a table with an index on a `char` column. (Bug #39243)
 - Running `ALTER TABLE PARTITION` on a `Maria` table would lead to a crash. (Bug #39227)
 - Using `Maria`, executing `FLUSH TABLES WITH READ LOCK` after a `LOCK TABLES` statement would lead to a crash. (Bug #39226)
 - Running multiple `SELECT`, `INSERT`, `UPDATE` and `DELETE` queries on the `Maria` table could lead to a deadlock. (Bug #39210)
 - A server built using `yaSSL` for SSL support would crash if configured to use an RSA key and a client sent a cipher list containing a non-RSA key as acceptable. (Bug #39178)
 - When built with `Valgrind`, the server failed to access tables created with the `DATA DIRECTORY` or `INDEX DIRECTORY` table option. (Bug #39102)
 - With binary logging enabled `CREATE VIEW` was subject to possible buffer overwrite and a server crash. (Bug #39040)
 - The fast mutex implementation was subject to excessive lock contention. (Bug #38941)
 - Use of `InnoDB` monitoring (`SHOW ENGINE INNODB STATUS` or one of the `InnoDB` Monitor tables) could cause a server crash due to invalid access to a shared variable in a concurrent environment. (Bug #38883)
 - `InnoDB` could fail to generate `AUTO_INCREMENT` values after an `UPDATE` statement for the table. (Bug #38839)
 - If delayed insert failed to upgrade the lock, it did not free the temporary memory storage used to keep newly constructed `BLOB` values in memory, resulting in a memory leak. (Bug #38693)

- On Windows, a five-second delay occurred at shutdown of applications that used the embedded server. (Bug #38522)
- On Solaris, a scheduling policy applied to the main server process could be unintentionally overwritten in client-servicing threads. (Bug #38477)
- Building MySQL on FreeBSD would result in a failure during the `gen_lex_hash` phase of the build. (Bug #38364)
- On Windows, the embedded server would crash in `mysql_library_init()` if the language file was missing. (Bug #38293)
- A mix of `TRUNCATE TABLE` with `LOCK TABLES` and `UNLOCK TABLES` for an InnoDB could cause a server crash. (Bug #38231)
- The `ExtractValue()` function did not work correctly with XML documents containing a `DOCTYPE` declaration. (Bug #38227)
- Queries with a `HAVING` clause could return a spurious row. (Bug #38072)
- The Event Scheduler no longer logs “started in thread” or “executed” successfully messages to the error log. (Bug #38066)
- Use of spatial data types in prepared statements could cause memory leaks or server crashes. (Bug #37956, Bug #37671)
- An error in a debugging check caused crashes in debug servers. (Bug #37936)
- A `SELECT` with a `NULL NOT IN` condition containing a complex subquery from the same table as in the outer select caused an assertion failure. (Bug #37894)
- The presence of a `/* ... */` comment preceding a query could cause InnoDB to use unnecessary gap locks. (Bug #37885)
- Use of an uninitialized constant in `EXPLAIN` evaluation caused an assertion failure. (Bug #37870)
- When using `ALTER TABLE` on an InnoDB table, the `AUTO_INCREMENT` value could be changed to an incorrect value. (Bug #37788)
- Primary keys were treated as part of a covering index even if only a prefix of a key column was used. (Bug #37742)
- Renaming an `ARCHIVE` table to the same name with different lettercase and then selecting from it could cause a server crash. (Bug #37719)
- The `MONTHNAME()` and `DAYNAME()` functions returned a binary string, so that using `LOWER()` or `UPPER()` had no effect. Now `MONTHNAME()` and `DAYNAME()` return a value in `character_set_connection` character set. (Bug #37575)
- `TIMEDIFF()` was erroneously treated as always returning a positive result. Also, `CAST()` of `TIME` values to `DECIMAL` dropped the sign of negative values. (Bug #37553)

References: See also: Bug #42525.

- `SHOW PROCESSLIST` displayed “copy to tmp table” when no such copy was occurring. (Bug #37550)
- `mysqlcheck` used `SHOW FULL TABLES` to get the list of tables in a database. For some problems, such as an empty `.frm` file for a table, this failed and `mysqlcheck` then would neglect to check other tables in the database. (Bug #37527)

- Updating a view with a subquery in the `CHECK` option could cause an assertion failure. (Bug #37460)
- Statements that displayed the value of system variables (for example, `SHOW VARIABLES`) expect variable values to be encoded in `character_set_system`. However, variables set from the command line such as `basedir` or `datadir` were encoded using `character_set_filesystem` and not converted correctly. (Bug #37339)
- `CREATE INDEX` could crash with `InnoDB` plugin 1.0.1. (Bug #37284)
- Certain boolean-mode `FULLTEXT` searches that used the truncation operator did not return matching records and calculated relevance incorrectly. (Bug #37245)
- On a 32-bit server built without big tables support, the offset argument in a `LIMIT` clause might be truncated due to a 64-bit to 32-bit cast. (Bug #37075)
- For an `InnoDB` table with a `FOREIGN KEY` constraint, `TRUNCATE TABLE` may be performed using row by row deletion. If an error occurred during this deletion, the table would be only partially emptied. Now if an error occurs, the truncation operation is rolled back and the table is left unchanged. (Bug #37016)
- The code for the `ut_usecetime()` function in `InnoDB` did not handle errors from the `gettimeofday()` system call. Now it retries `gettimeofday()` several times and updates the value of the `Innodb_row_lock_time_max` status variable only if `ut_usecetime()` was successful. (Bug #36819)
- Use of `CONVERT()` with `GROUP BY` to convert numeric values to `CHAR` could return truncated results. (Bug #36772)
- The `mysql` client, when built with Visual Studio 2005, did not display Japanese characters. (Bug #36279)
- `CREATE INDEX` for `InnoDB` tables could under very rare circumstances cause the server to crash.. (Bug #36169)
- A read past the end of the string could occur while parsing the value of the `--innodb-data-file-path` option. (Bug #36149)
- Setting the `slave_compressed_protocol` system variable to `DEFAULT` failed in the embedded server. (Bug #35999)
- For upgrades to MySQL 5.1 or higher, `mysql_upgrade` did not re-encode database or table names that contained nonalphanumeric characters. (They would still appear after the upgrade with the `#mysql50#` prefix described in [Mapping of Identifiers to File Names](#).) To correct this problem, it was necessary to run `mysqlcheck --all-databases --check-upgrade --fix-db-names --fix-table-names` manually. `mysql_upgrade` now runs that command automatically after performing the initial upgrade. (Bug #35934)
- `SHOW CREATE TABLE` did not display a printable value for the default value of `BIT` columns. (Bug #35796)
- The columns that store character set and collation names in several `INFORMATION_SCHEMA` tables were lengthened because they were not long enough to store some possible values: `SCHEMATA`, `TABLES`, `COLUMNS`, `CHARACTER_SETS`, `COLLATIONS`, and `COLLATION_CHARACTER_SET_APPLICABILITY`. (Bug #35789)
- The `max_length` metadata value was calculated incorrectly for the `FORMAT()` function, which could cause incorrect result set metadata to be sent to clients. (Bug #35558)
- `InnoDB` was not updating the `Handler_delete` or `Handler_update` status variables. (Bug #35537)

- `InnoDB` could fail to generate `AUTO_INCREMENT` values if rows previously had been inserted containing literal values for the `AUTO_INCREMENT` column. (Bug #35498, Bug #36411, Bug #39830)
- The `CREATE_OPTIONS` column for `INFORMATION_SCHEMA.TABLES` did not display the `KEY_BLOCK_SIZE` option. (Bug #35275)
- Selecting from an `INFORMATION_SCHEMA` table into an incorrectly defined `MERGE` table caused an assertion failure. (Bug #35068)
- `perror` on Windows did not know about Win32 system error codes. (Bug #34825)
- `EXPLAIN EXTENDED` evaluation of aggregate functions that required a temporary table caused a server crash. (Bug #34773)
- `SHOW GLOBAL STATUS` shows values that aggregate the session status values for all threads. This did not work correctly for the embedded server. (Bug #34517)
- `mysqldumpslow` did not aggregate times. (Bug #34129)
- `mysql_config` did not output `-ldl` (or equivalent) when needed for `--libmysqld-libs`, so its output could be insufficient to build applications that use the embedded server. (Bug #34025)
- The `mysql` client incorrectly parsed statements containing the word “delimiter” in mid-statement. This fix is different from the one applied for this bug in MySQL 5.1.26. (Bug #33812)

References: See also: Bug #38158.

- For a stored procedure containing a `SELECT * ... RIGHT JOIN` query, execution failed for the second call. (Bug #33811)
- Previously, use of index hints with views (which do not have indexes) produced the error `ERROR 1221 (HY000): Incorrect usage of USE/IGNORE INDEX and VIEW`. Now this produces `ERROR 1176 (HY000): Key '...' doesn't exist in table '...'`, the same error as for base tables without an appropriate index. (Bug #33461)
- Three conditions were discovered that could cause an upgrade from MySQL 5.0 to 5.1 to fail: 1) Triggers associated with a table that had a `#mysql50#` prefix in the name could cause assertion failure. 2) `ALTER DATABASE ... UPGRADE DATA DIRECTORY NAME` failed for databases that had a `#mysql50#` prefix if there were triggers in the database. 3) `mysqlcheck --fix-table-name` did not use `utf8` as the default character set, resulting in parsing errors for tables with nonlatin symbols in their names and trigger definitions. (Bug #33094, Bug #41385)
- Execution of a prepared statement that referred to a system variable caused a server crash. (Bug #32124)
- Some division operations produced a result with incorrect precision. (Bug #31616)
- Queries executed using join buffering of `BIT` columns could produce incorrect results. (Bug #31399)
- `ALTER TABLE CONVERT TO CHARACTER SET` did not convert `TINYTEXT` or `MEDIUMTEXT` columns to a longer text type if necessary when converting the column to a different character set. (Bug #31291)
- Server variables could not be set to their current values on Linux platforms. (Bug #31177)

References: See also: Bug #6958.

- For installation on Solaris using `pkgadd` packages, the `mysql_install_db` script was generated in the `scripts` directory, but the temporary files used during the process were left there and not deleted. (Bug #31052)

- Static storage engines and plugins that were disabled and dynamic plugins that were installed but disabled were not listed in the `INFORMATION_SCHEMA` appropriate `PLUGINS` or `ENGINES` table. (Bug #29263)
- Some `SHOW` statements and retrievals from the `INFORMATION_SCHEMA` `TRIGGERS` and `EVENTS` tables used a temporary table and incremented the `Created_tmp_disk_tables` status variable, due to the way that `TEXT` columns are handled. The `TRIGGERS.SQL_MODE`, `TRIGGERS.DEFINER`, and `EVENTS.SQL_MODE` columns now are `VARCHAR` to avoid this problem. (Bug #29153)
- For several read only system variables that were viewable with `SHOW VARIABLES`, attempting to view them with `SELECT @@var_name` or set their values with `SET` resulted in an `unknown system variable` error. Now they can be viewed with `SELECT @@var_name` and attempting to set their values results in a message indicating that they are read only. (Bug #28234)
- On Windows, Visual Studio does not take into account some x86 hardware limitations, which led to incorrect results converting large `DOUBLE` values to unsigned `BIGINT` values. (Bug #27483)
- SSL support was not included in some “generic” RPM packages. (Bug #26760)
- The `Questions` status variable is intended as a count of statements sent by clients to the server, but was also counting statements executed within stored routines. (Bug #24289)
- Setting the session value of the `max_allowed_packet` or `net_buffer_length` system variable was permitted but had no effect. The session value of these variables is now read only. (Bug #22891)

References: See also: Bug #32223.

- A race condition between the `mysqld.exe` server and the Windows service manager could lead to inability to stop the server from the service manager. (Bug #20430)
- On Windows, moving an `InnoDB .ibd` file and then symlinking to it in the database directory using a `.sym` file caused a server crash. (Bug #11894)

Changes in MySQL 5.1.30 (2008-11-14, General Availability)

Bugs Fixed

- **Partitioning:** A `SELECT` using a range `WHERE` condition with an `ORDER BY` on a partitioned table caused a server crash. (Bug #40494)
- **Replication:** Executing `SHOW BINLOG EVENTS` increased the value of `max_allowed_packet` applying to the session that executed the statement. (Bug #55322)
- **Replication:** Row-based replication failed with nonpartitioned `MyISAM` tables having no indexes. (Bug #40004)
- With statement-based binary logging format and a transaction isolation level of `READ COMMITTED` or stricter, `InnoDB` printed an error because statement-based logging might lead to inconsistency between master and slave databases. However, this error was printed even when binary logging was not enabled (in which case, no such inconsistency can occur). (Bug #40360)
- The `CHECK TABLE ... FOR UPGRADE` statement did not check for incompatible collation changes made in MySQL 5.1.24 (Bug #27877). This also affects `mysqlcheck` and `mysql_upgrade`, which cause that statement to be executed. See [Checking Whether Tables or Indexes Must Be Rebuilt](#).

Prior to this fix, a binary upgrade (performed without dumping tables with `mysqldump` before the upgrade and reloading the dump file after the upgrade) would corrupt tables that have indexes that use the `utf8_general_ci` or `ucs2_general_ci` collation for columns that contain 'ß' LATIN SMALL

LETTER SHARP S (German). After the fix, `CHECK TABLE ... FOR UPGRADE` properly detects the problem and warns about tables that need repair.

However, the fix is not backward compatible and can result in a downgrading problem under these circumstances:

1. Perform a binary upgrade to a version of MySQL that includes the fix.
2. Run `CHECK TABLE ... FOR UPGRADE` (or `mysqlcheck` or `mysql_upgrade`) to upgrade tables.
3. Perform a binary downgrade to a version of MySQL that does not include the fix.

The solution is to dump tables with `mysqldump` before the downgrade and reload the dump file after the downgrade. Alternatively, drop and recreate affected indexes. (Bug #40053)

References: See also: Bug #27877.

- Some recent releases for Solaris 10 were built on Solaris 10 U5, which included a new version of `libnsl.so` that does not work on U4 or earlier. To correct this, Solaris 10 builds now are created on machines that do not have that upgraded `libnsl.so`, so that they will work on Solaris 10 installations both with and without the upgraded `libnsl.so`. (Bug #39074)
- With binary logging enabled, `CREATE TABLE ... SELECT` and `INSERT INTO ... SELECT` failed if the source table was a log table. (Bug #34306)
- XA transaction rollbacks could result in corrupted transaction states and a server crash. (Bug #28323)
- `ALTER TABLE` for an `ENUM` column could change column values. (Bug #23113)

Changes in MySQL 5.1.29 (2008-10-11)

- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

Functionality Added or Changed

- **Important Change:** The `--skip-thread-priority` option is now deprecated such that the server will not change the thread priorities by default. Giving threads different priorities might yield marginal improvements in some platforms (where it actually works), but it might instead cause significant degradation depending on the thread count and number of processors. Meddling with the thread priorities is a not a safe bet as it is very dependent on the behavior of the CPU scheduler and system where MySQL is being run. (Bug #35164, Bug #37536)
- **Important Change:** The `--log` option now is deprecated and will be removed (along with the `log` system variable) in the future. Instead, use the `--general_log` option to enable the general query log and the `--general_log_file=file_name` option to set the general query log file name. The values of these options are available in the `general_log` and `general_log_file` system variables, which can be changed at runtime.

Similar changes were made for the `--log-slow-queries` option and `log_slow_queries` system variable. You should use the `--slow_query_log` and `--slow_query_log_file=file_name` options instead (and the `slow_query_log` and `slow_query_log_file` system variables).

- The `BUILD/compile-solaris-*` scripts now compile MySQL with the `mtmalloc` library rather than `malloc`. (Bug #38727)

Bugs Fixed

- **Incompatible Change; Replication:** The default binary logging mode has been changed from `MIXED` to `STATEMENT` for compatibility with MySQL 5.0. (Bug #39812)
- **Incompatible Change:** `CHECK TABLE ... FOR UPGRADE` did not check for incompatible collation changes made in MySQL 5.1.21 (Bug #29499) and 5.1.23 (Bug #27562, Bug #29461). This also affects `mysqlcheck` and `mysql_upgrade`, which cause that statement to be executed. See [Checking Whether Tables or Indexes Must Be Rebuilt](#). (Bug #39585)

References: See also: Bug #40984, Bug #27562, Bug #29461, Bug #29499.

- **Incompatible Change:** In connection with view creation, the server created `arc` directories inside database directories and maintained useless copies of `.frm` files there. Creation and renaming procedures of those copies as well as creation of `arc` directories has been discontinued.

This change does cause a problem when downgrading to older server versions which manifests itself under these circumstances:

1. Create a view `v_orig` in MySQL 5.1.29 or higher.
2. Rename the view to `v_new` and then back to `v_orig`.
3. Downgrade to an older 5.1.x server and run `mysql_upgrade`.
4. Try to rename `v_orig` to `v_new` again. This operation fails.

As a workaround to avoid this problem, use either of these approaches:

- Dump your data using `mysqldump` before downgrading and reload the dump file after downgrading.
- Instead of renaming a view after the downgrade, drop it and recreate it.

The downgrade problem introduced by the fix for this bug has been addressed as Bug #40021. (Bug #17823)

References: See also: Bug #40021.

- **Important Change; Replication:** The `SUPER` privilege is now required to change the session value of `binlog_format` as well as its global value. For more information about `binlog_format`, see [Replication Formats](#). (Bug #39106)
- **Partitioning; Replication:** Replication to partitioned `MyISAM` tables could be slow with row-based binary logging. (Bug #35843)
- **Partitioning:** If an error occurred when evaluating a column of a partitioned table for the partitioning function, the row could be inserted anyway. (Bug #38083)
- **Partitioning:** Using `INSERT ... SELECT` to insert records into a partitioned `MyISAM` table could fail if some partitions were empty and others are not. (Bug #38005)
- **Partitioning:** Ordered range scans on partitioned tables were not always handled correctly. In some cases this caused some rows to be returned twice. The same issue also caused `GROUP BY` query results to be aggregated incorrectly. (Bug #30573, Bug #33257, Bug #33555)
- **Replication:** Server code used in binary logging could in some cases be invoked even though binary logging was not actually enabled, leading to asserts and other server errors. (Bug #38798)
- **Replication:** Replication of `BLACKHOLE` tables did not work with row-based binary logging. (Bug #38360)

- **Replication:** In some cases, a replication master sent a special event to a reconnecting slave to keep the slave's temporary tables, but they still had references to the “old” slave SQL thread and used them to access that thread's data. (Bug #38269)
- **Replication:** Replication filtering rules were inappropriately applied when executing `BINLOG` pseudo-queries. One way in which this problem showed itself was that, when replaying a binary log with `mysqlbinlog`, RBR events were sometimes not executed if the `--replicate-do-db` option was specified. Now replication rules are applied only to those events executed by the slave SQL thread. (Bug #36099)
- **Replication:** For a `CREATE TABLE ... SELECT` statement that creates a table in a database other than the current one, the table could be created in the wrong database on replication slaves if row-based binary logging is used. (Bug #34707)
- **Replication:** A statement did not always commit or roll back correctly when the server was shut down; the error could be triggered by having a failing `UPDATE` or `INSERT` statement on a transactional table, causing an implicit rollback. (Bug #32709)

References: See also: Bug #38262.

- The Sun Studio compiler failed to build debug versions of the server due to use of features specific to `gcc`. (Bug #39451)
- For a `TIMESTAMP` column in an `InnoDB` table, testing the column with multiple conditions in the `WHERE` clause caused a server crash. (Bug #39353)
- References to local variables in stored procedures are replaced with `NAME_CONST(name, value)` when written to the binary log. However, an “illegal mix of collation” error might occur when executing the log contents if the value's collation differed from that of the variable. Now information about the variable collation is written as well. (Bug #39182)
- Queries of the form `SELECT ... REGEXP BINARY NULL` could lead to a hung or crashed server. (Bug #39021)
- Statements of the form `INSERT ... SELECT .. ON DUPLICATE KEY UPDATE col_name = DEFAULT` could result in a server crash. (Bug #39002)
- Column names constructed due to wild-card expansion done inside a stored procedure could point to freed memory if the expansion was performed after the first call to the stored procedure. (Bug #38823)
- Repeated `CREATE TABLE ... SELECT` statements, where the created table contained an `AUTO_INCREMENT` column, could lead to an assertion failure. (Bug #38821)
- For deadlock between two transactions that required a timeout to resolve, all server tables became inaccessible for the duration of the deadlock. (Bug #38804)
- When inserting a string into a duplicate-key error message, the server could improperly interpret the string, resulting in a crash. (Bug #38701)
- A race condition between threads sometimes caused unallocated memory to be addressed. (Bug #38692)
- A server crash resulted from concurrent execution of a multiple-table `UPDATE` that used a `NATURAL` or `USING` join together with `FLUSH TABLES WITH READ LOCK` or `ALTER TABLE` for the table being updated. (Bug #38691)
- On ActiveState Perl, `mysql-test-run.pl --start-and-exit` started but did not exit. (Bug #38629)

- An uninitialized variable in the query profiling code was corrected (detected by Valgrind). (Bug #38560)
- A server crash resulted from execution of an `UPDATE` that used a derived table together with `FLUSH TABLES`. (Bug #38499)
- Stored procedures involving substrings could crash the server on certain platforms due to invalid memory reads. (Bug #38469)
- The handlerton-to-plugin mapping implementation did not free handler plugin references when the plugin was uninstalled, resulting in a server crash after several install/uninstall cycles. Also, on Mac OS X, the server crashed when trying to access an `EXAMPLE` table after the `EXAMPLE` plugin was installed. (Bug #37958)
- The server crashed if an argument to a stored procedure was a subquery that returned more than one row. (Bug #37949)
- When analyzing the possible index use cases, the server was incorrectly reusing an internal structure, leading to a server crash. (Bug #37943)
- Access checks were skipped for `SHOW PROCEDURE STATUS` and `SHOW FUNCTION STATUS`, which could lead to a server crash or insufficient access checks in subsequent statements. (Bug #37908)
- The `<=>` operator could return incorrect results when comparing `NULL` to `DATE`, `TIME`, or `DATETIME` values. (Bug #37526)
- The combination of a subquery with a `GROUP BY`, an aggregate function calculated outside the subquery, and a `GROUP BY` on the outer `SELECT` could cause the server to crash. (Bug #37348)
- The `NO_BACKSLASH_ESCAPES` SQL mode was ignored for `LOAD DATA INFILE` and `SELECT INTO ... OUTFILE`. The setting is taken into account now. (Bug #37114)
- In some cases, references to views were confused with references to anonymous tables and privilege checking was not performed. (Bug #36086)
- For crash reports on Windows, symbol names in stack traces were not correctly resolved. (Bug #35987)
- `ALTER EVENT` changed the `PRESERVE` attribute of an event even when `PRESERVE` was not specified in the statement. (Bug #35981)
- Host name values in SQL statements were not being checked for `'@'`, which is illegal according to RFC952. (Bug #35924)
- `mysql_install_db` failed on machines that had the host name set to `localhost`. (Bug #35754)
- Dynamic plugins failed to load on i5/OS. (Bug #35743)
- With the `PAD_CHAR_TO_FULL_LENGTH` SQL mode enabled, a `ucs2 CHAR` column returned additional garbage after trailing space characters. (Bug #35720)
- A trigger for an `InnoDB` table activating multiple times could lead to `AUTO_INCREMENT` gaps. (Bug #31612)
- `mysqldump` could fail to dump views containing a large number of columns. (Bug #31434)
- The server could improperly type user-defined variables used in the select list of a query. (Bug #26020)
- For access to the `INFORMATION_SCHEMA.VIEWS` table, the server did not check the `SHOW VIEW` and `SELECT` privileges, leading to inconsistency between output from that table and the `SHOW CREATE VIEW` statement. (Bug #22763)

- `mysqld_safe` would sometimes fail to remove the pid file for the old `mysql` process after a crash. As a result, the server failed to start due to a false `A mysqld process already exists...` error. (Bug #11122)

Changes in MySQL 5.1.28 (2008-08-28)

- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

Functionality Added or Changed

- **Important Change:** `mysqlbinlog` now supports `--verbose` and `--base64-output=DECODE-ROWS` options to display row events as commented SQL statements. (The default otherwise is to display row events encoded as base-64 strings using `BINLOG` statements.) See [mysqlbinlog Row Event Display](#). (Bug #31455)
- MySQL source distributions are now available in Zip format. (Bug #27742)
- Added the `SHOW PROFILES` and `SHOW PROFILE` statements to display statement profile data, and the accompanying `INFORMATION_SCHEMA.PROFILING` table. Profiling is controlled using the `profiling` and `profiling_history_size` session variables. see [SHOW PROFILES Syntax](#), and [The INFORMATION_SCHEMA PROFILING Table](#). (Community contribution by Jeremy Cole)

The profiling feature is enabled using the `--enable-community-features` and `--enable-profiling` options to `configure`. These options are enabled by default; to disable them, use `--disable-community-features` and `--disable-profiling`. (Bug #24795)

Bugs Fixed

- **Security Fix; Important Change:** Additional corrections were made for the symlink-related privilege problem originally addressed in MySQL 5.1.24. The original fix did not correctly handle the data directory path name if it contained symlinked directories in its path, and the check was made only at table-creation time, not at table-opening time later.



Note

Additional fixes were made in MySQL 5.1.41.

(Bug #32167, CVE-2008-2079)

References: See also: Bug #39277.

- **Security Enhancement:** The server consumed excess memory while parsing statements with hundreds or thousands of nested boolean conditions (such as `OR (OR ... (OR ...))`). This could lead to a server crash or incorrect statement execution, or cause other client statements to fail due to lack of memory. The latter result constitutes a denial of service. (Bug #38296)
- **Performance; Incompatible Change:** Some performance problems of `SHOW ENGINE INNODB STATUS` were reduced by removing `used cells` and `Total number of lock structs in row lock hash table` from the output. Now these values are present only if the `UNIV_DEBUG` symbol is defined at MySQL build time. (Bug #36941, Bug #36942)
- **Performance:** Over-aggressive lock acquisition by `InnoDB` when calculating free space for tablespaces could result in performance degradation when multiple threads were executing statements on multi-core machines. (Bug #38185)

- **Incompatible Change:** There were some problems using `DllMain()` hook functions on Windows that automatically do global and per-thread initialization for `libmysqld.dll`:
 - Per-thread initialization: MySQL internally counts the number of active threads, which causes a delay in `my_end()` if not all threads have exited. But there are threads that can be started either by Windows internally (often in TCP/IP scenarios) or by users. Those threads do not necessarily use `libmysql.dll` functionality but still contribute to the open-thread count. (One symptom is a five-second delay in times for PHP scripts to finish.)
 - Process-initialization: `my_init()` calls `WSAStartup` that itself loads DLLs and can lead to a deadlock in the Windows loader.

To correct these problems, DLL initialization code now is not invoked from `libmysql.dll` by default. To obtain the previous behavior (DLL initialization code will be called), set the `LIBMYSQL_DLLINIT` environment variable to any value. This variable exists only to prevent breakage of existing Windows-only applications that do not call `mysql_thread_init()` and work okay today. Use of `LIBMYSQL_DLLINIT` is discouraged and is removed in MySQL 6.0. (Bug #37226, Bug #33031)

- **Incompatible Change:** `SHOW STATUS` took a lot of CPU time for calculating the value of the `Innodb_buffer_pool_pages_latched` status variable. Now this variable is calculated and included in the output of `SHOW STATUS` only if the `UNIV_DEBUG` symbol is defined at MySQL build time. (Bug #36600)
- **Incompatible Change:** An additional correction to the original MySQL 5.1.23 fix was made to normalize directory names before adding them to the list of directories. This prevents `/etc/` and `/etc` from being considered different, for example. (Bug #20748)

References: See also: Bug #38180.

- **Partitioning:** When a partitioned table had a `TIMESTAMP` column defined with `CURRENT_TIMESTAMP` as the default but with no `ON UPDATE` clause, the column's value was incorrectly set to `CURRENT_TIMESTAMP` when updating across partitions. (Bug #38272)
- **Partitioning:** `myisamchk` failed with an assertion error when analyzing a partitioned `MyISAM` table. (Bug #37537)
- **Partitioning:** A `LIST` partitioned `MyISAM` table returned erroneous results when an index was present on a column in the `WHERE` clause and `NOT IN` was used on that column.

Searches using the index were also much slower then if the index were not present. (Bug #35931)

- **Partitioning:** `SELECT COUNT(*)` was not correct for some partitioned tables using a storage engine that did not support `HA_STATS_RECORDS_IS_EXACT`. Tables using the `ARCHIVE` storage engine were known to be affected.

This was because `ha_partition::records()` was not implemented, and so the default `handler::records()` was used in its place. However, this is not correct behavior if the storage engine does not support `HA_STATS_RECORDS_IS_EXACT`.

The solution was to implement `ha_partition::records()` as a wrapper around the underlying partition records.

As a result of this fix, the `rows` column in the output of `EXPLAIN PARTITIONS` now includes the total number of records in the partitioned table. (Bug #35745)

- **Partitioning:** `MyISAM` recovery enabled with the `--myisam-recover` option did not work for partitioned `MyISAM` tables. (Bug #35161)

- **Partitioning:** When one user was in the midst of a transaction on a partitioned table, a second user performing an `ALTER TABLE` on this table caused the server to hang. (Bug #34604)
- **Partitioning:** Attempting to execute an `INSERT DELAYED` statement on a partitioned table produced the error `Table storage engine for 'table' doesn't have this option`, which did not reflect the source of the error accurately. The error message returned in such cases has been changed to `DELAYED option not supported for table 'table'`. (Bug #31210)
- **Replication:** Some kinds of internal errors, such as `Out of memory` errors, could cause the server to crash when replicating statements with user variables.

certain internal errors. (Bug #37150)

- **Replication:** Row-based replication did not correctly copy `TIMESTAMP` values from a big-endian storage engine to a little-endian storage engine. (Bug #37076)
- **Replication:** `INSTALL PLUGIN` and `UNINSTALL PLUGIN` caused row-based replication to fail.

**Note**

These statements are not replicated; however, when using row-based logging, the changes they introduce in the `mysql` system tables are written to the binary log.

(Bug #35807)

- Server-side cursors were not initialized properly, which could cause a server crash. (Bug #38486)
- A server crash or Valgrind warnings could result when a stored procedure selected from a view that referenced a function. (Bug #38291)
- A failure to clean up binary log events was corrected (detected by Valgrind). (Bug #38290)
- Incorrect handling of aggregate functions when loose index scan was used caused a server crash. (Bug #38195)
- Queries containing a subquery with `DISTINCT` and `ORDER BY` could cause a server crash. (Bug #38191)
- The fix for Bug #20748 caused a problem such that on Unix, MySQL programs looked for options in `~/my.cnf` rather than the standard location of `~/.my.cnf`. (Bug #38180)

References: See also: Bug #20748.

- If the table definition cache contained tables with many `BLOB` columns, much memory could be allocated to caching `BLOB` values. Now a size limit on the cached `BLOB` values is enforced. (Bug #38002)
- For `InnoDB` tables, `ORDER BY ... DESC` sometimes returned results in ascending order. (Bug #37830)
- If a table has a `BIT NOT NULL` column `c1` with a length shorter than 8 bits and some additional `NOT NULL` columns `c2`, ..., and a `SELECT` query has a `WHERE` clause of the form `(c1 = constant) AND c2 ...`, the query could return an unexpected result set. (Bug #37799)
- The server returned unexpected results if a right side of the `NOT IN` clause consisted of the `NULL` value and some constants of the same type. For example, this query might return 3, 4, 5, and so forth if a table contained those values:

```
SELECT * FROM t WHERE NOT t.id IN (NULL, 1, 2);
```

(Bug #37761)

- Setting the session value of the `innodb_table_locks` system variable caused a server crash. (Bug #37669)
- Nesting of `IF ()` inside of `SUM()` could cause an extreme server slowdown. (Bug #37662)
- Killing a query that used an `EXISTS` subquery as the argument to `SUM()` or `AVG()` caused a server crash. (Bug #37627)
- When using indexed `ORDER BY` sorting, incorrect query results could be produced if the optimizer switched from a covering index to a noncovering index. (Bug #37548)
- After `TRUNCATE TABLE` for an `InnoDB` table, inserting explicit values into an `AUTO_INCREMENT` column could fail to increment the counter and result in a duplicate-key error for subsequent insertion of `NULL`. (Bug #37531)
- Within stored programs or prepared statements, `REGEXP` could return incorrect results due to improper initialization. (Bug #37337)
- For a `MyISAM` table with `CHECKSUM = 1` and `ROW_FORMAT = DYNAMIC` table options, a data consistency check (maximum record length) could fail and cause the table to be marked as corrupted. (Bug #37310)
- The `max_length` result set metadata value was calculated incorrectly under some circumstances. (Bug #37301)
- If the length of a field was 3, internal `InnoDB` to integer type conversion didn't work on big-endian machines in the `row_search_autoinc_column()` function. (Bug #36793)
- A query which had an `ORDER BY DESC` clause that is satisfied with a reverse range scan could cause a server crash for some specific CPU/compiler combinations. (Bug #36639)
- The `CSV` storage engine returned success even when it failed to open a table's data file. (Bug #36638)
- `SELECT DISTINCT` from a simple view on an `InnoDB` table, where all selected columns belong to the same unique index key, returned incorrect results. (Bug #36632)
- Dumping information about locks in use by sending a `SIGHUP` signal to the server or by invoking the `mysqladmin debug` command could lead to a server crash in debug builds or to undefined behavior in production builds. (Bug #36579)
- If initialization of an `INFORMATION_SCHEMA` plugin failed, `INSTALL PLUGIN` freed some internal plugin data twice. (Bug #36399)
- For `InnoDB` tables, the `DATA_FREE` column of the `INFORMATION_SCHEMA.TABLES` displayed free space in kilobytes rather than bytes. Now it displays bytes. (Bug #36278)
- When the fractional part in a multiplication of `DECIMAL` values overflowed, the server truncated the first operand rather than the longest. Now the server truncates so as to produce more precise multiplications. (Bug #36270)
- The `mysql` client failed to recognize comment lines consisting of `--` followed by a newline. (Bug #36244)
- The server could crash with an assertion failure (or cause the client to get a "Packets out of order" error) when the expected query result was that it should terminate with a "Subquery returns more than 1 row" error. (Bug #36135)

- The `UUID()` function returned UUIDs with the wrong time; this was because the offset for the time part in UUIDs was miscalculated. (Bug #35848)
- The `configure` script did not permit `utf8_hungarian_ci` to be specified as the default collation. (Bug #35808)
- On 64-bit systems, assigning values of $2^{63} - 1$ or larger to `key_buffer_size` caused memory overruns. (Bug #35616)
- For `InnoDB` tables, `REPLACE` statements used “traditional” style locking, regardless of the setting of `innodb_autoinc_lock_mode`. Now `REPLACE` works the same way as “simple inserts” instead of using the old locking algorithm. (`REPLACE` statements are treated in the same way as `INSERT` statements.) (Bug #35602)
- Freeing of an internal parser stack during parsing of complex stored programs caused a server crash. (Bug #35577, Bug #37269, Bug #37228)
- `mysqlbinlog` left temporary files on the disk after shutdown, leading to the pollution of the temporary directory, which eventually caused `mysqlbinlog` to fail. This caused problems in testing and other situations where `mysqlbinlog` might be invoked many times in a relatively short period of time. (Bug #35543)
- Index scans performed with the `sort_union()` access method returned wrong results, caused memory to be leaked, and caused temporary files to be deleted when the limit set by `sort_buffer_size` was reached. (Bug #35477, Bug #35478)
- Table checksum calculation could cause a server crash for `FEDERATED` tables with `BLOB` columns containing `NULL` values. (Bug #34779)
- A significant slowdown occurred when many `SELECT` statements that return many rows from `InnoDB` tables were running concurrently. (Bug #34409)
- `mysql_install_db` failed if the server was running with an SQL mode of `TRADITIONAL`. This program now resets the SQL mode internally to avoid this problem. (Bug #34159)
- Changes to build files were made to enable the MySQL distribution to compile on Microsoft Visual C++ Express 2008. (Bug #33907)
- Fast `ALTER TABLE` operations were not fast for columns that used multibyte character sets. (Bug #33873)
- The internal functions `my_getsystime()`, `my_micro_time()`, and `my_micro_time_and_time()` did not work correctly on Windows. One symptom was that uniqueness of `UUID()` values could be compromised. (Bug #33748)
- Cached queries that used 256 or more tables were not properly cached, so that later query invalidation due to a `TRUNCATE TABLE` for one of the tables caused the server to hang. (Bug #33362)
- `mysql_upgrade` attempted to use the `/proc` file system even on systems that do not have it. (Bug #31605)
- `mysql_install_db` failed if the default storage engine was `NDB`. Now it explicitly uses `MyISAM` as the storage engine when running `mysqld --bootstrap`. (Bug #31315)
- Several MySQL programs could fail if the `HOME` environment variable had an empty value. (Bug #30394)
- On NetWare, `mysql_install_db` could appear to execute normally even if it failed to create the initial databases. (Bug #30129)

- The Serbian translation for the `ER_INCORRECT_GLOBAL_LOCAL_VAR` error was corrected. (Bug #29738)
- `TRUNCATE TABLE` for InnoDB tables returned a count showing too many rows affected. Now the statement returns 0 for InnoDB tables. (Bug #29507)
- The `BUILD/check-cpu` build script failed if `gcc` had a different name (such as `gcc.real` on Debian). (Bug #27526)
- In some cases, the parser interpreted the `;` character as the end of input and misinterpreted stored program definitions. (Bug #26030)
- The `FLUSH PRIVILEGES` statement did not produce an error when it failed. (Bug #21226)
- After executing a prepared statement that accesses a stored function, the next execution failed to find the function if the stored function cache was flushed in the meantime. (Bug #12093, Bug #21294)

Changes in MySQL 5.1.27 (Not released)

- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

Functionality Added or Changed

- `mysqltest` now installs signal handlers and generates a stack trace if it crashes. (Bug #37003)
- `mysql-test-run.pl` now supports `--client-bindir` and `--client-libdir` options for specifying the directory where client binaries and libraries are located. (Bug #34995)

Bugs Fixed

- **Incompatible Change; Partitioning:** On Mac OS X, with `lower_case_table_names = 2`, the server could not read partitioned tables whose names contained uppercase letters.

Partitioned tables using mixed case names should be renamed or dropped before upgrading to this version of the server on Mac OS X. (Bug #37402)

- **Important Change; Partitioning:** The statements `ANALYZE TABLE`, `CHECK TABLE`, `OPTIMIZE TABLE`, and `REPAIR TABLE` are now supported for partitioned tables.

Also as a result of this fix, the following statements which were disabled in MySQL 5.1.24 have been re-enabled:

- `ALTER TABLE ... ANALYZE PARTITION`
- `ALTER TABLE ... CHECK PARTITION`
- `ALTER TABLE ... OPTIMIZE PARTITION`
- `ALTER TABLE ... REPAIR PARTITION`

(Bug #20129)

References: See also: Bug #39434.

- **Replication:** Issuing a `DROP DATABASE` while any temporary tables were open caused the server to switch to statement-based mode. (Bug #38773)

- **Replication:** The `--replicate-*-table` options were not evaluated correctly when replicating multi-table updates.

As a result of this fix, replication of multi-table updates no longer fails when an update references a missing table but does not update any of its columns. (Bug #37051)

- The fix for Bug #33812 had the side effect of causing the `mysql` client not to be able to read some dump files produced with `mysqldump`. To address this, that fix was reverted. (Bug #38158)

References: Reverted patches: Bug #33812.

Changes in MySQL 5.1.26 (2008-06-30)

- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

Functionality Added or Changed

- **Incompatible Change; Important Change:** The `FEDERATED` storage engine is now disabled by default in binary distributions. The engine is still available and can be enabled by starting the server with the `--federated` option. (Bug #37069)
- `mysqltest` was changed to be more robust in the case of a race condition that can occur for rapid disconnect/connect sequences with the server. The account used by `mysqltest` could reach its permitted simultaneous-sessions user limit if the connect attempt occurred before the server had fully processed the preceding disconnect. `mysqltest` now checks specifically for a user-limits error when it connects; if that error occurs, it delays briefly before retrying. (Bug #23921)

Bugs Fixed

- **Replication:** Row-based replication broke for `utf8 CHAR` columns longer than 85 characters. (Bug #37426)
- **Replication:** Performing an insert on a table having an `AUTO_INCREMENT` column and an `INSERT` trigger that was being replicated from a master running MySQL 5.0 or any version of MySQL 5.1 up to and including MySQL 5.1.11 to a slave running MySQL 5.1.12 or later caused the replication slave to crash. (Bug #36443)

References: See also: Bug #33029.

- Some binary distributions had a duplicate “-64bit” suffix in the file name. (Bug #37623)
- `NOT IN` subqueries that selected `MIN()` or `MAX()` values but produced an empty result could cause a server crash. (Bug #37004)
- `ha_innodb.so` was incorrectly installed in the `lib/mysql` directory rather than in `lib/mysql/plugin`. (Bug #36434)
- An empty bit-string literal (`b''`) caused a server crash. Now the value is parsed as an empty bit value (which is treated as an empty string in string context or 0 in numeric context). (Bug #35658)
- The code for detecting a byte order mark (BOM) caused `mysql` to crash for empty input. (Bug #35480)
- The `mysql` client incorrectly parsed statements containing the word “delimiter” in mid-statement.

The fix for this bug had the side effect of causing the problem reported in Bug #38158, so it was reverted in MySQL 5.1.27. (Bug #33812)

References: See also: Bug #38158.

Changes in MySQL 5.1.25 (2008-05-28)

- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

Functionality Added or Changed

- **Incompatible Change:** A change has been made to the way that the server handles prepared statements. This affects prepared statements processed at the SQL level (using the `PREPARE` statement) and those processed using the binary client/server protocol (using the `mysql_stmt_prepare()` C API function).

Previously, changes to metadata of tables or views referred to in a prepared statement could cause a server crash when the statement was next executed, or perhaps an error at execute time with a crash occurring later. For example, this could happen after dropping a table and recreating it with a different definition.

Now metadata changes to tables or views referred to by prepared statements are detected and cause automatic reparation of the statement when it is next executed.

The server attempts reparation up to three times. An error occurs if all attempts fail.

Metadata changes occur for DDL statements such as those that create, drop, alter, rename, or truncate tables, or that analyze, optimize, or repair tables. Reparation also occurs after referenced tables or views are flushed from the table definition cache, either implicitly to make room for new entries in the cache, or explicitly due to `FLUSH TABLES`.

Table content changes (for example, with `INSERT` or `UPDATE`) do not cause reparation, nor do `SELECT` statements.

Reparation is automatic, but to the extent that it occurs, performance of prepared statements is diminished.

Reparation uses the default database and SQL mode that were in effect for the original preparation.

A status variable, `Com_stmt_reprepare`, has been introduced to track the number of reparations.

Along with this change to statement reparation, the default value of the `table_definition_cache` system variable has been increased from 128 to 256. The purpose of this increase is to lessen the chance that prepared statements will need reparation due to referred-to tables/views having been flushed from the cache to make room for new entries.

An incompatibility with previous versions of MySQL is that a prepared statement may now return a different set of columns or different column types from one execution to the next. For example, if the prepared statement is `SELECT * FROM t1`, altering `t1` to contain a different number of columns causes the next execution to return a number of columns different from the previous execution.

Older versions of the client library cannot handle this change in behavior. For applications that use prepared statements with the new server, an upgrade to the new client library is strongly recommended. (Bug #27420, Bug #27430, Bug #27690)

- **Important Change:** Some changes were made to `CHECK TABLE ... FOR UPGRADE` and `REPAIR TABLE` with respect to detection and handling of tables with incompatible `.frm` files (files created with a

different version of the MySQL server). These changes also affect `mysqlcheck` because that program uses `CHECK TABLE` and `REPAIR TABLE`, and thus also `mysql_upgrade` because that program invokes `mysqlcheck`.

- If your table was created by a different version of the MySQL server than the one you are currently running, `CHECK TABLE ... FOR UPGRADE` indicates that the table has an `.frm` file with an incompatible version. In this case, the result set returned by `CHECK TABLE` contains a line with a `Msg_type` value of `error` and a `Msg_text` value of `Table upgrade required. Please do "REPAIR TABLE `tbl_name`" to fix it!`
- `REPAIR TABLE` without `USE_FRM` upgrades the `.frm` file to the current version.
- If you use `REPAIR TABLE ...USE_FRM` and your table was created by a different version of the MySQL server than the one you are currently running, `REPAIR TABLE` will not attempt to repair the table. In this case, the result set returned by `REPAIR TABLE` contains a line with a `Msg_type` value of `error` and a `Msg_text` value of `Failed repairing incompatible .FRM file`.

Previously, use of `REPAIR TABLE ...USE_FRM` with a table created by a different version of the MySQL server risked the loss of all rows in the table.

(Bug #36055)

- `mysql_upgrade` now has a `--tmpdir` option to enable the location of temporary files to be specified. (Bug #36469)
- `mysqldump` now adds the `LOCAL` qualifier to the `FLUSH TABLES` statement that is sent to the server when the `--master-data` option is enabled. This prevents the `FLUSH TABLES` statement from replicating to slaves, which is disadvantageous because it would cause slaves to block while the statement executes. (Bug #35157)

References: See also: Bug #38303.

Bugs Fixed

- **Important Change:** The server no longer issues warnings for truncation of excess spaces for values inserted into `CHAR` columns. This reverts a change in the previous release that caused warnings to be issued. (Bug #30059)
- **Replication:** `CREATE PROCEDURE` and `CREATE FUNCTION` statements containing extended comments were not written to the binary log correctly, causing parse errors on the slave. (Bug #36570)

References: See also: Bug #32575.

- **Replication:** When flushing tables, there was a slight chance that the flush occurred between the processing of one table map event and the next. Since the tables were opened one by one, subsequent locking of tables would cause the slave to crash. This problem was observed when replicating `NDBCLUSTER` or `InnoDB` tables, when executing multi-table updates, and when a trigger or a stored routine performed an (additional) insert on a table so that two tables were effectively being inserted into in the same statement. (Bug #36197)
- **Replication:** `CREATE VIEW` statements containing extended comments were not written to the binary log correctly, causing parse errors on the slave. Now, all comments are stripped from such statements before being written to the binary log. (Bug #32575)

References: See also: Bug #36570.

- On Windows 64-bit systems, temporary variables of `long` types were used to store `ulong` values, causing key cache initialization to receive distorted parameters. The effect was that setting

`key_buffer_size` to values of 2GB or more caused memory exhaustion to due allocation of too much memory. (Bug #36705)

- Multiple-table `UPDATE` statements that used a temporary table could fail to update all qualifying rows or fail with a spurious duplicate-key error. (Bug #36676)
- A `REGEXP` match could return incorrect rows when the previous row matched the expression and used `CONCAT()` with an empty string. (Bug #36488)
- `mysqltest` ignored the value of `--tmpdir` in one place. (Bug #36465)
- When updating an existing instance (for example, from MySQL 5.0 to 5.1, or 5.1 to 6.0), the Instance Configuration Wizard unnecessarily prompted for a `root` password when there was an existing `root` password. (Bug #36305)
- Conversion of a `FLOAT ZEROFILL` value to string could cause a server crash if the value was `NULL`. (Bug #36139)
- On Windows, the installer attempted to use JScript to determine whether the target data directory already existed. On Windows Vista x64, this resulted in an error because the installer was attempting to run the JScript in a 32-bit engine, which wasn't registered on Vista. The installer no longer uses JScript but instead relies on a native WiX command. (Bug #36103)
- `mysqltest` was performing escape processing for the `--replace_result` command, which it should not have been. (Bug #36041)
- An error in calculation of the precision of zero-length items (such as `NULL`) caused a server crash for queries that employed temporary tables. (Bug #36023)
- For `EXPLAIN EXTENDED`, execution of an uncorrelated `IN` subquery caused a crash if the subquery required a temporary table for its execution. (Bug #36011)
- The `MERGE` storage engine did a table scan for `SELECT COUNT(*)` statements when it could calculate the number of records from the underlying tables. (Bug #36006)
- The server crashed inside `NOT IN` subqueries with an impossible `WHERE` or `HAVING` clause, such as `NOT IN (SELECT ... FROM t1, t2, ... WHERE 0)`. (Bug #36005)
- The Event Scheduler was not designed to work under the embedded server. It is now disabled for the embedded server, and the `event_scheduler` system variable is not displayed. (Bug #35997)
- Grouping or ordering of long values in unindexed `BLOB` or `TEXT` columns with the `gbk` or `big5` character set crashed the server. (Bug #35993)
- `SET GLOBAL debug= ''` resulted in a Valgrind warning in `DbugParse()`, which was reading beyond the end of the control string. (Bug #35986)
- The “prefer full scan on clustered primary key over full scan of any secondary key” optimizer rule introduced by Bug #26447 caused a performance regression for some queries, so it has been disabled. (Bug #35850)

References: See also: Bug #26447.

- The server ignored any covering index used for `ref` access of a table in a query with `ORDER BY` if this index was incompatible with the `ORDER BY` list and there was another covering index compatible with this list. As a result, suboptimal execution plans were chosen for some queries that used an `ORDER BY` clause. (Bug #35844)

- `mysql_upgrade` did not properly update the `mysql.event` table. (Bug #35824)
- An incorrect error and message was produced for attempts to create a `MyISAM` table with an index (`.MYI`) file name that was already in use by some other `MyISAM` table that was open at the same time. For example, this might happen if you use the same value of the `INDEX DIRECTORY` table option for tables belonging to different databases. (Bug #35733)
- Enabling the `read_only` system variable while `autocommit` mode was enabled caused `SELECT` statements for transactional storage engines to fail. (Bug #35732)
- The combination of `GROUP_CONCAT()`, `DISTINCT`, and `LEFT JOIN` could crash the server when the right table is empty. (Bug #35298)
- Some binaries produced stack corruption messages due to being built with versions of `bison` older than 2.1. Builds are now created using `bison` 2.3. (Bug #34926)
- The `log_output` system variable could be set to an illegal value. (Bug #34820)
- On Windows 64-bit builds, an apparent compiler bug caused memory overruns for code in `innobase/mem/*`. Removed optimizations so as not to trigger this problem. (Bug #34297)
- Several additional configuration scripts in the `BUILD` directory now are included in source distributions. These may be useful for users who wish to build MySQL from source. (See [Installing MySQL Using a Development Source Tree](#), for information about what they do.) (Bug #34291)
- Executing a `FLUSH PRIVILEGES` statement after creating a temporary table in the `mysql` database with the same name as one of the MySQL system tables caused the server to crash.

**Note**

While it is possible to shadow a system table in this way, the temporary table exists only for the current user and connection, and does not effect any user privileges.

(Bug #33275)

- `UNION` constructs cannot contain `SELECT ... INTO` except in the final `SELECT`. However, if a `UNION` was used in a subquery and an `INTO` clause appeared in the top-level query, the parser interpreted it as having appeared in the `UNION` and raised an error. (Bug #32858)
- Assignment of relative path names to `general_log_file` or `slow_query_log_file` did not always work. (Bug #32748)
- The `mysql.servers` table was not created during installation on Windows. (Bug #28680, Bug #32797)
- The `jp` test suite was not working. (Bug #28563)
- The internal `init_time()` library function was renamed to `my_init_time()` to avoid conflicts with external libraries. (Bug #26294)
- The parser used signed rather than unsigned values in some cases that caused legal lengths in column declarations to be rejected. (Bug #15776)

Changes in MySQL 5.1.24 (2008-04-08)

- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

Functionality Added or Changed

- **Important Change; MySQL Cluster; Packaging:** Beginning with this release, standard MySQL 5.1 binaries are no longer built with support for the [NDBCLUSTER](#) storage engine, and the [NDBCLUSTER](#) code included in 5.1 mainline sources is no longer guaranteed to be maintained or supported. Those using MySQL Cluster in MySQL 5.1.23 and earlier MySQL 5.1 mainline releases should upgrade to MySQL Cluster NDB 6.2.15 or a later MySQL Cluster NDB 6.2 or 6.3 release. (Bug #36193)
- **Important Change:** The [FEDERATED](#) storage engine is not included in binary distributions of MySQL 5.1.24. (It will be included again in 5.1.25.)
- **Replication:** Introduced the [slave_exec_mode](#) system variable to control whether idempotent or strict mode is used for replication conflict resolution. Idempotent mode suppresses duplicate-key, no-key-found, and some other errors, and is needed for circular replication, multi-master replication, and some other complex replication setups when using MySQL Cluster, where idempotent mode is the default. However, strict mode is the default for storage engines other than [NDB](#). (Bug #31609)
- **Replication:** When running the server with [--binlog-format=MIXED](#) or [--binlog-format=STATEMENT](#), a query that referred to a system variable used the slave's value when replayed on the slave. This meant that, if the value of a system variable was inserted into a table, the slave differed from the master. Now, statements that refer to a system variable are marked as “unsafe”, which means that:
 - When the server is using [--binlog-format=MIXED](#), the row-based format is used automatically to replicate these statements.
 - When the server is using [--binlog-format=STATEMENT](#), these statements produce a warning. (Bug #31168)

References: See also: Bug #34732.

- The [PROCESS](#) privilege now is required to start or stop the [InnoDB](#) monitor tables (see [InnoDB Monitors](#)). Previously, no privilege was required. (Bug #34053)
- For binary [.tar.gz](#) packages, [mysqld](#) and other binaries now are compiled with debugging symbols included to enable easier use with a debugger. If you do not need debugging symbols and are short on disk space, you can use [strip](#) to remove the symbols from the binaries. (Bug #33252)
- Formerly, when the MySQL server crashed, the generated stack dump was numeric and required external tools to properly resolve the names of functions. This is not very helpful to users having a limited knowledge of debugging techniques. In addition, the generated stack trace contained only the names of functions and was formatted differently for each platform due to different stack layouts.

Now it is possible to take advantage of newer versions of the GNU C Library provide a set of functions to obtain and manipulate stack traces from within the program. On systems that use the ELF binary format, the stack trace contains important information such as the shared object where the call was generated, an offset into the function, and the actual return address. Having the function name also makes possible the name demangling of C++ functions.

The library generates meaningful stack traces on the following platforms: i386, x86_64, PowerPC, IA64, Alpha, and S390. On other platforms, a numeric stack trace is still produced, and the use of the [resolve_stack_dump](#) utility is still required. (Bug #31891)

- [mysqltest](#) now has [mkdir](#) and [rmdir](#) commands for creating and removing directories. (Bug #31004)
- The server uses less memory when loading privileges containing table grants. (Patch provided by Google.) (Bug #25175)

- Added the `Uptime_since_flush_status` status variable, which indicates the number of seconds since the most recent `FLUSH STATUS` statement. (Community contribution by Jeremy Cole) (Bug #24822)
- `SHOW OPEN TABLES` now supports `FROM` and `LIKE` clauses. (Bug #12183)
- The new read-only global system variables `report_host`, `report_password`, `report_port`, and `report_user` system variables provide runtime access to the values of the corresponding `--report-host`, `--report-password`, `--report-port`, and `--report-user` options.
- The `innodb_flush_method` value, `fdatasync`, has been renamed to `fsync`. This change is to avoid confusing the `fdatasync` option name with the `fdatasync()` system call, which is no longer used by InnoDB. As of MySQL 3.23.41, InnoDB uses an `fsync()` system call instead of `fdatasync()` as the default InnoDB flush method.
- The use of InnoDB hash indexes now can be controlled by setting the new `innodb_adaptive_hash_index` system variable at server startup. By default, this variable is enabled. See [Adaptive Hash Indexes](#).

Bugs Fixed

- **Security Fix; Important Change:** It was possible to circumvent privileges through the creation of `MyISAM` tables employing the `DATA DIRECTORY` and `INDEX DIRECTORY` options to overwrite existing table files in the MySQL data directory. Use of the MySQL data directory in `DATA DIRECTORY` and `INDEX DIRECTORY` is no longer permitted. This is now also true of these options when used with partitioned tables and individual partitions of such tables.



Note

Additional fixes were made in MySQL 5.1.28, 5.1.41.

(Bug #32167, CVE-2008-2079)

References: See also: Bug #39277.

- **Security Fix:** A client that connects to a malicious server could be tricked by the server into sending files from the client host to the server. This occurs because the `libmysqlclient` client library would respond to a `FETCH LOCAL FILE` request from the server even if the request is sent for statements from the client other than `LOAD DATA LOCAL INFILE`. The client library has been modified to respond to a `FETCH LOCAL FILE` request from the server only if it is sent in response to a `LOAD DATA LOCAL INFILE` statement from the client.

The client library now also checks whether `CLIENT_LOCAL_FILE` is set and refuses to send a local file if not.



Note

Binary distributions ship with the `local-infile` capability enabled. Applications that do not use this functionality should disable it to be safe.

(Bug #29605)

- **Security Enhancement; Important Change:** On Windows Vista and Windows Server 2008, a user without administrative privileges does not have write permissions to the `Program Files` directory where MySQL and the associated data files are normally installed. Using data files located in the standard `Program Files` installation directory could therefore cause MySQL to fail, or lead to potential security issues in an installed instance.

To address the problem, on Windows XP, Windows Vista and Windows Server 2008, the datafiles and data file configuration are now set to the Microsoft recommended [AppData](#) folder. The [AppData](#) folder is typically located within the user's home directory.

**Important**

When upgrading an existing 5.1.23 or 6.0.4 installation of MySQL you must take a backup of your data and configuration file ([my.ini](#)) before installing the new version. To migrate your data, either extract the data and re-import (using [mysqldump](#), then upgrade and re-import using [mysql](#)), or back up your data, upgrade to the new version, and copy your existing data files from your old [datadir](#) directory to the new directory located within [AppData](#).

Failure to back up your data and follow these procedures may lead to data loss.

(Bug #34593)

- **Performance:** [InnoDB](#) adaptive hash latches could be held too long during filesort operations, resulting in a server crash. Now the hash latch is released when a query on [InnoDB](#) tables performs a filesort. This eliminates the crash and may provide significant performance improvements on systems on which many queries using filesorts with temporary tables are being performed. (Bug #32149)
- **Performance:** [InnoDB](#) exhibited thread thrashing with more than 50 concurrent connections under an update-intensive workload. (Bug #22868)
- **Incompatible Change:** In MySQL 5.1.23, the [last_errno](#) and [last_error](#) members of the [NET](#) structure in [mysql_com.h](#) were renamed to [client_last_errno](#) and [client_last_error](#). This was found to cause problems for connectors that use the internal [NET](#) structure for error handling. The change has been reverted. (Bug #34655)

References: See also: Bug #12713.

- **Incompatible Change:** It was possible to use [FRAC_SECOND](#) as a synonym for [MICROSECOND](#) with [DATE_ADD\(\)](#), [DATE_SUB\(\)](#), and [INTERVAL](#); now, using [FRAC_SECOND](#) with anything other than [TIMESTAMPADD\(\)](#) or [TIMESTAMPDIFF\(\)](#) produces a syntax error.

It is now possible (and preferable) to use [MICROSECOND](#) with [TIMESTAMPADD\(\)](#) and [TIMESTAMPDIFF\(\)](#), and [FRAC_SECOND](#) is now deprecated. (Bug #33834)

- **Incompatible Change:** The [UPDATE](#) statement permitted [NULL](#) to be assigned to [NOT NULL](#) columns (the implicit default value for the column data type was assigned). This was changed so that on error occurs.

This change was reverted, because the original report was determined not to be a bug: Assigning [NULL](#) to a [NOT NULL](#) column in an [UPDATE](#) statement should produce an error only in strict SQL mode and set the column to the implicit default with a warning otherwise, which was the original behavior. See [Data Type Default Values](#), and Bug #39265. (Bug #33699)

References: See also: Bug #39265.

- **Incompatible Change:** For packages that are built within their own prefix (for example, [/usr/local/mysql](#)) the plugin directory will be [lib/plugin](#). For packages that are built to be installed into a system-wide prefix (such as RPM packages with a prefix of [/usr](#)), the plugin directory will be [lib/mysql/plugin](#) to ensure a clean [/usr/lib](#) hierarchy. In both cases, the [\\$pkglibdir](#) configuration setting is used at build time to set the plugin directory.

The current plugin directory location is available as the value of the `plugin_dir` system variable as before, but the `mysql_config` script now has a `--pluginindir` option that can be used externally to the server by third-party plugin writers to obtain the default plugin directory path name and configure their installation directory appropriately. (Bug #31736)

- **Incompatible Change:** The `-`, `*`, and `/` operators and the functions `POW()` and `EXP()` could misbehave when used with floating-point numbers. Previously they might return `+INF`, `-INF`, or `NaN` in cases of numeric overflow (including that caused by division by zero) or when invalid arguments were used. Now `NULL` is returned in all such cases. (Bug #31236)
- **Incompatible Change:** Previously, the parser accepted the ODBC `{ OJ ... LEFT OUTER JOIN ... }` syntax for writing left outer joins. The parser now permits `{ OJ ... }` to be used to write other types of joins, such as `INNER JOIN` or `RIGHT OUTER JOIN`. This helps with compatibility with some third-party applications, but is not official ODBC syntax.

A consequence of this change is that the parser no longer permits nested `{ OJ ... }` constructs (which are not legal ODBC syntax, anyway). Queries that use such constructs should be rewritten. For example, this query is now produces an error:

```
SELECT * FROM
{OJ
  {OJ a LEFT OUTER JOIN b ON a.a1=b.a1}
  LEFT OUTER JOIN c ON b.b1 = c.b1};
```

That can be replaced by any of the following rewrites:

```
SELECT * FROM
  {OJ a LEFT OUTER JOIN b
    LEFT OUTER JOIN c ON b.b1 = c.b1 ON a.a1=b.a1};

SELECT * FROM
  {OJ a LEFT OUTER JOIN b ON a.a1=b.a1
    LEFT OUTER JOIN c ON b.b1 = c.b1};

SELECT * FROM
  a LEFT OUTER JOIN b ON a.a1=b.a1 LEFT OUTER JOIN c ON b.b1 = c.b1;
```

The first two are legal according to ODBC, and you nest the joins inside a single `{ OJ ... }` clause. The third is standard SQL syntax, without ODBC decoration. It can be used with parentheses to emphasize the evaluation order:

```
SELECT * FROM
  ((a LEFT OUTER JOIN b ON a.a1=b.a1)
   LEFT OUTER JOIN c ON b.b1 = c.b1);
```

(Bug #28317)

- **Incompatible Change:** The `utf8_general_ci` and `ucs2_general_ci` collations did not sort the letter "U+00DF SHARP S" equal to 's'.

As a result of this bug fix, indexes must be rebuilt for columns that use the `utf8_general_ci` or `ucs2_general_ci` collation for columns that contain SHARP S. See [Checking Whether Tables or Indexes Must Be Rebuilt](#). (Bug #27877)

References: See also: Bug #37046.

- **Important Change; Partitioning:** The following statements did not function correctly with corrupted or crashed tables and have been disabled:

- `ALTER TABLE ... ANALYZE PARTITION`
- `ALTER TABLE ... CHECK PARTITION`
- `ALTER TABLE ... OPTIMIZE PARTITION`
- `ALTER TABLE ... REPAIR PARTITION`

`ALTER TABLE ... REBUILD PARTITION` is unaffected by this change and continues to be available. This statement and `ALTER TABLE ... REORGANIZE PARTITION` may be used to analyze and optimize partitioned tables, since these operations cause the partition files to be rebuilt. (Bug #20129)

References: See also: Bug #39434.

- **Important Change; Replication:** When the master crashed during an update on a transactional table while in `autocommit` mode, the slave failed. This fix causes every transaction (including `autocommit` transactions) to be recorded in the binary log as starting with a `BEGIN` and ending with a `COMMIT` or `ROLLBACK`.



Note

The current fix does *not* cause nontransactional changes to be wrapped in `BEGIN ... COMMIT` or `BEGIN ... ROLLBACK` when written to the binary log. For this purpose, any statements affecting tables using a nontransactional storage engine such as `MyISAM` are regarded as nontransactional, even when `autocommit` is enabled.

(Bug #26395)

References: See also: Bug #29288, Bug #49522.

- **Important Change:** `InnoDB` free space information is now shown in the `Data_free` column of `SHOW TABLE STATUS` and in the `DATA_FREE` column of the `INFORMATION_SCHEMA.TABLES` table. (Bug #32440)

References: See also: Bug #11379.

- **Important Change:** The server handled truncation of values having excess trailing spaces into `CHAR`, `VARCHAR`, and `TEXT` columns in different ways. This behavior has now been made consistent for columns of all three of these types, and now follows the existing behavior of `VARCHAR` columns in this regard; that is, a `Note` is always issued whenever such truncation occurs.

This change does not affect columns of these three types when using a binary encoding; `BLOB` columns are also unaffected by the change, since they always use a binary encoding. (Bug #30059)

- **Important Change:** An `AFTER UPDATE` trigger was not invoked when the `UPDATE` did not make any changes to the table for which the trigger was defined. Now `AFTER UPDATE` triggers behave the same in this regard as do `BEFORE UPDATE` triggers, which are invoked whether the `UPDATE` makes any changes in the table or not. (Bug #23771)
- **Important Note; Replication:** Network timeouts between the master and the slave could result in corruption of the relay log. This fix rectifies a long-standing replication issue when using unreliable networks, including replication over wide area networks such as the Internet. If you experience reliability

issues and see many `You have an error in your SQL syntax` errors on replication slaves, we strongly recommend that you upgrade to a MySQL version which includes this fix. (Bug #26489)

- **Partitioning:** In some cases, matching rows from a partitioned `MyISAM` using a `BIT` column as the primary key were not found by queries. (Bug #34358)
- **Partitioning:** Enabling `innodb_file_per_table` produced problems with partitioning and tablespace operations on partitioned `InnoDB` tables, in some cases leading to corrupt partitions or causing the server to crash. (Bug #33429)
- **Partitioning:** A table defined using `PARTITION BY KEY` and having a `BIT` column referenced in the partitioning key did not behave correctly; some rows could be inserted into the wrong partition, causing wrong results to be returned from queries. (Bug #33379)
- **Partitioning:** For `InnoDB` tables, there was a race condition involving the data dictionary and repartitioning. (Bug #33349)
- **Partitioning:** When `ALTER TABLE DROP PARTITION` was executed on a table on which there was a trigger, the statement failed with an error. This occurred even if the trigger did not reference any tables. (Bug #32943)
- **Partitioning:** Currently, all partitions of a partitioned table must use the same storage engine. One may optionally specify the storage engine on a per-partition basis; however, where this is the done, the storage engine must be the same as used by the table as a whole. `ALTER TABLE` did not enforce these rules correctly, the result being that inaccurate error messages were shown when trying to use the statement to change the storage engine used by an individual partition or partitions. (Bug #31931)
- **Partitioning:** Using the `DATA DIRECTORY` and `INDEX DIRECTORY` options for partitions with `CREATE TABLE` or `ALTER TABLE` statements appeared to work on Windows, although they are not supported by MySQL on Windows systems, and subsequent attempts to use the tables referenced caused errors. Now these options are disabled on Windows, and attempting to use them generates a warning. (Bug #30459)
- **Replication:** The failure of a `CREATE TABLE ... ENGINE=InnoDB ... SELECT` statement caused the slave to lose data. (Bug #35762)
- **Replication:** When using row-based replication, a slave could crash at startup because it received a row-based replication event that `InnoDB` could not handle due to an incorrect test of the query string provided by MySQL, which was `NULL` for row-based replication events. (Bug #35226)
- **Replication:** `insert_id` was not written to the binary log for inserts into `BLACKHOLE` tables. (Bug #35178)
- **Replication:** When using statement-based replication and a `DELETE`, `UPDATE`, or `INSERT ... SELECT` statement using a `LIMIT` clause is encountered, a warning that the statement is not safe to replicate in statement mode is now issued; when using `MIXED` mode, the statement is now replicated using the row-based format. (Bug #34768)
- **Replication:** `mysqlbinlog` did not output the values of `auto_increment_increment` and `auto_increment_offset` when both were equal to their default values (for both of these variables, the default is 1). This meant that a binary log recorded by a client using the defaults for both variables and then replayed on another client using its own values for either or both of these variables produced erroneous results. (Bug #34732)

References: See also: Bug #31168.

- **Replication:** When the Windows version of `mysqlbinlog` read 4.1 binary logs containing `LOAD DATA INFILE` statements, it output backslashes as path separators, causing problems for client programs expecting forward slashes. In such cases, it now converts `\\` to `/` in directory paths. (Bug #34355)
- **Replication:** `SHOW SLAVE STATUS` failed when slave I/O was about to terminate. (Bug #34305)
- **Replication:** The character sets and collations used for constant identifiers in stored procedures were not replicated correctly. (Bug #34289)
- **Replication:** `mysqlbinlog` from a 5.1 or later MySQL distribution could not read binary logs generated by a 4.1 server when the logs contained `LOAD DATA INFILE` statements. (Bug #34141)

References: This issue is a regression of: Bug #32407.

- **Replication:** A `CREATE USER`, `DROP USER`, or `RENAME USER` statement that fails on the master, or that is a duplicate of any of these statements, is no longer written to the binary log; previously, either of these occurrences could cause the slave to fail. (Bug #33862)

References: See also: Bug #29749.

- **Replication:** `SHOW BINLOG EVENTS` could fail when the binary log contained one or more events whose size was close to the value of `max_allowed_packet`. (Bug #33413)
- **Replication:** An extraneous `ROLLBACK` statement was written to the binary log by a connection that did not use any transactional tables. (Bug #33329)
- **Replication:** `mysqlbinlog` failed to release all of its memory after terminating abnormally. (Bug #33247)
- **Replication:** When a stored routine or trigger, running on a master that used MySQL 5.0 or MySQL 5.1.11 or earlier, performed an insert on an `AUTO_INCREMENT` column, the `insert_id` value was not replicated correctly to a slave running MySQL 5.1.12 or later (including any MySQL 6.0 release). (Bug #33029)

References: See also: Bug #19630.

- **Replication:** The error message generated due to lack of a default value for an extra column was not sufficiently informative. (Bug #32971)
- **Replication:** When a user variable was used inside an `INSERT` statement, the corresponding binary log event was not written to the binary log correctly. (Bug #32580)
- **Replication:** When using row-based replication, deletes from a table with a foreign key constraint failed on the slave. (Bug #32468)
- **Replication:** The `--base64-output` option for `mysqlbinlog` was not honored for all types of events. This interfered in some cases with performing point-in-time recovery. (Bug #32407)

References: See also: Bug #46640, Bug #34777.

- **Replication:** SQL statements containing comments using `--` syntax were not replayable by `mysqlbinlog`, even though such statements replicated correctly. (Bug #32205)
- **Replication:** When using row-based replication from a master running MySQL 5.1.21 or earlier to a slave running 5.1.22 or later, updates of integer columns failed on the slave with `Error in Unknown event: row application failed`. (Bug #31583)

References: This issue is a regression of: Bug #21842.

- **Replication:** Replicating write, update, or delete events from a master running MySQL 5.1.15 or earlier to a slave running 5.1.16 or later caused the slave to crash. (Bug #31581)
- **Replication:** When using row-based replication, the slave stopped when attempting to delete nonexistent rows from a slave table without a primary key. In addition, no error was reported when this occurred. (Bug #31552)
- **Replication:** Errors due to server ID conflicts were reported only in the slave's error log; now these errors are also shown in the `Server_IO_State` column in the output of `SHOW SLAVE STATUS`. (Bug #31316)
- **Replication:** `STOP SLAVE` did not stop connection attempts properly. If the I/O slave thread was attempting to connect, `STOP SLAVE` waited for the attempt to finish, sometimes for a long period of time, rather than stopping the slave immediately. (Bug #31024)

References: See also: Bug #30932.

- **Replication:** Issuing a `DROP VIEW` statement caused replication to fail if the view did not actually exist. (Bug #30998)
- **Replication:** Replication of `LOAD DATA INFILE` could fail when `read_buffer_size` was larger than `max_allowed_packet`. (Bug #30435)
- **Replication:** Replication crashed with the NDB storage engine when `mysqld` was started with `--character-set-server=ucs2`. (Bug #29562)
- **Replication:** When using row-based logging, nontransactional updates were not written atomically to the binary log. If a nontransactional update was made concurrently with some other update, this could cause incorrect binary logging, and consequently the slave could diverge from the master. Now, nontransactional updates are always written atomically to the binary log. (Bug #29020)
- **Replication:** Setting `server_id` did not update its value for the current session. (Bug #28908)
- **Replication:** Some older servers wrote events to the binary log using different numbering from what is currently used, even though the file format number in the file is the same. Slaves running MySQL 5.1.18 and later could not read these binary logs properly. Binary logs from these older versions now are recognized and event numbers are mapped to the current numbering so that they can be interpreted properly. (Bug #27779, Bug #32434)

References: This issue is a regression of: Bug #22583.

- **Replication:** `MASTER_POS_WAIT()` did not return `NULL` when the server was not a slave. (Bug #26622)
- **Replication:** The nonspecific error message `Wrong parameters to function register_slave` resulted when `START SLAVE` failed to register on the master due to excess length of any the slave server options `--report-host`, `--report-user`, or `--report-password`. An error message specific to each of these options is now returned in such cases. The new error messages are:

- `Failed to register slave: too long 'report-host'`
 - `Failed to register slave: too long 'report-user'`
 - `Failed to register slave; too long 'report-password'`
- (Bug #22989)

References: See also: Bug #19328.

- **Replication:** `PURGE BINARY LOGS TO` and `PURGE BINARY LOGS BEFORE` did not handle missing binary log files correctly or in the same way. Now for both of these statements, if any files listed in

the `.index` file are missing from the file system, the statement fails with an error. (Bug #18199, Bug #18453)

- **Replication:** `START SLAVE UNTIL MASTER_LOG_POS=position` issued on a slave that was using `--log-slave-updates` and that was involved in circular replication would cause the slave to run and stop one event later than that specified by the value of `position`. (Bug #13861)
- Manually replacing a binary log file with a directory having the same name caused an error that was not handled correctly. (Bug #35675)
- Using `LOAD DATA INFILE` with a view could crash the server. (Bug #35469)
- Selecting from `INFORMATION_SCHEMA.REFERENTIAL_CONSTRAINTS` could cause a server crash. (Bug #35406)

References: See also: Bug #35108.

- For a `TEMPORARY` table, `DELETE` with no `WHERE` clause could fail when preceded by `DELETE` statements with a `WHERE` clause. (Bug #35392)
- If the server crashed with an `InnoDB` error due to unavailability of undo slots, errors could persist during rollback when the server was restarted: There are two `UNDO` slot caches (for `INSERT` and `UPDATE`). If all slots end up in one of the slot caches, a request for a slot from the other slot cache failed. This can happen if the request is for an `UPDATE` slot and all slots are in the `INSERT` slot cache, or vice versa. (Bug #35352)
- In some cases, when too many clients tried to connect to the server, the proper `SQLSTATE` code was not returned. (Bug #35289)
- Memory-allocation failures for attempts to set `key_buffer_size` to large values could result in a server crash. (Bug #35272)
- For `InnoDB` tables, `ALTER TABLE DROP` failed if the name of the column to be dropped began with "foreign". (Bug #35220)
- Queries could return different results depending on whether `ORDER BY` columns were indexed. (Bug #35206)
- When a view containing a reference to `DUAL` was created, the reference was removed when the definition was stored, causing some queries against the view to fail with invalid SQL syntax errors. (Bug #35193)
- `SELECT ... FROM INFORMATION_SCHEMA.REFERENTIAL_CONSTRAINTS` caused the server to crash if the table referenced by a foreign key had been dropped. This issue was observed on Windows platforms only. (Bug #35108)

References: See also: Bug #35406.

- Debugging symbols were missing for some executables in Windows binary distributions. (Bug #35104)
- Nonconnection threads were being counted in the value of the `Max_used_connections` status variable. (Bug #35074)
- A query that performed a `ref_or_null` join where the second table used a key having one or columns that could be `NULL` and had a column value that was `NULL` caused the server to crash. (Bug #34945)

References: This issue is a regression of: Bug #12144.

- For some queries, the optimizer used an ordered index scan for `GROUP BY` or `DISTINCT` when it was supposed to use a loose index scan, leading to incorrect results. (Bug #34928)
- Creating a foreign key on an `InnoDB` table that was created with an explicit `AUTO_INCREMENT` value caused that value to be reset to 1. (Bug #34920)

References: This issue is a regression of: Bug #23313.

- `mysqldump` failed to return an error code when using the `--master-data` option without binary logging being enabled on the server. (Bug #34909)
- Under some circumstances, the value of `mysql_insert_id()` following a `SELECT ... INSERT` statement could return an incorrect value. This could happen when the last `SELECT ... INSERT` did not involve an `AUTO_INCREMENT` column, but the value of `mysql_insert_id()` was changed by some previous statements. (Bug #34889)
- Table and database names were mixed up in some places of the subquery transformation procedure. This could affect debugging trace output and further extensions of that procedure. (Bug #34830)
- If `fsync()` returned `ENOLCK`, `InnoDB` could treat this as fatal and cause abnormal server termination. `InnoDB` now retries the operation. (Bug #34823)
- `CREATE SERVER` and `ALTER SERVER` could crash the server if out-of-memory conditions occurred. (Bug #34790)
- `DROP SERVER` does not release memory cached for server structures created by `CREATE SERVER`, so repeated iterations of these statements resulted in a memory leak. `FLUSH PRIVILEGES` now releases the memory allocated for `CREATE SERVER`. (Bug #34789)
- A malformed URL used for a `FEDERATED` table's `CONNECTION` option value in a `CREATE TABLE` statement was not handled correctly and could crash the server. (Bug #34788)
- Queries such as `SELECT ROW(1, 2) IN (SELECT t1.a, 2) FROM t1 GROUP BY t1.a` (combining row constructors and subqueries in the `FROM` clause) could lead to assertion failure or unexpected error messages. (Bug #34763)
- Using `NAME_CONST()` with a negative number and an aggregate function caused MySQL to crash. This could also have a negative impact on replication. (Bug #34749)
- A memory-handling error associated with use of `GROUP_CONCAT()` in subqueries could result in a server crash. (Bug #34747)
- For an indexed integer column `col_name` and a value `N` that is one greater than the maximum value permitted for the data type of `col_name`, conditions of the form `WHERE col_name < N` failed to return rows where the value of `col_name` is `N - 1`. (Bug #34731)
- A server running with the `--debug` option could attempt to dereference a null pointer when opening tables, resulting in a crash. (Bug #34726)
- Assigning an “incremental” value to the `debug` system variable did not add the new value to the current value. For example, if the current `debug` value was `'T'`, the statement `SET debug = '+P'` resulted in a value of `'P'` rather than the correct value of `'P:T'`. (Bug #34678)
- For debug builds, reading from `INFORMATION_SCHEMA.TABLES` or `INFORMATION_SCHEMA.COLUMNS` could cause assertion failures. This could happen under rare circumstances when `INFORMATION_SCHEMA` fails to get information about a table (for example, when a connection is killed). (Bug #34656)

- Executing a `TRUNCATE TABLE` statement on a table having both a foreign key reference and a `DELETE` trigger crashed the server. (Bug #34643)
- Some subqueries using an expression that included an aggregate function could fail or in some cases lead to a crash of the server. (Bug #34620)
- Dangerous pointer arithmetic crashed the server on some systems. (Bug #34598)
- Creating a view inside a stored procedure could lead to a crash of the MySQL Server. (Bug #34587)
- A server crash could occur if `INFORMATION_SCHEMA` tables built in memory were swapped out to disk during query execution. (Bug #34529)
- `CAST(AVG(arg) AS DECIMAL)` produced incorrect results for non-`DECIMAL` arguments. (Bug #34512)
- The per-thread debugging settings stack was not being deallocated before thread termination, resulting in a stack memory leak. (Bug #34424)
- Executing an `ALTER VIEW` statement on a table crashed the server. (Bug #34337)
- `InnoDB` could crash if overflow occurred for an `AUTO_INCREMENT` column. (Bug #34335)
- For `InnoDB`, exporting and importing a table could corrupt `TINYBLOB` columns, and a subsequent `ALTER TABLE` could corrupt `TINYTEXT` columns as well. (Bug #34300)
- `DEFAULT 0` was not permitted for the `YEAR` data type. (Bug #34274)
- Under some conditions, a `SET GLOBAL innodb_commit_concurrency` or `SET GLOBAL innodb_autoextend_increment` statement could fail. (Bug #34223)

References: This issue is a regression of: Bug #31177.

- `mysqldump` attempts to set the `character_set_results` system variable after connecting to the server. This failed for pre-4.1 servers that have no such variable, but `mysqldump` did not account for this and 1) failed to dump database contents; 2) failed to produce any error message alerting the user to the problem. (Bug #34192)
- Use of stored functions in the `WHERE` clause for `SHOW OPEN TABLES` caused a server crash. (Bug #34166)
- For a `FEDERATED` table with an index on a nullable column, accessing the table could crash a server, return an incorrect result set, or return `ERROR 1030 (HY000): Got error 1430 from storage engine`. (Bug #33946)
- Passing anything other than an integer argument to a `LIMIT` clause in a prepared statement would fail. (This limitation was introduced to avoid replication problems; for example, replicating the statement with a string argument would cause a parse failure in the slave). Now, arguments to the `LIMIT` clause are converted to integer values, and these converted values are used when logging the statement. (Bug #33851)
- An internal buffer in `mysql` was too short. Overextending it could cause stack problems or segmentation violations on some architectures. (This is not a problem that could be exploited to run arbitrary code.) (Bug #33841)
- A query using `WHERE (column1='string1' AND column2=constant1) OR (column1='string2' AND column2=constant2)`, where `coll` used a binary collation and `string1` matched `string2` except for case, failed to match any records even when

matches were found by a query using the equivalent clause `WHERE column2=constant1 OR column2=constant2`. (Bug #33833)

- Large unsigned integers were improperly handled for prepared statements, resulting in truncation or conversion to negative numbers. (Bug #33798)
- Reuse of prepared statements could cause a memory leak in the embedded server. (Bug #33796)
- The server crashed when executing a query that had a subquery containing an equality `X=Y` where `Y` referred to a named select list expression from the parent select. The server crashed when trying to use the `X=Y` equality for `ref`-based access. (Bug #33794)
- Some queries using a combination of `IN`, `CONCAT()`, and an implicit type conversion could return an incorrect result. (Bug #33764)
- In some cases a query that produced a result set when using `ORDER BY ASC` did not return any results when this was changed to `ORDER BY DESC`. (Bug #33758)
- Disabling concurrent inserts caused some cacheable queries not to be saved in the query cache. (Bug #33756)
- `ORDER BY ... DESC` sorts could produce misordered results. (Bug #33697)
- The server could crash when `REPEAT` or another control instruction was used in conjunction with labels and a `LEAVE` instruction. (Bug #33618)
- The parser permitted control structures in compound statements to have mismatched beginning and ending labels. (Bug #33618)
- `make_binary_distribution` passed the `--print-libgcc-file` option to the C compiler, but this does not work with the `ICC` compiler. (Bug #33536)
- Threads created by the event scheduler were incorrectly counted against the `max_connections` thread limit, which could lead to client lockout. (Bug #33507)
- Dropping a function after dropping the function's creator could cause the server to crash. (Bug #33464)
- Certain combinations of views, subselects with outer references and stored routines or triggers could cause the server to crash. (Bug #33389)
- `SET GLOBAL myisam_max_sort_file_size=DEFAULT` set `myisam_max_sort_file_size` to an incorrect value. (Bug #33382)

References: See also: Bug #31177.

- `ENUM`- or `SET`-valued plugin variables could not be set from the command line. (Bug #33358)
- Loading plugins using command-line options to `mysqld` could cause an assertion failure. (Bug #33345)
- `SLEEP(0)` failed to return on 64-bit Mac OS X due to a bug in `pthread_cond_timedwait()`. (Bug #33304)
- Using **Control+R** in the `mysql` client caused it to crash. (Bug #33288)
- For `MyISAM` tables, `CHECK TABLE` (non-`QUICK`) and any form of `REPAIR TABLE` incorreced treated rows as corrupted under the combination of the following conditions:
 - The table had dynamic row format

- The table had a `CHAR` (not `VARCHAR`) column longer than 127 bytes (for multibyte character sets this could be less than 127 characters)
- The table had rows with a significant length of more than 127 bytes significant length in that `CHAR` column (that is, a byte beyond byte position 127 must be a nonspace character)

This problem affected `CHECK TABLE`, `REPAIR TABLE`, `OPTIMIZE TABLE`, `ALTER TABLE`. `CHECK TABLE` reported and marked the table as crashed if any row was present that fulfilled the third condition. The other statements deleted these rows. (Bug #33222)

- Granting the `UPDATE` privilege on one column of a view caused the server to crash. (Bug #33201)
- For `DECIMAL` columns used with the `ROUND(X,D)` or `TRUNCATE(X,D)` function with a nonconstant value of `D`, adding an `ORDER BY` for the function result produced misordered output. (Bug #33143)

References: See also: Bug #33402, Bug #30617.

- The `CSV` engine did not honor update requests for `BLOB` columns when the new column value had the same length as the value to be updated. (Bug #33067)
- After receiving a `SIGHUP` signal, the server could crash, and user-specified log options were ignored when reopening the logs. (Bug #33065)
- When MySQL was built with OpenSSL, the SSL library was not properly initialized with information of which endpoint it was (server or client), causing connection failures. (Bug #33050)
- Under some circumstances a combination of aggregate functions and `GROUP BY` in a `SELECT` query over a view could lead to incorrect calculation of the result type of the aggregate function. This in turn could lead to incorrect results, or to crashes on debug builds of the server. (Bug #33049)
- For `DISTINCT` queries, MySQL 4.0 and 4.1 stopped reading joined tables as soon as the first matching row was found. However, this optimization was lost in MySQL 5.0, which instead read all matching rows. This fix for this regression may result in a major improvement in performance for `DISTINCT` queries in cases where many rows match. (Bug #32942)
- Repeated creation and deletion of views within prepared statements could eventually crash the server. (Bug #32890)

References: See also: Bug #34587.

- Incorrect assertions could cause a server crash for `DELETE` triggers for transactional tables. (Bug #32790)
- In some cases where setting a system variable failed, no error was sent to the client, causing the client to hang. (Bug #32757)
- Enabling the `PAD_CHAR_TO_FULL_LENGTH` SQL mode caused privilege-loading operations (such as `FLUSH PRIVILEGES`) to include trailing spaces from grant table values stored in `CHAR` columns. Authentication for incoming connections failed as a result. Now privilege loading does not include trailing spaces, regardless of SQL mode. (Bug #32753)
- The `SHOW ENGINE INNODB STATUS` and `SHOW ENGINE INNODB MUTEX` statements incorrectly required the `SUPER` privilege rather than the `PROCESS` privilege. (Bug #32710)
- Inserting strings with a common prefix into a table that used the `ucs2` character set corrupted the table. (Bug #32705)

- Tables in the `mysql` database that stored the current `sql_mode` value as part of stored program definitions were not updated with newer mode values (`NO_ENGINE_SUBSTITUTION`, `PAD_CHAR_TO_FULL_LENGTH`). This causes various problems defining stored programs if those modes were included in the current `sql_mode` value. (Bug #32633)
- A view created with a string literal for one of the columns picked up the connection character set, but not the collation. Comparison to that field therefore used the default collation for that character set, causing an error if the connection collation was not compatible with the default collation. The problem was caused by text literals in a view being dumped with a character set introducer even when this was not necessary, sometimes leading to a loss of collation information. Now the character set introducer is dumped only if it was included in the original query. (Bug #32538)

References: See also: Bug #21505.

- Queries using `LIKE` on tables having indexed `CHAR` columns using either of the `uicjpm` or `ujis` character sets did not return correct results. (Bug #32510)
- Executing a prepared statement associated with a materialized cursor sent to the client a metadata packet with incorrect table and database names. The problem occurred because the server sent the name of the temporary table used by the cursor instead of the table name of the original table.

The same problem occurred when selecting from a view, in which case the name of the table name was sent, rather than the name of the view. (Bug #32265)
- On Windows, `mysqltest_embedded.exe` did not properly execute the `send` command. (Bug #32044)
- A variable named `read_only` could be declared even though that is a reserved word. (Bug #31947)
- On Windows, the build process failed with four parallel build threads. (Bug #31929)
- Queries testing numeric constants containing leading zeros against `ZEROFILL` columns were not evaluated correctly. (Bug #31887)
- If an error occurred during file creation, the server sometimes did not remove the file, resulting in an unused file in the file system. (Bug #31781)
- The `mysqld` crash handler failed on Windows. (Bug #31745)
- When upgrading from MySQL 5.1.19 to any version between MySQL 5.1.20 to MySQL 5.1.23, the MySQL Instance Configuration Wizard failed to account for the change in name for the `mysqld-nt.exe` to `mysqld.exe`, causing MySQL to fail to start properly after the upgrade. (Bug #31674)
- The server returned the error message `Out of memory; restart server and try again` when the actual problem was that the sort buffer was too small. Now an appropriate error message is returned in such cases. (Bug #31590)
- A table having an index that included a `BLOB` or `TEXT` column, and that was originally created with a MySQL server using version 4.1 or earlier, could not be opened by a 5.1 or later server. (Bug #31331)
- The `mysql_change_user()` C API function caused global `Com_xxx` status variable values to be incorrect. (Bug #31222)
- When sorting privilege table rows, the server treated escaped wildcard characters (`\%` and `_`) the same as unescaped wildcard characters (`%` and `_`), resulting in incorrect row ordering. (Bug #31194)
- On Windows, `SHOW PROCESSLIST` could display process entries with a `State` value of `*** DEAD ***`. (Bug #30960)

- `ROUND(X,D)` or `TRUNCATE(X,D)` for nonconstant values of `D` could crash the server if these functions were used in an `ORDER BY` that was resolved using `filesort`. (Bug #30889)
- Resetting the query cache by issuing a `SET GLOBAL query_cache_size=0` statement caused the server to crash if it concurrently was saving a new result set to the query cache. (Bug #30887)
- Manifest problems prevented `MySQLInstanceConfig.exe` from running on Windows Vista. (Bug #30823)
- If an alias was used to refer to the value returned by a stored function within a subselect, the outer select recognized the alias but failed to retrieve the value assigned to it in the subselect. (Bug #30787)

References: This issue is a regression of: Bug #20777.

- Binary logging for a stored procedure differed depending on whether or not execution occurred in a prepared statement. (Bug #30604)
- An orphaned PID file from a no-longer-running process could cause `mysql.server` to wait for that process to exit even though it does not exist. (Bug #30378)
- The `Table_locks_waited` waited variable was not incremented in the cases that a lock had to be waited for but the waiting thread was killed or the request was aborted. (Bug #30331)
- The `Com_create_function` status variable was not incremented properly. (Bug #30252)
- View metadata returned from `INFORMATION_SCHEMA.VIEWS` was changed by the fix for Bug #11986, causing the information returned in MySQL 5.1 to differ from that returned in 5.0. (Bug #30217)

References: See also: Bug #11986.

- `mysqld` displayed the `--enable-pstack` option in its help message even if MySQL was configured without `--with-pstack`. (Bug #29836)
- The `mysql_config` command would output `CFLAGS` values that were incompatible with C++ for the HP-UX platform. (Bug #29645)
- Views were treated as insertable even if some base table columns with no default value were omitted from the view definition. (This is contrary to the condition for insertability that a view must contain all columns in the base table that do not have a default value.) (Bug #29477)
- `myisamchk` always reported the character set for a table as `latin1_swedish_ci (8)` regardless of the table's actual character set. (Bug #29182)
- `InnoDB` could return an incorrect rows-updated value for `UPDATE` statements. (Bug #29157)
- The MySQL preferences pane did not work to start or stop MySQL on Mac OS X 10.5 (Leopard). (Bug #28854)
- For upgrading to a new major version using RPM packages (such as 4.1 to 5.0), if the installation procedure found an existing MySQL server running, it could fail to shut down the old server, but also erroneously removed the server's socket file. Now the procedure checks for an existing server package from a different vendor or major MySQL version. In such case, it refuses to install the server and recommends how to safely remove the old packages before installing the new ones. (Bug #28555)
- `mysqlhotcopy` silently skipped databases with names consisting of two alphanumeric characters. (Bug #28460)
- No information was written to the general query log for the `COM_STMT_CLOSE`, `COM_STMT_RESET`, and `COM_STMT_SEND_LONG_DATA` commands. (These occur when a client invokes the

`mysql_stmt_close()`, `mysql_stmt_reset()` and `mysql_stmt_send_long_data()` C API functions.) (Bug #28386)

- The `FEDERATED` storage engine did not perform identifier quoting for column names that are reserved words when sending statements to the remote server. (Bug #28269)
- The SQL parser did not accept an empty `UNION=()` clause. This meant that, when there were no underlying tables specified for a `MERGE` table, `SHOW CREATE TABLE` and `mysqldump` both output statements that could not be executed.

Now it is possible to execute a `CREATE TABLE` or `ALTER TABLE` statement with an empty `UNION=()` clause. However, `SHOW CREATE TABLE` and `mysqldump` do not output the `UNION=()` clause if there are no underlying tables specified for a `MERGE` table. This also means it is now possible to remove the underlying tables for a `MERGE` table using `ALTER TABLE ... UNION=()`. (Bug #28248)

- It was possible to exhaust memory by repeatedly running `index_merge` queries and never performing any `FLUSH TABLES` statements. (Bug #27732)
- When `utf8` was set as the connection character set, using `SPACE()` with a non-Unicode column produced an error. (Bug #27580)

References: See also: Bug #23637.

- The parser rules for the `SHOW PROFILE` statement were revised to work with older versions of `bison`. (Bug #27433)
- `resolveip` failed to produce correct results for host names that begin with a digit. (Bug #27427)
- In `ORDER BY` clauses, mixing aggregate functions and nongrouping columns is not permitted if the `ONLY_FULL_GROUP_BY` SQL mode is enabled. However, in some cases, no error was thrown because of insufficient checking. (Bug #27219)
- For the `--record_log_pos` option, `mysqlhotcopy` now determines the slave status information from the result of `SHOW SLAVE STATUS` by using the `Relay_Master_Log_File` and `Exec_Master_Log_Pos` values rather than the `Master_Log_File` and `Read_Master_Log_Pos` values. This provides a more accurate indication of slave execution relative to the master. (Bug #27101)
- The MySQL Instance Configuration Wizard would not permit you to choose a service name, even though the criteria for the service name were valid. The code that checks the name has been updated to support the correct criteria of any string less than 256 character and not containing either a forward or backward slash character. (Bug #27013)
- Memory corruption, a crash of the MySQL server, or both, could take place if a low-level I/O error occurred while an `ARCHIVE` table was being opened. (Bug #26978)
- `DROP DATABASE` failed for attempts to drop databases with names that contained the legacy `#mysql50#` name prefix. (Bug #26703)
- `config-win.h` unconditionally defined `bool` as `BOOL`, causing problems on systems where `bool` is 1 byte and `BOOL` is 4 bytes. (Bug #26461)
- On Windows, for distributions built with debugging support, `mysql` could crash if the user typed **Control +C**. (Bug #26243)
- The XPath `boolean()` function did not cast string and nodeset values correctly in some cases. It now returns `TRUE` for any nonempty string or nodeset and 0 for a `NULL` string, as specified in the XPath standard.. (Bug #26051)

- When symbolic links were disabled, either with a server startup option or by enabling the `NO_DIR_IN_CREATE` SQL mode, `CREATE TABLE` silently ignored the `DATA DIRECTORY` and `INDEX DIRECTORY` table options. Now the server issues a warning if symbolic links are disabled when these table options are used. (Bug #25677)
- Attempting to create an index with a prefix on a `DECIMAL` column appeared to succeed with an inaccurate warning message. Now, this action fails with the error `Incorrect prefix 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 prefix keys`. (Bug #25426)
- `mysqlcheck -A -r` did not correctly identify all tables that needed repairing. (Bug #25347)
- On Windows, an error in `configure.js` caused installation of source distributions to fail. (Bug #25340)
- The `Qcache_free_blocks` status variable did not display a value of 0 if the query cache was disabled. (Bug #25132)
- The client library had no way to return an error if no connection had been established. This caused problems such as `mysql_library_init()` failing silently if no `errmsg.sys` file was available. (Bug #25097)
- On Mac OS X, the StartupItem for MySQL did not work. (Bug #25008)
- For Windows 64-bit builds, enabling shared-memory support caused client connections to fail. (Bug #24992)
- `mysql` did not use its completion table. Also, the table contained few entries. (Bug #24624)
- If a user installed MySQL Server and set a password for the `root` user, and then uninstalled and reinstalled MySQL Server to the same location, the user could not use the MySQL Instance Config wizard to configure the server because the uninstall operation left the previous data directory intact. The config wizard *assumed* that any new install (not an upgrade) would have the default data directory where the `root` user has no password. The installer now writes a registry key named `FoundExistingDataDir`. If the installer finds an existing data directory, the key will have a value of 1, otherwise it will have a value of 0. When `MySQLInstanceConfig.exe` is run, it will attempt to read the key. If it can read the key, and the value is 1 and there is no existing instance of the server (indicating a new installation), the Config Wizard will permit the user to input the old password so the server can be configured. (Bug #24215)
- Logging of statements to log tables was incorrect for statements that contained `utf8`-incompatible binary strings. Incompatible sequences are hex-encoded now. (Bug #23924)
- The MySQL header files contained some duplicate macro definitions that could cause compilation problems. (Bug #23839)
- `SHOW COLUMNS` on a `TEMPORARY` table caused locking issues. (Bug #23588)
- For distributions compiled with the bundled `libedit` library, there were difficulties using the `mysql` client to enter input for non-ASCII or multibyte characters. (Bug #23097)
- `percona` reported incomplete or inaccurate information. (Bug #23028, Bug #25177)
- After stopping and starting the event scheduler, disabled events could remain in the execution queue. (Bug #22738)
- The server produced a confusing error message when attempting to open a table that required a storage engine that was not loaded. (Bug #22708)

- For views or stored programs created with an invalid `DEFINER` value, the error message was confusing (did not tie the problem to the `DEFINER` clause) and has been improved. (Bug #21854)
 - Warnings for deprecated syntax constructs used in stored routines make sense to report only when the routine is being created, but they were also being reported when the routine was parsed for loading into the execution cache. Now they are reported only at routine creation time. (Bug #21801)
 - On Mac OS X, `mysqld` did not react to Control+C when run under `gdb`, even when run with the `--gdb` option. (Bug #21567)
 - `CREATE ... SELECT` did not always set `DEFAULT` column values in the new table. (Bug #21380)
 - `mysql_config` output did not include `-lmygcc` on some platforms when it was needed. (Bug #21158)
 - `mysql-stress-test.pl` and `mysqld_multi.server.sh` were missing from some binary distributions. (Bug #21023, Bug #25486)
 - The `BENCHMARK()` function, invoked with more than 2147483648 iterations (the size of a signed 32-bit integer), terminated prematurely. (Bug #20752)
 - `mysqldumpslow` returned a confusing error message when no configuration file was found. (Bug #20455)
 - `MySQLInstanceConfig.exe` could lose the `innodb_data_home_dir` setting when reconfiguring an instance. (Bug #19797)
 - `DROP DATABASE` did not drop orphaned `FOREIGN KEY` constraints. (Bug #18942)
 - `CREATE TABLE` permitted 0 as the default value for a `TIMESTAMP` column when the server was running in `NO_ZERO_DATE` mode. (Bug #18834)
 - A `SET` column whose definition specified 64 elements could not be updated using integer values. (Bug #15409)
 - If a `SELECT` calls a stored function in a transaction, and a statement within the function fails, that statement should roll back. Furthermore, if `ROLLBACK` is executed after that, the entire transaction should be rolled back. Before this fix, the failed statement did not roll back when it failed (even though it might ultimately get rolled back by a `ROLLBACK` later that rolls back the entire transaction). (Bug #12713)
- References: See also: Bug #34655.
- The parser incorrectly permitted `SQLSTATE '00000'` to be specified for a condition handler. (This is incorrect because the condition must be a failure condition and `'00000'` indicates success.) (Bug #8759)
 - `MySQLInstanceConfig.exe` did not save the `innodb_data_home_dir` value to the `my.ini` file under certain circumstances. (Bug #6627)

Changes in MySQL 5.1.23 (2008-01-29)

- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

Functionality Added or Changed

- **Security Fix; Important Change; Partitioning:** It was possible, by creating a partitioned table using the `DATA DIRECTORY` and `INDEX DIRECTORY` options to gain privileges on other tables having the

same name as the partitioned table. As a result of this fix, any table-level [DATA DIRECTORY](#) or [INDEX DIRECTORY](#) options are now ignored for partitioned tables. (Bug #32091, CVE-2007-5970)

References: See also: Bug #29325, Bug #32111.

- **Incompatible Change:** In MySQL 5.1.6, when log tables were implemented, the default log destination for the general query and slow query log was [TABLE](#). This default has been changed to [FILE](#), which is compatible with MySQL 5.0, but incompatible with earlier releases of MySQL 5.1 from 5.1.6 to 5.1.20. If you are upgrading from MySQL 5.0 to 5.1.21 or higher, no logging option changes should be necessary. However, if you are upgrading from 5.1.6 through 5.1.20 to 5.1.21 or higher and were using [TABLE](#) logging, use the `--log-output=TABLE` option explicitly to preserve your server's table-logging behavior.

The MySQL 5.1.23 fix is in addition to a fix in 5.1.21 because it turned out that the default was set in two places, only one of which was fixed the first time. (Bug #29993)

- **Incompatible Change**

The parser accepted statements that contained `/* ... */` that were not properly closed with `*/`, such as `SELECT 1 /* + 2`. Statements that contain unclosed `/*`-comments now are rejected with a syntax error.

This fix has the potential to cause incompatibilities. Because of Bug #26302, which caused the trailing `*/` to be truncated from comments in views, stored routines, triggers, and events, it is possible that objects of those types may have been stored with definitions that now will be rejected as syntactically invalid. Such objects should be dropped and re-created so that their definitions do not contain truncated comments. (Bug #28779)

- **MySQL Cluster:** The following improvements have been made in the `ndb_size.pl` utility:
 - The script can now be used with multiple databases; lists of databases and tables can also be excluded from analysis.
 - Schema name information has been added to index table calculations.
 - The database name is now an optional parameter, the exclusion of which causes all databases to be examined.
 - If selecting from [INFORMATION_SCHEMA](#) fails, the script now attempts to fall back to `SHOW TABLES`.
 - A `--real_table_name` option has been added; this designates a table to handle unique index size calculations.
 - The report title has been amended to cover cases where more than one database is being analyzed.

Support for a `--socket` option was also added.

For more information, see [ndb_size.pl — NDBCLUSTER Size Requirement Estimator](#). (Bug #28683, Bug #28253)

- **MySQL Cluster:** Mapping of [NDB](#) error codes to MySQL storage engine error codes has been improved. (Bug #28423)
- **MySQL Cluster:** The output of the `ndb_mgm` client `SHOW` and `STATUS` commands now indicates when the cluster is in single user mode. (Bug #27999)
- **MySQL Cluster:** The output from the cluster management client showing the progress of data node starts has been improved. (Bug #23354)

- **Partitioning:** Error messages for partitioning syntax errors have been made more descriptive. (Bug #29368)
 - **Replication:** Replication of the following SQL functions now switches to row-based logging in `MIXED` mode, and generates a warning in `STATEMENT` mode:
 - `USER()`
 - `CURRENT_USER()` and its alias `CURRENT_USER`
 - `FOUND_ROWS()`
 - `ROW_COUNT()`
- See [Mixed Binary Logging Format](#), for more information. (Bug #12092, Bug #28086, Bug #30244)
- `mysqldump` information at the top of the output now shows the same information as `mysqldump` invoked with the `-v` option, namely the `mysqldump` version number, the MySQL server version, and the distribution. (Bug #32350)
 - `mysqltest` now has a `change_user` command to change the user for the current connection. (It invokes the `mysql_change_user()` C API function.) (Bug #31608)
 - `mysql-test-run.pl` now permits a suite name prefix to be specified in command-line arguments that name test cases. The test name syntax now is `[suite_name.]test_name[.suffix]`. For example, `mysql-test-run.pl binlog.mytest` runs the `mytest.test` test in the `binlog` test suite. (Bug #31400)
 - The `--event-scheduler` option without a value disabled the event scheduler. Now it enables the event scheduler. (Bug #31332)
 - `mysqldump` produces a `-- Dump completed on DATE` comment at the end of the dump if `--comments` is given. The date causes dump files for identical data take at different times to appear to be different. The new options `--dump-date` and `--skip-dump-date` control whether the date is added to the comment. `--skip-dump-date` suppresses date printing. The default is `--dump-date` (include the date in the comment). (Bug #31077)
 - Server parser performance was improved for expression parsing by lowering the number of state transitions and reductions needed. (Bug #30625)
 - Server parser performance was improved for identifier lists, expression lists, and UDF expression lists. (Bug #30333)
 - Server parser performance was improved for boolean expressions. (Bug #30237)
 - The `LAST_EXECUTED` column of the `INFORMATION_SCHEMA.EVENTS` table now indicates when the event started executing rather than when it finished executing. As a result, the `ENDS` column is never less than `LAST_EXECUTED`. (Bug #29830)
 - The `mysql_odbc_escape_string()` C API function has been removed. It has multibyte character escaping issues, doesn't honor the `NO_BACKSLASH_ESCAPES` SQL mode and is not needed anymore by Connector/ODBC as of 3.51.17. (Bug #29592)

References: See also: Bug #41728.

- If a `MyISAM` table is created with no `DATA DIRECTORY` option, the `.MYD` file is created in the database directory. By default, if `MyISAM` finds an existing `.MYD` file in this case, it overwrites it. The same applies to `.MYI` files for tables created with no `INDEX DIRECTORY` option. To suppress this behavior, start

the server with the new `--keep_files_on_create` option, in which case `MyISAM` will not overwrite existing files and returns an error instead. (Bug #29325)

- The default value of the `connect_timeout` system variable was increased from 5 to 10 seconds. This might help in cases where clients frequently encounter errors of the form `Lost connection to MySQL server at 'XXX', system error: errno`. (Bug #28359)
- MySQL now can be compiled with `gcc` 4.2.x. There was a problem involving a conflict with the `min()` and `max()` macros in `my_global.h`. (Bug #28184)
- `SHOW COLUMNS` now returns `NULL` instead of the empty string for the `Default` value of columns that have no `DEFAULT` clause specified. (Bug #27747)
- `mysql-test-run.pl` now supports a `--combination` option for specifying options to the `mysqld` server. This option is similar to `--mysqld` but should be given two or more times. `mysql-test-run.pl` executes multiple test runs, using the options for each instance of `--combination` in successive runs.

For test runs specific to a given test suite, an alternative to the use of `--combination` is to create a `combinations` file in the suite directory. The file should contain a section of options for each test run.

- The argument for the `mysql-test-run.pl --do-test` and `--skip-test` options is now interpreted as a Perl regular expression if there is a pattern metacharacter in the argument value. This enables more flexible specification of which tests to perform or skip.

Bugs Fixed

- **Security Fix; Replication:** It was possible for any connected user to issue a `BINLOG` statement, which could be used to escalate privileges.

Use of the `BINLOG` statement now requires the `SUPER` privilege. (Bug #31611, CVE-2007-6313)

- **Security Fix:** Three vulnerabilities in `yaSSL` versions 1.7.5 and earlier were discovered that could lead to a server crash or execution of unauthorized code. The exploit requires a server with `yaSSL` enabled and TCP/IP connections enabled, but does not require valid MySQL account credentials. The exploit does not apply to `OpenSSL`.



Warning

The proof-of-concept exploit is freely available on the Internet. Everyone with a vulnerable MySQL configuration is advised to upgrade *immediately*.

(Bug #33814, CVE-2008-0226, CVE-2008-0227)

- **Security Fix:** Using `RENAME TABLE` against a table with explicit `DATA DIRECTORY` and `INDEX DIRECTORY` options can be used to overwrite system table information by replacing the symbolic link points. the file to which the symlink points.

MySQL will now return an error when the file to which the symlink points already exists. (Bug #32111, CVE-2007-5969)

- **Security Fix:** `ALTER VIEW` retained the original `DEFINER` value, even when altered by another user, which could enable that user to gain the access rights of the view. Now `ALTER VIEW` is permitted only to the original definer or users with the `SUPER` privilege. (Bug #29908)
- **Security Fix:** When using a `FEDERATED` table, the local server could be forced to crash if the remote server returned a result with fewer columns than expected. (Bug #29801)

- **Security Enhancement:** It was possible to force an error message of excessive length which could lead to a buffer overflow. This has been made no longer possible as a security precaution. (Bug #32707)
- **Performance:** If a `LIMIT` clause was present, the server could fail to consider indexes that could be used for `ORDER BY` or `GROUP BY`. (Bug #28404)
- **Incompatible Change:** It is no longer possible to create `CSV` tables with `NULL` columns. However, for backward compatibility, you can continue to use such tables that were created in previous MySQL releases. (Bug #32050)
- **Incompatible Change:** With `ONLY_FULL_GROUP_BY` SQL mode enabled, queries such as `SELECT a FROM t1 HAVING COUNT(*)>2` were not being rejected as they should have been.

This fix results in the following behavior:

- There is a check against mixing group and nongroup columns *only* when `ONLY_FULL_GROUP_BY` is enabled.
- This check is done both for the select list and for the `HAVING` clause if there is one.

This behavior differs from previous versions as follows:

- Previously, the `HAVING` clause was not checked when `ONLY_FULL_GROUP_BY` was enabled; now it is checked.
- Previously, the select list was checked even when `ONLY_FULL_GROUP_BY` was not enabled; now it is checked only when `ONLY_FULL_GROUP_BY` is enabled.

(Bug #31794)

- **Incompatible Change:** Inserting a row with a `NULL` value for a `DATETIME` column results in a `CSV` file that the storage engine cannot read.

All `CSV` tables now need to be defined with each column marked as `NOT NULL`. An error is raised if you try to create a `CSV` table with columns that are not defined with `NOT NULL`. (Bug #31473, Bug #32817)

- **Incompatible Change:** `SET PASSWORD` statements now cause an implicit commit, and thus are prohibited within stored functions and triggers. (Bug #30904)
- **Incompatible Change:** The `mysql_install_db` script could fail to locate some components (including `resolveip`) during execution if the `--basedir` option was specified on the command-line or within the `my.cnf` file. This was due to a conflict when comparing the compiled-in values and the supplied values.

The `--source-install` command-line option to the script has been removed and replaced with the `--srcdir` option. `mysql_install_db` now locates components either using the compiled-in options, the `--basedir` option or `--srcdir` option. (Bug #30759)

- **Incompatible Change:** Multiple-table `DELETE` statements containing ambiguous aliases could have unintended side effects such as deleting rows from the wrong table. Examples:

```
DELETE FROM t1 AS a2 USING t1 AS a1 INNER JOIN t2 AS a2;  
DELETE t1 AS a2 FROM t1 AS a1 INNER JOIN t2 AS a2;
```

To avoid ambiguity, declaration of aliases other than in the `table_references` part of the statement should be avoided:

```
DELETE FROM t1 USING t1 AS a1 INNER JOIN t2 AS a2;
```

```
DELETE t1 FROM t1 AS a1 INNER JOIN t2 AS a2;
```

For the [USING](#) variant of multiple-table [DELETE](#) syntax, alias declarations outside the [table_references](#) part of the statement now are disallowed. (In MySQL 5.5, alias declarations outside [table_references](#) are disallowed for all multiple-table [DELETE](#) statements.) Statements containing aliases that are no longer permitted must be rewritten. (Bug #30234)

References: See also: Bug #27525.

- **Incompatible Change:** Within a stored routine, it is no longer permissible to declare a cursor for a [SHOW](#) or [DESCRIBE](#) statement. This happened to work in some instances, but is no longer supported. In many cases, a workaround for this change is to use the cursor with a [SELECT](#) query to read from an [INFORMATION_SCHEMA](#) table that produces the same information as the [SHOW](#) statement. (Bug #29223)
- **Incompatible Change:** It was possible to create a view having a column whose name consisted of an empty string or space characters only.

One result of this bug fix is that aliases for columns in the view [SELECT](#) statement are checked to ensure that they are legal column names. In particular, the length must be within the maximum column length of 64 characters, not the maximum alias length of 256 characters. This can cause problems for replication or loading dump files. For additional information and workarounds, see [Restrictions on Views](#). (Bug #27695)

References: See also: Bug #31202.

- **Incompatible Change:** Several type-preserving functions and operators returned an incorrect result type that does not match their argument types: [COALESCE\(\)](#), [IF\(\)](#), [IFNULL\(\)](#), [LEAST\(\)](#), [GREATEST\(\)](#), [CASE](#). These now aggregate using the precise SQL types of their arguments rather than the internal type. In addition, the result type of the [STR_TO_DATE\(\)](#) function is now [DATETIME](#) by default. (Bug #27216)
- **Incompatible Change:** [GRANT](#) and [REVOKE](#) statements now cause an implicit commit, and thus are prohibited within stored functions and triggers. (Bug #21975, Bug #21422, Bug #17244)
- **Incompatible Change:** It was possible for option files to be read twice at program startup, if some of the standard option file locations turned out to be the same directory. Now duplicates are removed from the list of files to be read.

Also, users could not override system-wide settings using `~/my.cnf` because `SYSCONFDIR/my.cnf` was read last. The latter file now is read earlier so that `~/my.cnf` can override system-wide settings.

The fix for this problem had a side effect such that on Unix, MySQL programs looked for options in `~/my.cnf` rather than the standard location of `~/my.cnf`. That problem is addressed in Bug #38180. (Bug #20748)

References: See also: Bug #38180.

- **Incompatible Change:** A number of problems existed in the implementation of [MERGE](#) tables that could cause problems. The problems are summarized below:
 - Bug #26379: Combination of [FLUSH TABLE](#) and [REPAIR TABLE](#) corrupts a [MERGE](#) table. This was caused in a number of situations:
 1. A thread trying to lock a [MERGE](#) table performs busy waiting while [REPAIR TABLE](#) or a similar table administration task is ongoing on one or more of its [MyISAM](#) tables.

2. A thread trying to lock a `MERGE` table performs busy waiting until all threads that did `REPAIR TABLE` or similar table administration tasks on one or more of its `MyISAM` tables in `LOCK TABLES` segments do `UNLOCK TABLES`. The difference against problem #1 is that the busy waiting takes place after the administration task. It is terminated by `UNLOCK TABLES` only.
 3. Two `FLUSH TABLES` within a `LOCK TABLES` segment can invalidate the lock. This does not require a `MERGE` table. The first `FLUSH TABLES` can be replaced by any statement that requires other threads to reopen the table. In 5.0 and 5.1 a single `FLUSH TABLES` can provoke the problem.
- Bug #26867: Simultaneously executing `LOCK TABLES` and `REPAIR TABLE` on a `MERGE` table would result in memory/cpu hogging.

Trying DML on a `MERGE` table, which has a child locked and repaired by another thread, made an infinite loop in the server.

- Bug #26377: Deadlock with `MERGE` and `FLUSH TABLE`

Locking a `MERGE` table and its children in parent-child order and flushing the child deadlocked the server.

- Bug #25038: Waiting `TRUNCATE TABLE`

Truncating a `MERGE` child, while the `MERGE` table was in use, let the truncate fail instead of waiting for the table to become free.

- Bug #25700: `MERGE` base tables get corrupted by `OPTIMIZE TABLE`, `ANALYZE TABLE`, or `REPAIR TABLE`.

Repairing a child of an open `MERGE` table corrupted the child. It was necessary to `FLUSH` the child first.

- Bug #30275: `MERGE` tables: `FLUSH TABLES` or `UNLOCK TABLES` causes server to crash.

Flushing and optimizing locked `MERGE` children crashed the server.

- Bug #19627: temporary merge table locking

Use of a temporary `MERGE` table with nontemporary children could corrupt the children.

Temporary tables are never locked. Creation of tables with nontemporary children of a temporary `MERGE` table is now prohibited.

- Bug #27660: `Falcon`: `MERGE` table possible

It was possible to create a `MERGE` table with non-`MyISAM` children.

- Bug #30273: `MERGE` tables: Can't lock file (errno: 155)

This was a Windows-only bug. Table administration statements sometimes failed with "Can't lock file (errno: 155)".

The fix introduces the following changes in behavior:

- This patch changes the behavior of temporary `MERGE` tables. Temporary `MERGE` must have temporary children. The old behavior was wrong. A temporary table is not locked. Hence even nontemporary children were not locked. See Bug #19627.

- You cannot change the union list of a nontemporary [MERGE](#) table when [LOCK TABLES](#) is in effect. The following does *not* work:

```
CREATE TABLE m1 ... ENGINE=MRG_MYISAM ...;
LOCK TABLES t1 WRITE, t2 WRITE, m1 WRITE;
ALTER TABLE m1 ... UNION=(t1,t2) ...;
```

However, you can do this with a temporary [MERGE](#) table.

- You cannot create a [MERGE](#) table with [CREATE ... SELECT](#), neither as a temporary [MERGE](#) table, nor as a nontemporary [MERGE](#) table. For example, [CREATE TABLE m1 ... ENGINE=MRG_MYISAM ... SELECT ...;](#) causes the error message: `table is not BASE TABLE`.

(Bug #19627, Bug #25038, Bug #25700, Bug #26377, Bug #26379, Bug #26867, Bug #27660, Bug #30275, Bug #30491)

- Important Change; MySQL Cluster:** [AUTO_INCREMENT](#) columns had the following problems when used in [NDB](#) tables:
 - The [AUTO_INCREMENT](#) counter was not updated correctly when such a column was updated.
 - [AUTO_INCREMENT](#) values were not prefetched beyond statement boundaries.
 - [AUTO_INCREMENT](#) values were not handled correctly with [INSERT IGNORE](#) statements.
 - After being set, [ndb_autoincrement_prefetch_sz](#) showed a value of 1, regardless of the value it had actually been set to.

As part of this fix, the behavior of [ndb_autoincrement_prefetch_sz](#) has changed. Setting this to less than 32 no longer has any effect on prefetching within statements (where IDs are now always obtained in batches of 32 or more), but only between statements. The default value for this variable has also changed, and is now 1. (Bug #25176, Bug #31956, Bug #32055)
- Important Note; Partitioning:** An apostrophe or single quote character (') used in the [DATA DIRECTORY](#), [INDEX DIRECTORY](#), or [COMMENT](#) for a [PARTITION](#) clause caused the server to crash. When used as part of a [CREATE TABLE](#) statement, the crash was immediate. When used in an [ALTER TABLE](#) statement, the crash did not occur until trying to perform a [SELECT](#) or DML statement on the table. In either case, the server could not be completely restarted until the `.frm` file corresponding to the newly created or altered table was deleted.



Note

Upgrading to the current (or later) release solves this problem only for tables that are newly created or altered. Tables created or altered in previous versions of the server to include ' characters in [PARTITION](#) options must still be removed by deleting the corresponding `.frm` files and re-creating them afterward.

(Bug #30695)

- Important Note:** The [RENAME DATABASE](#) statement was removed and replaced with [ALTER DATABASE db_name UPGRADE DATA DIRECTORY NAME](#). The [RENAME DATABASE](#) statement was intended for upgrading database directory names to the encoding format used in 5.1 for representing identifiers in the file system (see [Mapping of Identifiers to File Names](#)). However, the statement was found to be dangerous because it could result in loss of database contents. See [RENAME DATABASE Syntax](#), and [ALTER DATABASE Syntax](#). (Bug #17565, Bug #21741, Bug #28360)

- **MySQL Cluster; Replication:** Row-based replication from or to a big-endian machine where the table used the [NDB](#) storage engine failed, if the same table on the other machine was either non-[NDB](#) or the other machine was little-endian. (Bug #29549, Bug #30790)

References: See also: Bug #24231, Bug #30024, Bug #30133, Bug #30134.

- **MySQL Cluster:** An improperly reset internal signal was observed as a hang when using events in the [NDB](#) API but could result in various errors. (Bug #33206)
- **MySQL Cluster:** Incorrectly handled parameters could lead to a crash in the Transaction Coordinator during a node failure, causing other data nodes to fail. (Bug #33168)
- **MySQL Cluster:** A memory leak occurred if a subscription start request was received by the subscription manager before the node making the request was fully connected to the cluster. (Bug #32652)
- **MySQL Cluster:** A local checkpoint could sometimes be started before the previous LCP was restorable from a global checkpoint. (Bug #32519)
- **MySQL Cluster:** High numbers of API nodes on a slow or congested network could cause connection negotiation to time out prematurely, leading to the following issues:
 - Excessive retries
 - Excessive CPU usage
 - Partially connected API nodes(Bug #32359)
- **MySQL Cluster:** When a [mysqld](#) acting as a cluster SQL node starts the [NDBCLUSTER](#) storage engine, there is a delay during which some necessary data structures cannot be initialized until after it has connected to the cluster, and all MySQL Cluster tables should be opened as read only. This worked correctly when the [NDB](#) binary log thread was running, but when it was not running, Cluster tables were not opened as read only even when the data structures had not yet been set up. (Bug #32275, Bug #33763)
- **MySQL Cluster:** The failure of a master node could lead to subsequent failures in local checkpointing. (Bug #32160)
- **MySQL Cluster:** The management server was slow to respond when no data nodes were connected to the cluster. This was most noticeable when running [SHOW](#) in the management client. (Bug #32023)
- **MySQL Cluster:** An error with an [if](#) statement in [sql/ha_ndbcluster.cc](#) could potentially lead to an infinite loop in case of failure when working with [AUTO_INCREMENT](#) columns in [NDB](#) tables. (Bug #31810)
- **MySQL Cluster:** The [NDB](#) storage engine code was not safe for strict-alias optimization in [gcc](#) 4.2.1. (Bug #31761)
- **MySQL Cluster:** It was possible in some cases for a node group to be “lost” due to missed local checkpoints following a system restart. (Bug #31525)
- **MySQL Cluster:** A query against a table with [TEXT](#) or [BLOB](#) columns that would return more than a certain amount of data failed with `Got error 4350 'Transaction already aborted' from NDBCLUSTER`. (Bug #31482)

References: This issue is a regression of: Bug #29102.

- **MySQL Cluster:** `NDB` tables having names containing nonalphanumeric characters (such as “\$”) were not discovered correctly. (Bug #31470)
- **MySQL Cluster:** A node failure during a local checkpoint could lead to a subsequent failure of the cluster during a system restart. (Bug #31257)
- **MySQL Cluster:** A cluster restart could sometimes fail due to an issue with table IDs. (Bug #30975)
- **MySQL Cluster:** When handling `BLOB` columns, the addition of read locks to the lock queue was not handled correctly. (Bug #30764)
- **MySQL Cluster:** Discovery of `NDB` tables did not work correctly with `INFORMATION_SCHEMA`. (Bug #30667)
- **MySQL Cluster:** A file system close operation could fail during a node or system restart. (Bug #30646)
- **MySQL Cluster:** Transaction timeouts were not handled well in some circumstances, leading to excessive number of transactions being aborted unnecessarily. (Bug #30379)
- **MySQL Cluster:** The cluster management client could not connect, and would hang instead. This issue affected Mac OS X 64-bit only. (Bug #30366)
- **MySQL Cluster:** Attempting to restore a backup made on a cluster host using one endian to a machine using the other endian could cause the cluster to fail. (Bug #29674)
- **MySQL Cluster:** Log event requests to `ndb_mgmd` could time out, causing it to fail. (Bug #29621)
- **MySQL Cluster:** In some cases, the cluster management server logged entries multiple times following a restart of `ndb_mgmd`. (Bug #29565)
- **MySQL Cluster:** `ndb_mgm --help` did not display any information about the `-a` option. (Bug #29509)
- **MySQL Cluster:** An interpreted program of sufficient size and complexity could cause all cluster data nodes to shut down due to buffer overruns. (Bug #29390)
- **MySQL Cluster:** `ndb_size.pl` failed on tables with `FLOAT` columns whose definitions included commas (for example, `FLOAT(6,2)`). (Bug #29228)
- **MySQL Cluster:** The error message for `NDB` error code 275 (`Out of transaction records for complete phase`) was missing. (Bug #29139)
- **MySQL Cluster:** Reads on `BLOB` columns were not locked when they needed to be to guarantee consistency. (Bug #29102)

References: See also: Bug #31482.

- **MySQL Cluster:** A query using joins between several large tables and requiring unique index lookups failed to complete, eventually returning `Unknown Error` after a very long period of time. This occurred due to inadequate handling of instances where the Transaction Coordinator ran out of `TransactionBufferMemory`, when the cluster should have returned `NDB` error code 4012 (`Request ndbd time-out`). (Bug #28804)
- **MySQL Cluster:** There was a short interval during the startup process prior to the beginning of heartbeat detection such that, were an API or management node to reboot or a network failure to occur, data nodes could not detect this, with the result that there could be a lingering connection. (Bug #28445)
- **MySQL Cluster:** The description of the `--print` option provided in the output from `ndb_restore --help` was incorrect. (Bug #27683)

- **MySQL Cluster:** Restoring a backup made on a cluster host using one endian to a machine using the other endian failed for `BLOB` and `DATETIME` columns. (Bug #27543, Bug #30024)
- **MySQL Cluster:** An invalid subselect on an `NDB` table could cause `mysqld` to crash. (Bug #27494)
- **MySQL Cluster:** An attempt to perform a `SELECT ... FROM INFORMATION_SCHEMA.TABLES` whose result included information about `NDB` tables for which the user had no privileges crashed the MySQL Server on which the query was performed. (Bug #26793)
- **MySQL Cluster:** Performing `DELETE` operations after a data node had been shut down could lead to inconsistent data following a restart of the node. (Bug #26450)
- **MySQL Cluster:** `UPDATE IGNORE` could sometimes fail on `NDB` tables due to the use of uninitialized data when checking for duplicate keys to be ignored. (Bug #25817)
- **MySQL Cluster:** The cluster log was formatted inconsistently and contained extraneous newline characters. (Bug #25064)
- **MySQL Cluster:** A restart of the cluster failed when more than 1 REDO phase was in use. (Bug #22696)
- **MySQL Cluster:** When inserting a row into an `NDB` table with a duplicate value for a nonprimary unique key, the error issued would reference the wrong key.

This improves on an initial fix for this issue made in MySQL 5.1.13. (Bug #21072)

- **MySQL Cluster:** An insufficiently descriptive and potentially misleading Error 4006 (`Connect failure - out of connection objects...`) was produced when either of the following two conditions occurred:
 1. There were no more transaction records in the transaction coordinator
 2. An `NDB` object in the NDB API was initialized with insufficient parallelismSeparate error messages are now generated for each of these two cases. (Bug #11313)
- **Partitioning; Replication:** Replication of partitioned tables using the `InnoDB` storage engine failed with `binlog-format=ROW` or `binlog-format=MIXED`. (Bug #28430)
- **Partitioning:** It was possible to partition a table to which a foreign key referred. (Bug #32948)
- **Partitioning:** A query of the form `SELECT col1 FROM table GROUP BY (SELECT col2 FROM table LIMIT 1);` against a partitioned `table` having a `SET` column crashed the server. (Bug #32772)
- **Partitioning:** `SHOW CREATE TABLE` misreported the value of `AUTO_INCREMENT` for partitioned tables using either of the `InnoDB` or `ARCHIVE` storage engines. (Bug #32247)
- **Partitioning:** Selecting from `INFORMATION_SCHEMA.PARTITIONS` while partition management statements (for example, `ALTER TABLE ... ADD PARTITION`) were executing caused the server to crash. (Bug #32178)
- **Partitioning:** An error in the internal function `mysql_unpack_partition()` led to a fatal error in subsequent calls to `open_table_from_share()`. (Bug #32158)
- **Partitioning:** Repeated updates of a table that was partitioned by `KEY` on a `TIMESTAMP` column eventually crashed the server. (Bug #32067)
- **Partitioning:** Changing the storage engine used by a table having subpartitions led to a server crash. (Bug #31893)
- **Partitioning:** `ORDER BY ... DESC` did not always work correctly when selecting from partitioned tables. (Bug #31890)

References: See also: Bug #31001.

- **Partitioning:** Selecting from a table partitioned by `KEY` on a `VARCHAR` column whose size was greater than 65530 caused the server to crash. (Bug #31705)
- **Partitioning:** `INSERT DELAYED` on a partitioned table crashed the server. The server now rejects the statement with an error. (Bug #31210)
- **Partitioning:** Using `ALTER TABLE` to partition an existing table having an `AUTO_INCREMENT` column could crash the server. (Bug #30878)

References: This issue is a regression of: Bug #27405.

- **Partitioning:** `ALTER TABLE ... COALESCE PARTITION` on a table partitioned by `[LINEAR] HASH` or `[LINEAR] KEY` caused the server to crash. (Bug #30822)
- **Partitioning:** `LIKE` queries on tables partitioned by `KEY` and using third-party storage engines could return incomplete results. (Bug #30480)

References: See also: Bug #29320, Bug #29493, Bug #30563.

- **Partitioning:** It was not possible to insert the greatest possible value for a given data type into a partitioned table. For example, consider a table defined as shown here:

```
CREATE TABLE t (c BIGINT UNSIGNED)
  PARTITION BY RANGE(c) (
    PARTITION p0 VALUES LESS THAN MAXVALUE
  );
```

The largest possible value for a `BIGINT UNSIGNED` column is 18446744073709551615, but the statement `INSERT INTO t VALUES (18446744073709551615);` failed, even though the same statement succeeded were `t` not a partitioned table.

In other words, `MAXVALUE` was treated as being equal to the greatest possible value, rather than as a least upper bound. (Bug #29258)

- **Replication; Cluster Replication:** A node failure during replication could lead to buckets out of order; now active subscribers are checked for, rather than empty buckets. (Bug #31701)
- **Replication; Cluster Replication:** Incorrect handling of `INSERT` plus `DELETE` operations with regard to local checkpoints caused data node failures in multi-master replication setups. (Bug #30914)
- **Replication:** When dropping a database containing a stored procedure while using row-based replication, the delete of the stored procedure from the `mysql.proc` table was recorded in the binary log following the `DROP DATABASE` statement. To correct this issue, `DROP DATABASE` now uses statement-based replication. (Bug #32435)
- **Replication:** It was possible for the name of the relay log file to exceed the amount of memory reserved for it, possibly leading to a crash of the server. (Bug #31836)

References: See also: Bug #28597.

- **Replication:** Corruption of log events caused the server to crash on 64-bit Linux systems having 4 GB or more of memory. (Bug #31793)
- **Replication:** Trying to replicate an update of a row that was missing on the slave led to a failure on the slave. (Bug #31702)

- **Replication:** Use of the @@hostname system variable in inserts in `mysql_system_tables_data.sql` did not replicate. The workaround is to select its value into a user variable (which does replicate) and insert that. (Bug #31167)
- **Replication:** Table names were displayed as binary “garbage” characters in slave error messages. The issue was observed on 64-bit Windows but may have effected other platforms. (Bug #30854)
- **Replication:** One thread could read uninitialized memory from the stack of another thread. This issue was only known to occur in a `mysqld` process acting as both a master and a slave. (Bug #30752)
- **Replication:** It was possible to set `SQL_SLAVE_SKIP_COUNTER` such that the slave would jump into the middle of a transaction. This fix improves on one made for this bug in MySQL 5.1.20; the previous fix insured that the slave could not be made to jump into the middle of an event group, but the slave failed to recognize that `BEGIN`, `COMMIT`, and `ROLLBACK` statements could begin or end an event group. (Bug #28618)

References: See also: Bug #12691.

- **Replication:** Due a previous change in how the default name and location of the binary log file were determined, replication failed following some upgrades. (Bug #28597, Bug #28603)

References: See also: Bug #31836. This issue is a regression of: Bug #20166.

- **Replication:** Stored procedures having `BIT` parameters were not replicated correctly. (Bug #26199)
- **Replication:** Issuing `SHOW SLAVE STATUS` as `mysqld` was shutting down could cause a crash. (Bug #26000)
- **Replication:** If a temporary error occurred inside an event group on an event that was not the first event of the group, the slave could get caught in an endless loop because the retry counter was reset whenever an event was executed successfully. (Bug #24860)

References: See also: Bug #12691, Bug #23171.

- **Replication:** An `UPDATE` statement using a stored function that modified a nontransactional table was not logged if it failed. This caused the copy of the nontransactional table on the master have a row that the copy on the slave did not.

In addition, when an `INSERT ... ON DUPLICATE KEY UPDATE` statement encountered a duplicate key constraint, but the `UPDATE` did not actually change any data, the statement was not logged. As a result of this fix, such statements are now treated the same for logging purposes as other `UPDATE` statements, and so are written to the binary log. (Bug #23333)

References: See also: Bug #12713.

- **Replication:** A replication slave sometimes failed to reconnect because it was unable to run `SHOW SLAVE HOSTS`. It was not necessary to run this statement on slaves (since the master should track connection IDs), and the execution of this statement by slaves was removed. (Bug #21132)

References: See also: Bug #13963, Bug #21869.

- **Replication:** A replication slave sometimes stopped for changes that were idempotent (that is, such changes should have been considered “safe”), even though it should have simply noted that the change was already done, and continued operation. (Bug #19958)
- **Replication:** Replicating from a master table to a slave table where the size of a `CHAR` or `VARCHAR` column was a different size would cause `mysqld` to crash. For more information on replicating with different column definitions, see [Replication with Differing Table Definitions on Master and Slave](#).

- **Cluster Replication:** A replication slave could return “garbage” data that was not in recognizable row format due to a problem with the internal `all_set()` method. (Bug #33375)
- **Cluster Replication:** Memory was mistakenly freed for `NdbBlob` objects when adding an index while replicating the cluster, which could cause `mysqld` to crash. (Bug #33142)

References: See also: Bug #18106.

- **Cluster Replication:** Under certain conditions, the slave stopped processing relay logs. This resulted in the logs never being cleared and the slave eventually running out of disk space. (Bug #31958)
- **Cluster Replication:** Replicating `NDB` tables with extra `VARCHAR` columns on the master caused the slave to fail. (Bug #31646)

References: See also: Bug #29549.

- **Cluster Replication:** When the master `mysqld` crashed or was restarted, no `LOST_EVENTS` entry was made in the binlog. (Bug #31484)

References: See also: Bug #21494.

- **Cluster Replication:** An issue with the `mysql.ndb_apply_status` table could cause `NDB` schema autodiscovery to fail in certain rare circumstances. (Bug #20872)
- **Cluster API:** A call to `CHECK_TIMEDOUT_RET()` in `mgmapi.cpp` should have been a call to `DEBUG_CHECK_TIMEDOUT_RET()`. (Bug #30681)
- **API:** When the language option was not set correctly, API programs calling `mysql_server_init()` crashed. This issue was observed only on Windows platforms. (Bug #31868)
- Corrected a typecast involving `bool` on Mac OS X 10.5 (Leopard), which evaluated differently from earlier Mac OS X versions. (Bug #38217)
- Use of uninitialized memory for `filesort` in a subquery caused a server crash. (Bug #33675)
- `CREATE TABLE ... SELECT` created tables that for date columns used the obsolete `Field_date` type instead of `Field_newdate`. (Bug #33256)
- Some valid `SELECT` statements could not be used as views due to incorrect column reference resolution. (Bug #33133)
- The fix for Bug #11230 and Bug #26215 introduced a significant input-parsing slowdown for the `mysql` client. This has been corrected. (Bug #33057)

References: See also: Bug #11230, Bug #26215.

- The correct data type for a `NULL` column resulting from a `UNION` could be determined incorrectly in some cases: 1) Not correctly inferred as `NULL` depending on the number of selects; 2) Not inferred correctly as `NULL` if one select used a subquery. (Bug #32848)
- For queries containing `GROUP_CONCAT(DISTINCT col_list ORDER BY col_list)`, there was a limitation that the `DISTINCT` columns had to be the same as `ORDER BY` columns. Incorrect results could be returned if this was not true. (Bug #32798)
- `SHOW EVENTS` and selecting from the `INFORMATION_SCHEMA.EVENTS` table failed if the current database was `INFORMATION_SCHEMA`. (Bug #32775)
- The `LAST_DAY()` function returns a `DATE` value, but internally the value did not have the time fields zeroed and calculations involving the value could return incorrect results. (Bug #32770)

- Use of the `cp932` character set with `CAST()` in an `ORDER BY` clause could cause a server crash. (Bug #32726)
- A subquery using an `IS NULL` check of a column defined as `NOT NULL` in a table used in the `FROM` clause of the outer query produced an invalid result. (Bug #32694)
- `mysqld_safe` looked for error messages in the wrong location. (Bug #32679)
- Specifying a nonexistent column for an `INSERT DELAYED` statement caused a server crash rather than producing an error. (Bug #32676)
- An issue with the `NO_ENGINE_SUBSTITUTION sql_mode` database can cause the creation of stored routines to fail. If you are having problems with creating stored routines while using this `sql_mode` value, remove this value from your `sql_mode` setting. (Bug #32633)
- Use of `CLIENT_MULTI_QUERIES` caused `libmysqld` to crash. (Bug #32624)
- The `INTERVAL()` function incorrectly handled `NULL` values in the value list. (Bug #32560)
- Use of a `NULL`-returning `GROUP BY` expression in conjunction with `WITH ROLLUP` could cause a server crash. (Bug #32558)

References: See also: Bug #31095.

- `ORDER BY UpdateXML(...)` caused the server to crash in queries where `UpdateXML()` returned `NULL`. (Bug #32557)
- A `SELECT ... GROUP BY bit_column` query failed with an assertion if the length of the `BIT` column used for the `GROUP BY` was not an integer multiple of 8. (Bug #32556)
- Using `SELECT INTO OUTFILE` with 8-bit `ENCLOSED BY` characters led to corrupted data when the data was reloaded using `LOAD DATA INFILE`. This was because `SELECT INTO OUTFILE` failed to escape the 8-bit characters. (Bug #32533)
- For `FLUSH TABLES WITH READ LOCK`, the server failed to properly detect write-locked tables when running with low-priority updates, resulting in a crash or deadlock. (Bug #32528)
- The rules for valid column names were being applied differently for base tables and views. (Bug #32496)
- A query of the form `SELECT @user_variable := constant AS alias FROM table GROUP BY alias WITH ROLLUP` crashed the server. (Bug #32482)
- Sending several `KILL QUERY` statements to target a connection running `SELECT SLEEP()` could freeze the server. (Bug #32436)
- `ssl-cipher` values in option files were not being read by `libmysqlclient`. (Bug #32429)
- Repeated execution of a query containing a `CASE` expression and numerous `AND` and `OR` relations could crash the server. The root cause of the issue was determined to be that the internal `SEL_ARG` structure was not properly initialized when created. (Bug #32403)
- Referencing within a subquery an alias used in the `SELECT` list of the outer query was incorrectly permitted. (Bug #32400)
- If a global read lock acquired with `FLUSH TABLES WITH READ LOCK` was in effect, executing `ALTER TABLE` could cause a server crash. (Bug #32395)
- An `ORDER BY` query on a view created using a `FEDERATED` table as a base table caused the server to crash. (Bug #32374)

- Comparison of a `BIGINT NOT NULL` column with a constant arithmetic expression that evaluated to NULL mistakenly caused the error `Column '...' cannot be null` (error 1048). (Bug #32335)
- Assigning a 65,536-byte string to a `TEXT` column (which can hold a maximum of 65,535 bytes) resulted in truncation without a warning. Now a truncation warning is generated. (Bug #32282)
- `MIN()` and `MAX()` could return incorrect results when an index was present if a loose index scan was used. (Bug #32268)
- Some uses of user variables in a query could result in a server crash. (Bug #32260)
- Memory corruption could occur due to large index map in `Range checked for each record` status reported by `EXPLAIN SELECT`. The problem was based in an incorrectly calculated length of the buffer used to store a hexadecimal representation of an index map, which could result in buffer overrun and stack corruption under some circumstances. (Bug #32241)
- Various test program cleanups were made: 1) `mytest` and `libmysqltest` were removed. 2) `bug25714` displays an error message when invoked with incorrect arguments or the `--help` option. 3) `mysql_client_test` exits cleanly with a proper error status. (Bug #32221)
- The default grant tables on Windows contained information for host `production.mysql.com`, which should not be there. (Bug #32219)
- Under certain conditions, the presence of a `GROUP BY` clause could cause an `ORDER BY` clause to be ignored. (Bug #32202)
- For comparisons of the form `date_col OP datetime_const` (where `OP` is `=`, `<`, `>`, `<=`, or `>=`), the comparison is done using `DATETIME` values, per the fix for Bug #27590. However that fix caused any index on `date_col` not to be used and compromised performance. Now the index is used again. (Bug #32198)

References: See also: Bug #27590.

- `DATETIME` arguments specified in numeric form were treated by `DATE_ADD()` as `DATE` values. (Bug #32180)
- Killing a statement could lead to a race condition in the server. (Bug #32148)
- `InnoDB` does not support `SPATIAL` indexes, but could crash when asked to handle one. Now an error is returned. (Bug #32125)
- The server crashed on optimizations involving a join of `INT` and `MEDIUMINT` columns and a system variable in the `WHERE` clause. (Bug #32103)
- `mysql-test-run.pl` used the `--user` option when starting `mysqld`, which produces warnings if the current user is not `root`. Now `--user` is added only for `root`. (Bug #32078)
- `mysqlslap` was missing from the MySQL 5.1.22 Linux RPM packages. (Bug #32077)
- With `lower_case_table_names` set, `CREATE TABLE LIKE` was treated differently by `libmysqld` than by the nonembedded server. (Bug #32063)
- Within a subquery, `UNION` was handled differently than at the top level, which could result in incorrect results or a server crash. (Bug #32036, Bug #32051)
- On 64-bit platforms, assignments of values to enumeration-valued storage engine-specific system variables were not validated and could result in unexpected values. (Bug #32034)

- A `DELETE` statement with a subquery in the `WHERE` clause would sometimes ignore an error during subquery evaluation and proceed with the delete operation. (Bug #32030)
- Using dates in the range `'0000-00-01'` to `'0000-00-99'` range in the `WHERE` clause could result in an incorrect result set. (These dates are not in the supported range for `DATE`, but different results for a given query could occur depending on position of records containing the dates within a table.) (Bug #32021)
- User-defined functions are not loaded if the server is started with the `--skip-grant-tables` option, but the server did not properly handle this case and issued an `Out of memory` error message instead. (Bug #32020)
- If a user-defined function was used in a `SELECT` statement, and an error occurred during UDF initialization, the error did not terminate execution of the `SELECT`, but rather was converted to a warning. (Bug #32007)
- `hour()`, `minute()`, and `second()` could return nonzero values for `DATE` arguments. (Bug #31990)
- Changing the SQL mode to cause dates with “zero” parts to be considered invalid (such as `'1000-00-00'`) could result in indexed and nonindexed searches returning different results for a column that contained such dates. (Bug #31928)
- The server used unnecessarily large amounts of memory when user variables were used as an argument to `concat()` or `concat_ws()`. (Bug #31898)
- In debug builds, testing the result of an `IN` subquery against `NULL` caused an assertion failure. (Bug #31884)
- `mysql-test-run.pl` sometimes set up test scenarios in which the same port number was passed to multiple servers, causing one of them to be unable to start. (Bug #31880)
- `SHOW CREATE TRIGGER` caused a server crash. (Bug #31866)
- The server crashed after insertion of a negative value into an `AUTO_INCREMENT` column of an `InnoDB` table. (Bug #31860)
- For `libmysqld` applications, handling of `mysql_change_user()` calls left some pointers improperly updated, leading to server crashes. (Bug #31850)
- Using `ORDER BY` led to the wrong result when using the `ARCHIVE` on a table with a `BLOB` when the table cache was full. The table could also be reported as crashed after the query had completed, even though the table data was intact. (Bug #31833)
- Comparison results for `BETWEEN` were different from those for operators like `<` and `>` for `DATETIME`-like values with trailing extra characters such as `'2007-10-01 00:00:00 GMT-6'`. `BETWEEN` treated the values as `DATETIME`, whereas the other operators performed a binary-string comparison. Now they all uniformly use a `DATETIME` comparison, but generate warnings for values with trailing garbage. (Bug #31800)
- Name resolution for correlated subqueries and `HAVING` clauses failed to distinguish which of two was being performed when there was a reference to an outer aliased field. This could result in error messages about a `HAVING` clause for queries that had no such clause. (Bug #31797)
- The server could crash during `filesort` for `ORDER BY` based on expressions with `inet_ntoa()` or `oct()` if those functions returned `NULL`. (Bug #31758)
- For tables with certain definitions, `UPDATE` statements could fail to find the correct record to update and report an error when the record did in fact exist. (Bug #31747)

- For a fatal error during a filesort in `find_all_keys()`, the error was returned without the necessary handler uninitialization, causing an assertion failure. (Bug #31742)
- `mysqlslap` failed to commit after the final record load. (Bug #31704)
- The examined-rows count was not incremented for `const` queries. (Bug #31700)
- The server crashed if a thread was killed while locking the `general_log` table at the beginning of statement processing. (Bug #31692)
- The `mysql_change_user()` C API function was subject to buffer overflow. (Bug #31669)
- For `SELECT ... INTO OUTFILE`, if the `ENCLOSED BY` string is empty and the `FIELDS TERMINATED BY` string started with a special character (one of `n`, `t`, `r`, `b`, `0`, `Z`, or `N`), every occurrence of the character within field values would be duplicated. (Bug #31663)
- `SHOW COLUMNS` and `DESCRIBE` displayed `null` as the column type for a view with no valid definer. This caused `mysqldump` to produce a nonreloadable dump file for the view. (Bug #31662)
- The `mysqlbug` script did not include the correct values of `CFLAGS` and `CXXFLAGS` that were used to configure the distribution. (Bug #31644)
- Queries that include a comparison of an `INFORMATION_SCHEMA` table column to `NULL` caused a server crash. (Bug #31633)
- `EXPLAIN EXTENDED` for `SELECT` from `INFORMATION_SCHEMA` tables caused an assertion failure. (Bug #31630)
- `ucs2` does not work as a client character set, but attempts to use it as such were not rejected. Now `character_set_client` cannot be set to `ucs2`. This also affects statements such as `SET NAMES` and `SET CHARACTER SET`. (Bug #31615)
- A buffer used when setting variables was not dimensioned to accommodate the trailing `'\0'` byte, so a single-byte buffer overrun was possible. (Bug #31588)
- `HAVING` could treat lettercase of table aliases incorrectly if `lower_case_table_names` was enabled. (Bug #31562)
- Spurious duplicate-key errors could occur for multiple-row inserts into an `InnoDB` table that activate a trigger. (Bug #31540)
- Using `ALTER EVENT` to rename a disabled event caused it to become enabled. (Bug #31539)
- The fix for Bug #24989 introduced a problem such that a `NULL` thread handler could be used during a rollback operation. This problem is unlikely to be seen in practice. (Bug #31517)
- The length of the result from `IFNULL()` could be calculated incorrectly because the sign of the result was not taken into account. (Bug #31471)
- Queries that used the `ref` access method or index-based subquery execution over indexes that have `DECIMAL` columns could fail with an error `Column col_name cannot be null`. (Bug #31450)
- `InnoDB` now tracks locking and use of tables by MySQL only after a table has been successfully locked on behalf of a transaction. Previously, the locked flag was set and the table in-use counter was updated before checking whether the lock on the table succeeded. A subsequent failure in obtaining a lock on the table led to an inconsistent state as the table was neither locked nor in use. (Bug #31444)
- `SELECT 1 REGEX NULL` caused an assertion failure for debug servers. (Bug #31440)

- The `UpdateXML()` function did not check for the validity of all its arguments; in some cases, this could lead to a crash of the server. (Bug #31438)
- The `mysql_change_user()` C API function caused advisory locks (obtained with `GET_LOCK()`) to malfunction. (Bug #31418)
- NDB libraries and include files were missing from some binary `tar` file distributions. (Bug #31414)
- Executing `RENAME` while tables were open for use with `HANDLER` statements could cause a server crash. (Bug #31409)
- `mysql-test-run.pl` tried to create files in a directory where it could not be expected to have write permission. `mysqltest` created `.reject` files in a directory other than the one where test results go. (Bug #31398)
- For a table that had been opened with `HANDLER` and marked for reopening after being closed with `FLUSH TABLES, DROP TABLE` did not properly discard the handler. (Bug #31397)
- Automatically allocated memory for string options associated with a plugin was not freed if the plugin did not get installed. (Bug #31382)
- `INFORMATION_SCHEMA.TABLES` was returning incorrect information. (Bug #31381)
- `DROP USER` caused an increase in memory usage. (Bug #31347)
- For `InnoDB` tables with `READ COMMITTED` isolation level, semi-consistent reads used for `UPDATE` statements skipped rows locked by another transaction, rather than waiting for the locks to be released. Consequently, rows that possibly should have been updated were never examined. (Bug #31310)
- For an almost-full `MyISAM` table, an insert that failed could leave the table in a corrupt state. (Bug #31305)
- `myisamchk --unpack` could corrupt a table that when unpacked has static (fixed-length) row format. (Bug #31277)
- `CONVERT(val, DATETIME)` failed on invalid input, but processing was not aborted for the `WHERE` clause, leading to a server crash. (Bug #31253)
- Allocation of an insufficiently large group-by buffer following creation of a temporary table could lead to a server crash. (Bug #31249)
- Use of `DECIMAL(n, n) ZEROFILL` in `GROUP_CONCAT()` could cause a server crash. (Bug #31227)
- When a `TIMESTAMP` with a nonzero time part was converted to a `DATE` value, no warning was generated. This caused index lookups to assume that this is a valid conversion and was returning rows that match a comparison between a `TIMESTAMP` value and a `DATE` keypart. Now a warning is generated so that `TIMESTAMP` with a nonzero time part will not match `DATE` values. (Bug #31221)
- Server variables could not be set to their current values on Linux platforms. (Bug #31177)
References: See also: Bug #6958.
- With small values of `myisam_sort_buffer_size`, `REPAIR TABLE` for `MyISAM` tables could cause a server crash. (Bug #31174)
- If `MAKETIME()` returned `NULL` when used in an `ORDER BY` that was evaluated using `filesort`, a server crash could result. (Bug #31160)
- Data in `BLOB` or `GEOMETRY` columns could be cropped when performing a `UNION` query. (Bug #31158)

- `LAST_INSERT_ID()` execution could be handled improperly in subqueries. (Bug #31157)
- An assertion designed to detect a bug in the `ROLLUP` implementation would incorrectly be triggered when used in a subquery context with noncacheable statements. (Bug #31156)
- Selecting spatial types in a `UNION` could cause a server crash. (Bug #31155)
- Use of `GROUP_CONCAT(DISTINCT bit_column)` caused an assertion failure. (Bug #31154)
- The server crashed in the parser when running out of memory. Memory handling in the parser has been improved to gracefully return an error when out-of-memory conditions occur in the parser. (Bug #31153)
- MySQL declares a `UNIQUE` key as a `PRIMARY` key if it doesn't have `NULL` columns and is not a partial key, and the `PRIMARY` key must always be the first key. However, in some cases, a nonfirst key could be reported as `PRIMARY`, leading to an assert failure by `InnoDB`. This is fixed by correcting the key sort order. (Bug #31137)
- `mysqldump` failed to handle databases containing a '-' character in the name. (Bug #31113)
- Starting the server with `read_only` enabled and with the Event Scheduler enabled caused it to crash.

**Note**

This issue occurred only when the server had been built with certain nonstandard combinations of `configure` options.

(Bug #31111)

- `GROUP BY NULL WITH ROLLUP` could cause a server crash. (Bug #31095)
- References: See also: Bug #32558.
- A rule to prefer `filesort` over an indexed `ORDER BY` when accessing all rows of a table was being used even if a `LIMIT` clause was present. (Bug #31094)
 - `REGEXP` operations could cause a server crash for character sets such as `ucs2`. Now the arguments are converted to `utf8` if possible, to permit correct results to be produced if the resulting strings contain only 8-bit characters. (Bug #31081)
 - Expressions of the form `WHERE col NOT IN (col, ...)`, where the same column was named both times, could cause a server crash in the optimizer. (Bug #31075)
 - Internal conversion routines could fail for several multibyte character sets (`big5`, `cp932`, `euckr`, `gb2312`, `sjis`) for empty strings or during evaluation of `SOUNDS LIKE`. (Bug #31069, Bug #31070)
 - Many nested subqueries in a single query could lead to excessive memory consumption and possibly a crash of the server. (Bug #31048)
 - Using `ORDER BY` with `ARCHIVE` tables caused a server crash. (Bug #31036)
 - A server crash could occur when a non-`DETERMINISTIC` stored function was used in a `GROUP BY` clause. (Bug #31035)
 - The `MOD()` function and the `%` operator crashed the server for a divisor less than 1 with a very long fractional part. (Bug #31019)
 - Transactions were committed prematurely when `LOCK TABLE` and `SET autocommit = 0` were used together. (Bug #30996)
 - On Windows, the `pthread_mutex_trylock()` implementation was incorrect. (Bug #30992)

- A character set introducer followed by a hexadecimal or bit-value literal did not check its argument and could return an ill-formed result for invalid input. (Bug #30986)
- `CHAR(str USING charset)` did not check its argument and could return an ill-formed result for invalid input. (Bug #30982)
- The result from `CHAR(str USING ucs2)` did not add a leading 0x00 byte for input strings with an odd number of bytes. (Bug #30981)
- The `GeomFromText()` function could cause a server crash if the first argument was `NULL` or the empty string. (Bug #30955)
- `MAKEDATE()` incorrectly moved year values in the 100 to 200 range into the 1970 to 2069 range. (This is legitimate for 00 to 99, but three-digit years should be used unchanged.) (Bug #30951)
- When invoked with constant arguments, `STR_TO_DATE()` could use a cached value for the format string and return incorrect results. (Bug #30942)
- `GROUP_CONCAT()` returned `' , '` rather than an empty string when the argument column contained only empty strings. (Bug #30897)
- For `MEMORY` tables, lookups for `NULL` values in `BTREE` indexes could return incorrect results. (Bug #30885)
- A server crash could occur if a stored function that contained a `DROP TEMPORARY TABLE` statement was invoked by a `CREATE TEMPORARY TABLE` statement that created a table of the same name. (Bug #30882)
- Calling `NAME_CONST()` with nonconstant arguments triggered an assertion failure. Nonconstant arguments are no longer permitted. (Bug #30832)
- For a spatial column with a regular (non-`SPATIAL`) index, queries failed if the optimizer tried to use the index. (Bug #30825)
- Values for the `--tc-heuristic-recover` option incorrectly were treated as values for the `--myisam-stats-method` option. (Bug #30821)
- `INFORMATION_SCHEMA.SCHEMATA` was returning incorrect information. (Bug #30795)
- The optimizer incorrectly optimized conditions out of the `WHERE` clause in some queries involving subqueries and indexed columns. (Bug #30788)
- Improper calculation of `CASE` expression results could lead to value truncation. (Bug #30782)
- On Windows, the `pthread_mutex_trylock()` implementation was incorrect. One symptom was that invalidating the query cache could cause a server crash. (Bug #30768)
- A multiple-table `UPDATE` involving transactional and nontransactional tables caused an assertion failure. (Bug #30763)
- User-supplied names foreign key names might not be set to the right key, leading to foreign keys with no name. (Bug #30747)
- Under some circumstances, `CREATE TABLE ... SELECT` could crash the server or incorrectly report that the table row size was too large. (Bug #30736)
- Using the `MIN()` or `MAX()` function to select one part of a multi-part key could cause a crash when the function result was `NULL`. (Bug #30715)
- The embedded server did not properly check column-level privileges. (Bug #30710)

- `INFORMATION_SCHEMA.VIEWS.VIEW_DEFINITION` was incorrect for views that were defined to select from other `INFORMATION_SCHEMA` tables. (Bug #30689)
- Issuing an `ALTER SERVER` statement to update the settings for a `FEDERATED` server would cause the `mysqld` to crash. (Bug #30671)
- The optimizer could ignore `ORDER BY` in cases when the result set is ordered by `filesort`, resulting in rows being returned in incorrect order. (Bug #30666)
- A different execution plan was displayed for `EXPLAIN` than would actually have been used for the `SELECT` because the test of sort keys for `ORDER BY` did not consider keys mentioned in `IGNORE KEYS FOR ORDER BY`. (Bug #30665)
- The `thread_handling` system variable was treated as having a `SESSION` value and as being settable at runtime. Now it has only a `GLOBAL` read-only value. (Bug #30651)
- On Windows, `LIMIT` arguments greater than 2^{32} did not work correctly. (Bug #30639)
- `MyISAM` tables could not exceed $4294967295 (2^{32} - 1)$ rows on Windows. (Bug #30638)
- A failed `HANDLER ... READ` operation could leave the table in a locked state. (Bug #30632)
- `mysql-test-run.pl` could not run `mysqld` with `root` privileges. (Bug #30630)
- The `mysqld_safe` script contained a syntax error. (Bug #30624)
- The optimization that uses a unique index to remove `GROUP BY` did not ensure that the index was actually used, thus violating the `ORDER BY` that is implied by `GROUP BY`. (Bug #30596)
- `SHOW STATUS LIKE 'Ssl_cipher_list'` from a MySQL client connected using SSL returned an empty string rather than a list of available ciphers. (Bug #30593)
- For `MEMORY` tables, `DELETE` statements that remove rows based on an index read could fail to remove all matching rows. (Bug #30590)
- Using `GROUP BY` on an expression of the form `timestamp_col DIV number` caused a server crash due to incorrect calculation of number of decimals. (Bug #30587)
- Executing a `SELECT COUNT(*)` query on an `InnoDB` table partitioned by `KEY` that used a `DOUBLE` column as the partitioning key caused the server to crash. (Bug #30583)
- The options available to the `CHECK TABLE` statement were also permitted in `OPTIMIZE TABLE` and `ANALYZE TABLE` statements, but caused corruption during their execution. These options were never supported for these statements, and an error is now raised if you try to apply these options to these statements. (Bug #30495)
- A self-referencing trigger on a partitioned table caused the server to crash instead of failing with an error. (Bug #30484)
- The `mysql_change_user()` C API function did not correctly reset the character set variables to the values they had just after initially connecting. (Bug #30472)
- When expanding a `*` in a `USING` or `NATURAL` join, the check for table access for both tables in the join was done using only the grant information of the first table. (Bug #30468)
- When casting a string value to an integer, cases where the input string contained a decimal point and was long enough to overrun the `unsigned long long` type were not handled correctly. The position of the decimal point was not taken into account which resulted in miscalculated numbers and incorrect truncation to appropriate SQL data type limits. (Bug #30453)

- Versions of `mysqldump` from MySQL 4.1 or higher tried to use `START TRANSACTION WITH CONSISTENT SNAPSHOT` if the `--single-transaction` and `--master-data` options were given, even with servers older than 4.1 that do not support consistent snapshots. (Bug #30444)
- With `libmysqld`, use of prepared statements and the query cache at the same time caused problems. (Bug #30430)
- Issuing a `DELETE` statement having both an `ORDER BY` clause and a `LIMIT` clause could cause `mysqld` to crash. (Bug #30385)
- For `CREATE ... SELECT ... FROM`, where the resulting table contained indexes, adding `SQL_BUFFER_RESULT` to the `SELECT` part caused index corruption in the table. (Bug #30384)
- The `Last_query_cost` status variable value can be computed accurately only for simple “flat” queries, not complex queries such as those with subqueries or `UNION`. However, the value was not consistently being set to 0 for complex queries. (Bug #30377)
- The optimizer made incorrect assumptions about the value of the `is_member` value for user-defined functions, sometimes resulting in incorrect ordering of UDF results. (Bug #30355)
- Queries that had a `GROUP BY` clause and selected `COUNT(DISTINCT bit_column)` returned incorrect results. (Bug #30324)
- Some valid `euc-kr` characters having the second byte in the ranges `[0x41..0x5A]` and `[0x61..0x7A]` were rejected. (Bug #30315)
- When loading a dynamic plugin on FreeBSD, the plugin failed to load. This was due to a build error where the required symbols would be not exported correctly. (Bug #30296)
- Simultaneous `ALTER TABLE` statements for `BLACKHOLE` tables caused 100% CPU use due to locking problems. (Bug #30294)
- Setting certain values on a table using a spatial index could cause the server to crash. (Bug #30286)
- Tables with a `GEOMETRY` column could be marked as corrupt if you added a non-`SPATIAL` index on a `GEOMETRY` column. (Bug #30284)
- Flushing a merge table between the time it was opened and its child table were actually attached caused the server to crash. (Bug #30273)

References: This issue is a regression of: Bug #26379.

- The query cache does not support retrieval of statements for which column level access control applies, but the server was still caching such statements, thus wasting memory. (Bug #30269)
- Using `DISTINCT` or `GROUP BY` on a `BIT` column in a `SELECT` statement caused the column to be cast internally as an integer, with incorrect results being returned from the query. (Bug #30245)
- `GROUP BY` on `BIT` columns produced incorrect results. (Bug #30219)
- Short-format `mysql` commands embedded within `/*! ... */` comments were parsed incorrectly by `mysql`, which discarded the rest of the comment including the terminating `*/` characters. The result was a malformed (unclosed) comment. Now `mysql` does not discard the `*/` characters. (Bug #30164)
- If the server crashed during an `ALTER TABLE` statement, leaving a temporary file in the database directory, a subsequent `DROP DATABASE` statement failed due to the presence of the temporary file. (Bug #30152)

- When `mysqldump` wrote `DROP DATABASE` statements within version-specific comments, it included the terminating semicolon in the wrong place, causing following statements to fail when the dump file was reloaded. (Bug #30126)
- It was not possible for client applications to distinguish between auto-set and auto-updated `TIMESTAMP` column values.

To rectify this problem, a new `ON_UPDATE_NOW_FLAG` flag is set by `Field_timestamp` constructors whenever a column should be set to `NOW` on `UPDATE`, and the `get_schema_column_record()` function now reports whether a timestamp column is set to `NOW` on `UPDATE`. In addition, such columns now display `on update CURRENT_TIMESTAMP` in the `Extra` column in the output from `SHOW COLUMNS`. (Bug #30081)

- Some `INFORMATION_SCHEMA` tables are intended for internal use, but could be accessed by using `SHOW` statements. (Bug #30079)
- On some 64-bit systems, inserting the largest negative value into a `BIGINT` column resulted in incorrect data. (Bug #30069)
- `mysqlslap` did not properly handle multiple result sets from stored procedures. (Bug #29985)
- Statements within stored procedures ignored the value of the `low_priority_updates` system variable. (Bug #29963)

References: See also: Bug #26162.

- With auto-reconnect enabled, row fetching for a prepared statement could crash after reconnect occurred because loss of the statement handler was not accounted for. (Bug #29948)
- `mysqldump --skip-events --all-databases` dumped data from the `mysqld.event` table, and when restoring from this dump, events were created in spite of the `--skip-events` option. (Bug #29938)
- When `mysqlslap` was given a query to execute from a file using a `--query=file_name` option, it executed the query one too many times. (Bug #29803)
- Specifying the `--without-geometry` option for `configure` caused server compilation to fail. (Bug #29772)
- `configure` did not find `nss` on some Linux platforms. (Bug #29658)
- It was possible when creating a partitioned table using `CREATE TABLE ... SELECT` to refer in the `PARTITION BY` clause to columns in the table being selected from, which could cause the server to crash. An example of such a statement is:

```
CREATE TABLE t1 (b INT)
PARTITION BY RANGE(t2.b) (
    PARTITION p1 VALUES LESS THAN (10),
    PARTITION p2 VALUES LESS THAN (20)
)
SELECT * FROM t2;
```

The fix is to disallow references in `PARTITION BY` clauses to columns not in the table being created. (Bug #29444)

- If a view used a function in its `SELECT` statement, the columns from the view were not inserted into the `INFORMATION_SCHEMA.COLUMNS` table. (Bug #29408)
- The `mysql` client program now ignores Unicode byte order mark (BOM) characters at the beginning of input files. Previously, it read them and sent them to the server, resulting in a syntax error.

Presence of a BOM does not cause `mysql` to change its default character set. To do that, invoke `mysql` with an option such as `--default-character-set=utf8`. (Bug #29323)

- For transactional tables, an error during a multiple-table `DELETE` statement did not roll back the statement. (Bug #29136)
- The `log` and `log_slow_queries` system variables were displayed by `SHOW VARIABLES` but could not be accessed in expressions as `@@log` and `@@log_slow_queries`. Also, attempting to set them with `SET` produced an incorrect `Unknown system variable` message. Now these variables are treated as synonyms for `general_log` and `slow_query_log`, which means that they can be accessed in expressions and their values can be changed with `SET`. (Bug #29131)
- Denormalized double-precision numbers cannot be handled properly by old MIPS processors. For IRIX, this is now handled by enabling a mode to use a software workaround. (Bug #29085)
- `SHOW VARIABLES` did not display the `relay_log`, `relay_log_index`, or `relay_log_info_file` system variables. (Bug #28893)
- When doing a `DELETE` on a table that involved a `JOIN` with `MyISAM` or `MERGE` tables and the `JOIN` referred to the same table, the operation could fail reporting `ERROR 1030 (HY000): Got error 134 from storage engine`. This was because scans on the table contents would change because of rows that had already been deleted. (Bug #28837)
- Killing an SSL connection on platforms where MySQL is compiled with `-DSIGNAL_WITH_VIO_CLOSE` (Windows, Mac OS X, and some others) could crash the server. (Bug #28812)
- `SHOW VARIABLES` did not correctly display the value of the `thread_handling` system variable. (Bug #28785)
- On Windows, `mysql_upgrade` created temporary files in `C:\` and did not clean them up. (Bug #28774)
- Index hints specified in view definitions were ignored when using the view to select from the base table. (Bug #28702)
- Views do not have indexes, so index hints do not apply. Use of index hints when selecting from a view is no longer permitted. (Bug #28701)
- After changing the SQL mode to a restrictive value that would make already inserted dates in a column be considered invalid, searches returned different results depending on whether the column was indexed. (Bug #28687)
- When running the MySQL Instance Configuration Wizard, a race condition could exist that failed to connect to a newly configured instance. This was because `mysqld` had not completed the startup process before the next stage of the installation process. (Bug #28628)
- A `SELECT` in one connection could be blocked by `INSERT ... ON DUPLICATE KEY UPDATE` in another connection even when `low_priority_updates` is set. (Bug #28587)
- `mysql_upgrade` could run binaries dynamically linked against incorrect versions of shared libraries. (Bug #28560)
- The result from `CHAR()` was incorrectly assumed in some contexts to return a single-byte result. (Bug #28550)
- `mysqldump` reversed the event name and program name in one of its error messages. (Bug #28535)

- The parser confused user-defined function (UDF) and stored function creation for `CREATE FUNCTION` and required that there be a default database when creating UDFs, although there is no such requirement. (Bug #28318, Bug #29816)
- Fast-mutex locking was not thread-safe and optimization-safe on some platforms, which could cause program failures such as out-of-memory errors. (Bug #28284)
- The result of a comparison between `VARBINARY` and `BINARY` columns differed depending on whether the `VARBINARY` column was indexed. (Bug #28076)
- The metadata in some `MYSQL_FIELD` members could be incorrect when a temporary table was used to evaluate a query. (Bug #27990)
- Partition pruning was not used for queries having `<=` or `>=` conditions in the `WHERE` clause on a table using `TO_DAYS()` in the partitioning expression. (Bug #27927)
- `mysqlbinlog` produced incorrectly formatted `DATETIME` and `TIMESTAMP` values. (Bug #27894)
- Failure to log to the `general_log` or `slow_log` log tables were not logged to the error log at all or were logged incorrectly. (Bug #27858)
- An `ORDER BY` at the end of a `UNION` affected individual `SELECT` statements rather than the overall query result. (Bug #27848)
- `comp_err` created files with permissions such that they might be inaccessible during `make install` operations. (Bug #27789)
- With recent versions of `DBD::mysql`, `mysqlhotcopy` generated table names that were doubly qualified with the database name. (Bug #27694)
- The anonymous accounts were not being created during MySQL installation. (Bug #27692)
- Some `SHOW` statements and `INFORMATION_SCHEMA` queries could expose information not permitted by the user's access privileges. An implication of this change is that `SHOW TRIGGERS` and the `INFORMATION_SCHEMA.TRIGGERS` table require the `TRIGGER` privilege, not `SUPER`. (Bug #27629)
- `ALTER TABLE tbl_name ROW_FORMAT=format_type` did not cause the table to be rebuilt. (Bug #27610)
- A race condition between killing a statement and the thread executing the statement could lead to a situation such that the binary log contained an event indicating that the statement was killed, whereas the statement actually executed to completion. (Bug #27571)
- Some character mappings in the `ascii.xml` file were incorrect.

As a result of this bug fix, indexes must be rebuilt for columns that use the `ascii_general_ci` collation for columns that contain any of these characters: `'`, `[`, `\`, `]`, `~`. See [Checking Whether Tables or Indexes Must Be Rebuilt](#). (Bug #27562)
- Some queries using the `NAME_CONST()` function failed to return either a result or an error to the client, causing it to hang. This was due to the fact that there was no check to insure that both arguments to this function were constant expressions. (Bug #27545, Bug #32559)
- With the `read_only` system variable enabled, `CREATE DATABASE` and `DROP DATABASE` were permitted to users who did not have the `SUPER` privilege. (Bug #27440)
- For an event with an `ON COMPLETION` value of `PRESERVE`, an `ALTER EVENT` statement that specified no `ON COMPLETION` option caused the value to become `NOT PRESERVE`. (Bug #27407)

- MySQL failed to generate or retrieve an `AUTO_INCREMENT` primary key for `InnoDB` tables with user-defined partitioning. (Bug #27405)
- Changes to the `sql_mode` system variable were not tracked by `INSERT DELAYED`. (Bug #27358)
- A `SELECT` with more than 31 nested dependent subqueries returned an incorrect result. (Bug #27352)
- The `ExtractValue()` and `UpdateXML()` functions performed extremely slowly for large amounts of XML data (greater than 64 KB). These functions now execute approximately 2000 times faster than previously. (Bug #27287)
- On Windows, writes to the debug log were using `freopen()` instead of `fflush()`, resulting in slower performance. (Bug #27099)
- For a table that used different full-text parsers for different `FULLTEXT` indexes, `SHOW CREATE TABLE` displayed the first parser name for all of them. (Bug #27040)
- `STR_TO_DATE()` displayed an error message that referred to `STR_TO_TIME()`. (Bug #27014)
- The `mysql_insert_id()` C API function sometimes returned different results for `libmysqld` and `libmysqlclient`. (Bug #26921)
- Symbolic links on Windows could fail to work. (Bug #26811)
- `mysqld` sometimes miscalculated the number of digits required when storing a floating-point number in a `CHAR` column. This caused the value to be truncated, or (when using a debug build) caused the server to crash. (Bug #26788)

References: See also: Bug #12860.

- `LOAD DATA INFILE` ran very slowly when reading large files into partitioned tables. (Bug #26527)
- It makes no sense to attempt to use `ALTER TABLE ... ORDER BY` to order an `InnoDB` table if there is a user-defined clustered index, because rows are always ordered by the clustered index. Such attempts now are ignored and produce a warning.

Also, in some cases, `InnoDB` incorrectly used a secondary index when the clustered index would produce a faster scan. `EXPLAIN` output now indicates use of the clustered index (for tables that have one) as lines with a `type` value of `index`, a `key` value of `PRIMARY`, and without `Using index` in the `Extra` value. (Bug #26447)

References: See also: Bug #35850.

- Using `HANDLER` to open a table having a storage engine not supported by `HANDLER` properly returned an error, but also improperly prevented the table from being dropped by other connections. (Bug #25856)
- For a prepared statement `stmt`, changing the default database following `PREPARE stmt` but before `EXECUTE stmt` caused `stmt` to be recorded incorrectly in the binary log. (Bug #25843)
- `CREATE TABLE LIKE` did not work when the source table was an `INFORMATION_SCHEMA` table. (Bug #25629)
- Threads that were calculating the estimated number of records for a range scan did not respond to the `KILL` statement. That is, if a `range` join type is possible (even if not selected by the optimizer as a join type of choice and thus not shown by `EXPLAIN`), the query in the `statistics` state (shown by the `SHOW PROCESSLIST`) did not respond to the `KILL` statement. (Bug #25421)
- For `InnoDB` tables, `CREATE TABLE a AS SELECT * FROM A` failed. (Bug #25164)
- For `mysql --show-warnings`, warnings were in some cases not displayed. (Bug #25146)

- The `returns` column of the `mysql.proc` table was `CHAR(64)`, which is not long enough to store long data types such as `ENUM` types. The column has been changed to `LONGBLOB` and a warning is generated if truncation occurs when storing a row into the `proc` table. (Bug #24923)
- If the expected precision of an arithmetic expression exceeded the maximum precision supported by MySQL, the precision of the result was reduced by an unpredictable or arbitrary amount, rather than to the maximum precision. In some cases, exceeding the maximum supported precision could also lead to a crash of the server. (Bug #24907)
- For Vista installs, `MySQLInstanceConfig.exe` did not add the default MySQL port to the firewall exceptions. It now provides a check box that enables the user a choice of whether to do this. (Bug #24853)
- A `CREATE TRIGGER` statement could cause a deadlock or server crash if it referred to a table for which a table lock had been acquired with `LOCK TABLES`. (Bug #23713)
- For storage engines that do not redefine `handler::index_next_same()` and are capable of indexes, statements that include a `WHERE` clause might select incorrect data. (Bug #22351)
- The parser treated the `INTERVAL()` function incorrectly, leading to situations where syntax errors could result depending on which side of an arithmetic operator the function appeared. (Bug #22312)
- Using `FLUSH TABLES` in one connection while another connection is using `HANDLER` statements caused a server crash.

**Note**

This fix supersedes a fix made previously in MySQL 5.1.15 and reverted in MySQL 5.1.22.

(Bug #21587)

References: See also: Bug #29474.

- Entries in the general query log were truncated at 1000 characters. (Bug #21557)
- A memory leak occurred when `CREATE TEMPORARY TABLE ... SELECT` was invoked from a stored function that in turn was called from `CREATE TABLE ... SELECT`. (Bug #21136)
- It was possible to execute `CREATE TABLE t1 ... SELECT ... FROM t2` with the `CREATE` privilege for `t1` and `SELECT` privilege for `t2`, even in the absence of the `INSERT` privilege for `t1`. (Bug #20901)
- Worked around an `icc` problem with an incorrect machine instruction being generated in the context of software pre-fetching after a subroutine got in-lined. (Upgrading to `icc` 10.0.026 makes the workaround unnecessary.) (Bug #20803)
- If a column selected by a view referred to a stored function, the data type reported for the column in `INFORMATION_SCHEMA.COLUMNS` could be incorrect. (Bug #20550)
- The `mysql_change_user()` C API function changed the value of the `sql_big_selects` session variable. (Bug #20023)

References: See also: Bug #40363.

- Host names sometimes were treated as case sensitive in account-management statements (`CREATE USER`, `GRANT`, `REVOKE`, and so forth). (Bug #19828)
- Issuing an SQL `KILL` of the active connection caused an error on Mac OS X. (Bug #19723)

- The `readline` library has been updated to version 5.2. This addresses issues in the `mysql` client where history and editing within the client failed to work as expected. (Bug #18431)
- 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)
- `MySQLInstanceConfig.exe` failed to grant certain privileges to the `'root'@'%'` account. (Bug #17303)
- The `Aborted_clients` status variable was incremented twice if a client exited without calling `mysql_close()`. (Bug #16918)
- Use of `GRANT` statements with grant tables from an old version of MySQL could cause a server crash. (Bug #16470)
- Clients were ignoring the TCP/IP port number specified as the default port using the `--with-tcp-port` configuration option. (Bug #15327)
- Parameters of type `DATETIME` or `DATE` in stored procedures were silently converted to `VARBINARY`. (Bug #13675)
- Zero-padding of exponent values was not the same across platforms. (Bug #12860)
- Values of types `REAL ZEROFILL`, `DOUBLE ZEROFILL`, `FLOAT ZEROFILL`, were not zero-filled when converted to a character representation in the C prepared statement API. (Bug #11589)
- `mysql` stripped comments from statements sent to the server. Now the `--comments` or `--skip-comments` option can be used to control whether to retain or strip comments. The default is `--skip-comments`. (Bug #11230, Bug #26215)
- Grant table checks failed in `libmysqld`.
- Several buffer-size system variables were either being handled incorrectly for large values (for settings larger than 4GB, they were truncated to values less than 4GB without a warning), or were limited unnecessarily to 4GB even on 64-bit systems. The following changes were made:
 - For `key_buffer_size`, values larger than 4GB are permitted on 64-bit platforms.
 - For `join_buffer_size`, `sort_buffer_size`, and `myisam_sort_buffer_size`, values larger than 4GB are permitted on 64-bit platforms (except Windows, for which large values are truncated to 4GB with a warning).

In addition, settings for `read_buffer_size` and `read_rnd_buffer_size` are limited to 2GB on all platforms. Larger values are truncated to 2GB with a warning. (Bug #5731, Bug #29419, Bug #29446)

- Executing `DISABLE KEYS` and `ENABLE KEYS` on a nonempty table would cause the size of the index file for the table to grow considerable. This was because the `DISABLE KEYS` operation would only mark the existing index, without deleting the index blocks. The `ENABLE KEYS` operation would re-create the index, adding new blocks, while the previous index blocks would remain. Existing indexes are now dropped and recreated when the `ENABLE KEYS` statement is executed. (Bug #4692)

Changes in MySQL 5.1.22 (2007-09-24, Release Candidate)

- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

Functionality Added or Changed

- There is a new `innodb_autoinc_lock_mode` system variable to configure the locking behavior that InnoDB uses for generating auto-increment values. The default behavior now is slightly different from before, which involves a minor incompatibility for multiple-row inserts that specify an explicit value for the auto-increment column in some but not all rows. See [AUTO_INCREMENT Handling in InnoDB](#).

Bugs Fixed

- **MySQL Cluster; Replication:** (Replication): Multi-master replication setups did not handle `--log-slave-updates` correctly. (Bug #30017)
- **MySQL Cluster:** Backups of `TIMESTAMP` columns made with `ndb_restore` on a MySQL Cluster using data nodes hosts of one endian could not be used to restore the cluster's data to data node hosts of the other endian. (Bug #30134)
- **Replication:** Row-based replication from a pre-5.1.22 MySQL Server to a MySQL 5.1.22 was unstable due to an uninitialized variable. (Bug #31076)
- **Replication:** Operations that used the time zone replicated the time zone only for successful operations, but did not replicate the time zone for errors that need to know it. (Bug #29536)
- For an InnoDB table if a `SELECT` was ordered by the primary key and also had a `WHERE field = value` clause on a different field that was indexed, a `DESC` order instruction would be ignored. (Bug #31001)
- `mysql_install_db` could fail to find its message file. (Bug #30678)
- Memory corruption occurred for some queries with a top-level `OR` operation in the `WHERE` condition if they contained equality predicates and other sargable predicates in disjunctive parts of the condition. (Bug #30396)
- `CONNECTION_ID()` always returned 0 for the embedded server (`libmysqld`). (Bug #30389)
- The server created temporary tables for filesort operations in the working directory, not in the directory specified by the `tmpdir` system variable. (Bug #30287)
- Using `KILL QUERY` or `KILL CONNECTION` to kill a `SELECT` statement caused a server crash if the query cache was enabled. (Bug #30201)
- `mysqldump` from the MySQL 5.1.21 distribution could not be used to create a dump from a MySQL 5.1.20 or older server. (Bug #30123)
- Under some circumstances, a UDF initialization function could be passed incorrect argument lengths. (Bug #29804)
- When using a combination of `HANDLER... READ` and `DELETE` on a table, MySQL continued to open new copies of the table every time, leading to an exhaustion of file descriptors. (Bug #29474)

References: This issue is a regression of: Bug #21587.

- The `mysql_list_fields()` C API function incorrectly set `MYSQL_FIELD::decimals` for some view columns. (Bug #29306)
- Tables using the InnoDB storage engine incremented `AUTO_INCREMENT` values incorrectly with `ON DUPLICATE KEY UPDATE`. (Bug #28781)
- Nonrange queries of the form `SELECT ... FROM ... WHERE keypart1=constant, ..., keypartN=constant ORDER BY ... FOR UPDATE` sometimes were unnecessarily blocked waiting for a lock if another transaction was using `SELECT ... FOR UPDATE` on the same table. (Bug #28570)

- On Windows, symbols for yaSSL and taocrypt were missing from `mysqlclient.lib`, resulting in unresolved symbol errors for clients linked against that library. (Bug #27861)
- Read lock requests that were blocked by a pending write lock request were not permitted to proceed if the statement requesting the write lock was killed. (Bug #21281)

Changes in MySQL 5.1.21 (2007-08-16)

This is a new Beta development release, fixing recently discovered bugs.



Note

This Beta 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.



Note

Subsequent to release, it was discovered that on some platforms, `mysql_install_db` could fail to find its message file, resulting in error messages of the following form:

```
shell> mysql_install_db
Installing MySQL system tables...
070830  9:33:24 [ERROR] Can't find messagefile 'path/share/english/errmsg.sys'
070830  9:33:24 [ERROR] Aborting
```

To deal with this problem, specify a `--language` option to specify the proper path name to the language file directory. For example:

```
shell> mysql_install_db --language=/path/to/share/english/
```

This problem is corrected in MySQL 5.1.22.


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 Enterprise* (a commercial MySQL offering). For more details, please see (<http://www.mysql.com/products/enterprise>).

- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

Functionality Added or Changed

- **Incompatible Change:** In MySQL 5.1.6, when log tables were implemented, the default log destination for the general query and slow query log was `TABLE`. This default has been changed to `FILE`, which is compatible with MySQL 5.0, but incompatible with earlier releases of MySQL 5.1 from 5.1.6 to 5.1.20. If you are upgrading from MySQL 5.0 to 5.1.21 or higher, no logging option changes should be necessary. However, if you are upgrading from 5.1.6 through 5.1.20 to 5.1.21 or higher and were using `TABLE` logging, use the `--log-output=TABLE` option explicitly to preserve your server's table-logging behavior.

A further fix for this issue was made in MySQL 5.1.23. (Bug #29993)

- **Incompatible Change:** The `innodb_log_arch_dir` system variable (which has been deprecated since MySQL 5.0.24) has been removed and should no longer be used.
- **Incompatible Change:** On Windows only, the `mysqld-nt` has been removed from this release and all future releases. The `mysqld` server now includes named-pipe support as standard, and you do not have to use the `mysqld-nt` version to enable named-pipe support.
- **Important Change:** The default `mysqld_safe` logging behavior now is `--skip-syslog` rather than `--syslog`, which is compatible with the default behavior of writing an error log file for releases prior to 5.1.20.
- **Replication:** The SQL thread on a slave now is always permitted to enter `InnoDB` even if this would exceed the limit imposed by the `innodb_thread_concurrency` system variable. In cases of high load on the slave server (when `innodb_thread_concurrency` is reached), this change helps the slave stay more up to date with the master; in the previous behavior, the SQL thread was competing for resources with all client threads active on the slave server. (Bug #25078)
- **Replication:** Replication between master and slaves now supports different column numbers within a table on both master and slave. The rules for replication where the table definitions are different has also changed. This supersedes the functionality for replication from the master table to a slave table with more columns that was added in MySQL 5.1.12. For more information, see [Replication with Differing Table Definitions on Master and Slave](#).
- Several programs now accept `--debug-check` and `--debug-info` options: `mysql`, `mysqladmin`, `mysqlbinlog`, `mysqlcheck`, `mysqldump`, `mysqlimport`, `mysqlshow`, `mysqlslap`, `mysqltest`, `mysql_upgrade`. (Note: `mysql`, `mysqladmin`, `mysqlcheck`, `mysqldump`, `mysqlimport`, `mysqlshow`, and `mysqltest` already accepted `--debug-info`.) `--debug-check` prints debugging information at program exit. `--debug-info` is similar but also prints memory and CPU usage statistics. This patch also corrects a problem for `mysql` that `--debug-info` did not display statistics at exit time. (Bug #30127)
- The `--syslog` option that was introduced in 5.1.20 for `mysqld_safe` (to send error output to `syslog`) did not work correctly: Error output was buffered and not logged immediately. This has been corrected. In addition, some feature changes were made:
 -  **Important**
The default `mysqld_safe` logging behavior now is `--skip-syslog` rather than `--syslog`, which is compatible with the default behavior of writing an error log file for releases prior to 5.1.20.
 - A new option, `--syslog-tag=tag`, modifies the default tags written by `mysqld_safe` and `mysqld` to syslog to be `mysqld_safe-tag` and `mysqld-tag` rather than the default tags of `mysqld_safe` and `mysqld`.
(Bug #29992)
- Transaction support in the `FEDERATED` storage engine has been disabled due to issues with multiple active transactions and sessions on the same `FEDERATED` table. (Bug #29875)
- Previously, prepared statements processed using `PREPARE` and `EXECUTE` were not subject to caching in the query cache if they contained any `?` parameter markers. This limitation has been lifted. (Bug #29318)
- It is now possible to set `long_query_time` in microseconds or to 0. Setting this value to 0 causes all queries to be recorded in the slow query log.

Currently, fractional values can be used only when logging to files. We plan to provide this functionality for logging to tables when time-related data types are enhanced to support microsecond resolution. (Bug #25412)

- [INFORMATION_SCHEMA](#) implementation changes were made that optimize certain types of queries for [INFORMATION_SCHEMA](#) tables so that they execute more quickly. [INFORMATION_SCHEMA Optimization](#), provides guidelines on how to take advantage of these optimizations by writing queries that minimize the need for the server to access the file system to obtain the information contained in [INFORMATION_SCHEMA](#) tables. By writing queries that enable the server to avoid directory scans or opening table files, you will obtain better performance. (Bug #19588)
- Log table locking was redesigned, eliminating several lock-related problems:
 - Truncating `mysql.slow_log` in a stored procedure after use of a cursor caused the thread to lock.
 - Flushing a log table resulted in unnecessary warnings.
 - The server would hang when performing concurrent `ALTER TABLE` or `TRUNCATE TABLE` statements against the log tables.
 - Changing the value of the `general_log` system variable while a global read lock was in place resulted in deadlock.

The changes provide better-defined interface characteristics. See [Selecting General Query and Slow Query Log Output Destinations](#). (Bug #17876, Bug #23044, Bug #25422, Bug #29129)

- Added the `--commit`, `--detach`, `--post-system`, and `--pre-system` options for `mysqldump`.
- A new option, `--syslog-tag=tag`, modifies the default tags written by `mysqld_safe` and `mysqld` to syslog to be `mysqld_safe-tag` and `mysqld-tag` rather than the default tags of `mysqld_safe` and `mysqld`.
- Two options relating to slow query logging have been added for `mysqld`. `--log-slow-slave-statements` causes slow statements executed by a replication slave to be written to the slow query log; `min_examined_row_limit` can be used to cause queries which examine fewer than the stated number of rows not to be logged.

Bugs Fixed

- **Incompatible Change:** Failure to consider collation when comparing space characters could result in incorrect index entry order, leading to incorrect comparisons, inability to find some index values, misordered index entries, misordered `ORDER BY` results, or tables that `CHECK TABLE` reports as having corrupt indexes.

As a result of this bug fix, indexes must be rebuilt for columns that use any of these character sets: `eucjpms`, `euc_kr`, `gb2312`, `latin7`, `macce`, `ujis`. See [Checking Whether Tables or Indexes Must Be Rebuilt](#). (Bug #29461)

- **Incompatible Change:** Several issues were identified for stored programs (stored procedures and functions, triggers, and events) and views containing non-ASCII symbols. These issues involved conversion errors due to incomplete character set information when translating these objects to and from stored format, such as:
 - Parsing the original object definition so that it can be stored.
 - Compiling the stored definition into executable form when the object is invoked.

- Retrieval of object definitions from `INFORMATION_SCHEMA` tables.
- Displaying the object definition in `SHOW` statements. This issue also affected `mysqldump`, which uses `SHOW`.

The fix for the problems is to store character set information from the object creation context so that this information is available when the object needs to be used later. The context includes the client character set, the connection character set and collation, and the collation of the database with which the object is associated.

As a result of the patch, several tables have new columns:

- In the `mysql` database, the `proc` and `event` tables now have these columns: `character_set_client`, `collation_connection`, `db_collation`, `body_utf8`.
- In `INFORMATION_SCHEMA`, the `VIEWS` table now has these columns: `CHARACTER_SET_CLIENT`, `COLLATION_CONNECTION`. The `ROUTINES`, `TRIGGERS`, and `EVENTS` tables now have these columns: `CHARACTER_SET_CLIENT`, `COLLATION_CONNECTION`, `DATABASE_COLLATION`.

These columns store the session values of the `character_set_client` and `collation_connection` system variables, and the collation of the database with which the object is associated. The values are those in effect at object creation time. (The saved database collation is not the value of the `collation_database` system variable, which applies to the default database; the database that contains the object is not necessarily the default database.)

Several `SHOW` statements now display additional columns corresponding to the new table columns. These statements are: `SHOW CREATE EVENT`, `SHOW CREATE FUNCTION`, `SHOW CREATE PROCEDURE`, `SHOW CREATE VIEW`, `SHOW EVENTS`, `SHOW FUNCTION STATUS`, `SHOW PROCEDURE STATUS`, `SHOW TRIGGERS`.

A new statement, `SHOW CREATE TRIGGER` is introduced and is used by `mysqldump` for producing `CREATE TRIGGER` statements.

`mysqldump` output now can include `ALTER DATABASE` statements that change the database collation. These may be used when dumping stored programs to preserve their character encodings. To reload a dump file containing such statements, the `ALTER` privilege for the affected database is required.

Subsequent to the patch just described, it was discovered that the patch broke `mysql_upgrade`; this has been corrected.



Important

The fixes for the problems just describe affect *all* existing stored programs and views. (For example, you will see warnings about “no creation context.”) To avoid warnings from the server about the use of old definitions from any release prior to 5.1.21, you should dump stored programs and views with `mysqldump` after upgrading to 5.1.21, and then reload them to recreate them with new definitions. Invoke `mysqldump` with a `--default-character-set` option that names the non-ASCII character set that was used for the definitions when the objects were originally created, and the `--routines`, `--events`, and `--triggers` options to dump stored program definitions. For more information about upgrading stored programs, see [Changes Affecting Upgrades to 5.1](#).

(Bug #25221, Bug #21249, Bug #30027, Bug #16291, Bug #11986, Bug #25212, Bug #19443, Bug #30029)

- **MySQL Cluster; Replication:** (Replication): Inconsistencies could occur between the master and the slave when replicating Disk Data tables. (Bug #19259, Bug #19227)
- **MySQL Cluster:** `DELETE FROM table WHERE primary_key IN (value_list)`, where the `value_list` contained more than one value, called from an `AFTER DELETE` trigger on an `NDB` table, caused `mysqld` to crash. (Bug #30337)
- **MySQL Cluster:** When restarting a data node, queries could hang during that node's start phase 5, and continue only after the node had entered phase 6. (Bug #29364)
- **MySQL Cluster:** Replica redo logs were inconsistently handled during a system restart. (Bug #29354)
- **MySQL Cluster:** When a node failed to respond to a `COPY_GCI` signal as part of a global checkpoint, the master node was killed instead of the node that actually failed. (Bug #29331)
- **MySQL Cluster:** An invalid comparison made during `REDO` validation that could lead to an `Error while reading REDO log` condition. (Bug #29118)
- **MySQL Cluster:** The wrong data pages were sometimes invalidated following a global checkpoint. (Bug #29067)
- **MySQL Cluster:** If at least 2 files were involved in `REDO` invalidation, then file 0 of page 0 was not updated and so pointed to an invalid part of the redo log. (Bug #29057)
- **MySQL Cluster:** If a storage engine has its own logging capability, then any statement using both this engine and some other engine not having its own logging could not be correctly logged, due to the fact that entries from one engine could be logged before entries from the other engine were. This did not generate any error messages when it occurred.

Now, if multiple storage engines are used in a statement and at least one of them has its own logging capability, then an error message is generated and the statement is not executed.

**Note**

Currently, the only storage engine to have its own logging capability is `NDBCLUSTER`.

(Bug #28722)

- **MySQL Cluster:** Warnings and errors generated by `ndb_config --config-file=file` were sent to `stdout`, rather than to `stderr`. (Bug #25941)
- **MySQL Cluster:** When a cluster backup was terminated using the `ABORT BACKUP` command in the management client, a misleading error message `Backup aborted by application: Permanent error: Internal error` was returned. The error message returned in such cases now reads `Backup aborted by user request`. (Bug #21052)
- **MySQL Cluster:** Large file support did not work in AIX server binaries. (Bug #10776)
- **Replication:** The thread ID was not reset properly after execution of `mysql_change_user()`, which could cause replication failure when replicating temporary tables. (Bug #29734)
- **Replication:** Storage engine error conditions in row-based replication were not correctly reported to the user. (Bug #29570)
- **Replication:** `INSERT DELAYED` statements on a master server are replicated as non-`DELAYED` inserts on slaves (which is normal, to preserve serialization), but the inserts on the slave did not use concurrent inserts. Now `INSERT DELAYED` on a slave is converted to a concurrent insert when possible, and to a normal insert otherwise. (Bug #29152)

- **Replication:** An error that happened inside `INSERT`, `UPDATE`, or `DELETE` statements performed from within a stored function or trigger could cause inconsistency between master and slave servers. (Bug #27417)
- **Replication:** Slave servers could incorrectly interpret an out-of-memory error from the master and reconnect using the wrong binary log position. (Bug #24192)
- **Replication:** Using the `READ COMMITTED` transaction isolation level caused mixed and statement-based replication to fail. (Bug #23051)
- **Disk Data:** Performing Disk Data schema operations during a node restart could cause forced shutdowns of other data nodes. (Bug #29501)
- **Disk Data:** When dropping a page, the stack's bottom entry could sometime be left “cold” rather than “hot”, violating the rules for stack pruning. (Bug #29176)
- **Disk Data:** Disk data meta-information that existed in `ndbd` might not be visible to `mysqld`. (Bug #28720)
- **Disk Data:** The number of free extents was incorrectly reported for some tablespaces. (Bug #28642)
- **Cluster Replication:** When executing a statement where `binlog_format = statement`, the result of the statement was logged both as a statement and as rows. (Bug #29222)
- **Cluster Replication:** `mysqld` would segfault on startup when the `NDB` storage engine was enabled and the default character set was a strictly multibyte character set such as UCS2.

This issue does *not* apply to character sets that can contain single-byte characters in addition to multibyte characters such as UTF-8.

Additional issues remain with regard to the use of multibyte character sets in MySQL Cluster Replication; see [Known Issues in MySQL Cluster Replication](#), for more information. (Bug #27404)

- Prepared statements containing `CONNECTION_ID()` could be written improperly to the binary log. (Bug #30200)
- Use of local variables with non-ASCII names in stored procedures crashed the server. (Bug #30120)
- On Windows, client libraries lacked symbols required for linking. (Bug #30118)
- `--myisam-recover=''` (empty option value) did not disable `MyISAM` recovery. (Bug #30088)
- For the `SHOW TABLE TYPES` statement, the server sent incorrect output to clients, possibly causing them to crash. (Bug #30036)
- The `IS_UPDATABLE` column in the `INFORMATION_SCHEMA.VIEWS` table was not always set correctly. (Bug #30020)
- `SHOW` statements were being written to the slow query log that should not have been. (Bug #30000)
- `REPAIR TABLE ... USE_FRM` could corrupt tables. (Bug #29980)
- For `MyISAM` tables on Windows, `INSERT`, `DELETE`, or `UPDATE` followed by `ALTER TABLE` within `LOCK TABLES` could cause table corruption. (Bug #29957)
- `LOCK TABLES` did not pre-lock tables used in triggers of the locked tables. Unexpected locking behavior and statement failures similar to `failed: 1100: Table 'xx' was not locked with LOCK TABLES` could result. (Bug #29929)

- `INSERT ... VALUES(CONNECTION_ID(), ...)` statements were written to the binary log in such a way that they could not be properly restored. (Bug #29928)
- Adding `DISTINCT` could cause incorrect rows to appear in a query result. (Bug #29911)
- On Windows, the `CMake` build process did not produce the embedded server library or related binaries. (Bug #29903)
- Using the `DATE()` function in a `WHERE` clause did not return any records after encountering `NULL`. However, using `TRIM()` or `CAST()` produced the correct results. (Bug #29898)
- `SESSION_USER()` returned garbage data (rather than the correct value of the empty string) when executed by a slave SQL thread. (Bug #29878)
- Very long prepared statements in stored procedures could cause a server crash. (Bug #29856)
- If query execution involved a temporary table, `GROUP_CONCAT()` could return a result with an incorrect character set. (Bug #29850)
- If one thread was performing concurrent inserts, other threads reading from the same table using equality key searches could see the index values for new rows before the data values had been written, leading to reports of table corruption. (Bug #29838)
- Repeatedly accessing a view in a stored procedure (for example, in a loop) caused a small amount of memory to be allocated per access. Although this memory is deallocated on disconnect, it could be a problem for a long running stored procedures that make repeated access of views. (Bug #29834)
- `mysqldump` produced output that incorrectly discarded the `NO_AUTO_VALUE_ON_ZERO` value of the `sql_mode` variable after dumping triggers. (Bug #29788)
- An assertion failure occurred within yaSSL for very long keys. (Bug #29784)

References: See also: Bug #53463.

- For `MEMORY` tables, the `index_merge` union access method could return incorrect results. (Bug #29740)
- Comparison of `TIME` values using the `BETWEEN` operator led to string comparison, producing incorrect results in some cases. Now the values are compared as integers. (Bug #29739)
- For a table with a `DATE` column `date_col` such that selecting rows with `WHERE date_col = 'date_val 00:00:00'` yielded a nonempty result, adding `GROUP BY date_col` caused the result to be empty. (Bug #29729)
- In some cases, `INSERT INTO ... SELECT ... GROUP BY` could insert rows even if the `SELECT` by itself produced an empty result. (Bug #29717)
- Single-row inserts could report a row count greater than one. (Bug #29692)
- For the embedded server, the `mysql_stmt_store_result()` C API function caused a memory leak for empty result sets. (Bug #29687)
- `EXPLAIN` produced `Impossible where` for statements of the form `SELECT ... FROM t WHERE c=0`, where `c` was an `ENUM` column defined as a primary key. (Bug #29661)
- On Windows, `ALTER TABLE` hung if records were locked in share mode by a long-running transaction. (Bug #29644)
- `mysqld_safe` produced error messages and did not create the error log file under some circumstances. (Bug #29634)

- On 64-bit platforms, the filesort code (for queries with `GROUP BY` or `ORDER BY`) could crash due to an incorrect pointer size. (Bug #29610)
- A left join between two views could produce incorrect results. (Bug #29604)
- Certain statements with unions, subqueries, and joins could result in huge memory consumption. (Bug #29582)
- Clients using SSL could hang the server. (Bug #29579)
- A slave running with `--log-slave-updates` failed to write `INSERT DELAY IGNORE` statements to its binary log, resulting in different binary log contents on the master and slave. (Bug #29571)
- An incorrect result was returned when comparing string values that were converted to `TIME` values with `CAST()`. (Bug #29555)
- `gcov` coverage-testing information was not written if the server crashed. (Bug #29543)
- In the `ascii` character set, conversion of DEL (`0x7F`) to Unicode incorrectly resulted in QUESTION MARK (`0x3F`) rather than DEL. (Bug #29499)
- A field packet with `NULL` fields caused a `libmysqlclient` crash. (Bug #29494)
- On Windows, the `mysql` client died if the user entered a statement and Return after entering **Control+C**. (Bug #29469)
- The full-text parser could enter an infinite loop if it encountered an illegal multibyte sequence or a sequence that has no mapping to Unicode. (Bug #29464)
- Searching a `FULLTEXT` index for a word with the boolean mode truncation operator could cause an infinite loop. (Bug #29445)
- Corrupt data resulted from use of `SELECT ... INTO OUTFILE 'file_name' FIELDS ENCLOSED BY 'c'`, where `c` is a digit or minus sign, followed by `LOAD DATA INFILE 'file_name' FIELDS ENCLOSED BY 'c'`. (Bug #29442)
- Killing an `INSERT DELAYED` thread caused a server crash. (Bug #29431)
- Use of `SHOW BINLOG EVENTS` for a nonexistent log file followed by `PURGE BINARY LOGS` caused a server crash. (Bug #29420)
- Assertion failure could occur for grouping queries that employed `DECIMAL` user variables with assignments to them. (Bug #29417)
- For `CAST(expr AS DECIMAL(M,D))`, the limits of 65 and 30 on the precision (`M`) and scale (`D`) were not enforced. (Bug #29415)
- Deleting from a `CSV` table could corrupt it. (Bug #29411)
- Results for a select query that aliases the column names against a view could duplicate one column while omitting another. This bug could occur for a query over a multiple-table view that includes an `ORDER BY` clause in its definition. (Bug #29392)
- `mysqldump` created a stray file when a given a too-long file name argument. (Bug #29361)
- The special “zero” `ENUM` value was coerced to the normal empty string `ENUM` value during a column-to-column copy. This affected `CREATE ... SELECT` statements and `SELECT` statements with aggregate functions on `ENUM` columns in the `GROUP BY` clause. (Bug #29360)

- Inserting a negative number into a `CSV` table could corrupt it. (Bug #29353)
- Optimization of queries with `DETERMINISTIC` stored functions in the `WHERE` clause was ineffective: A sequential scan was always used. (Bug #29338)
- `MyISAM` corruption could occur with the `cp932_japanese_ci` collation for the `cp932` character set due to incorrect comparison for trailing space. (Bug #29333)
- For updates to `InnoDB` tables, a `TIMESTAMP` column with the `ON UPDATE CURRENT_TIMESTAMP` attribute could be updated even when no values actually changed. (Bug #29310)
- `FULLTEXT` indexes could be corrupted by certain `gbk` characters. (Bug #29299)
- `SELECT ... INTO OUTFILE` followed by `LOAD DATA` could result in garbled characters when the `FIELDS ENCLOSED BY` clause named a delimiter of `'0'`, `'b'`, `'n'`, `'r'`, `'t'`, `'N'`, or `'Z'` due to an interaction of character encoding and doubling for data values containing the enclosed-by character. (Bug #29294)
- Sort order of the collation wasn't used when comparing trailing spaces. This could lead to incorrect comparison results, incorrectly created indexes, or incorrect result set order for queries that include an `ORDER BY` clause. (Bug #29261)
- `CHECK TABLE` could erroneously report table corruption for a `CSV` table if multiple threads were modifying the table at the same time. (Bug #29253)
- Many threads accessing a `CSV` table simultaneously could cause an assertion failure. (Bug #29252)
- If an `ENUM` column contained `' '` as one of its members (represented with numeric value greater than 0), and the column contained error values (represented as 0 and displayed as `' '`), using `ALTER TABLE` to modify the column definition caused the 0 values to be given the numeric value of the nonzero `' '` member. (Bug #29251)
- Calling `mysql_options()` after `mysql_real_connect()` could cause clients to crash. (Bug #29247)
- `CHECK TABLE` for `ARCHIVE` tables could falsely report table corruption or cause a server crash. (Bug #29207)
- Mixing binary and `utf8` columns in a union caused field lengths to be calculated incorrectly, resulting in truncation. (Bug #29205)
- `AsText()` could fail with a buffer overrun. (Bug #29166)
- Under some circumstances, a `SELECT ... FROM mysql.event` could cause the server to crash. (Bug #29156)
- `InnoDB` refused to start on some versions of FreeBSD with LinuxThreads. This is fixed by enabling file locking on FreeBSD. (Bug #29155)
- `LOCK TABLES` was not atomic when more than one `InnoDB` tables were locked. (Bug #29154)
- `mysqld` failed to exit during shutdown. (Bug #29133)
- A network structure was initialized incorrectly, leading to embedded server crashes. (Bug #29117)
- An assertion failure occurred if a query contained a conjunctive predicate of the form `view_column = constant` in the `WHERE` clause and the `GROUP BY` clause contained a reference to a different view column. The fix also enables application of an optimization that was being skipped if a query contained a conjunctive predicate of the form `view_column = constant` in the `WHERE` clause and the `GROUP BY` clause contained a reference to the same view column. (Bug #29104)

- A maximum of 4TB `InnoDB` free space was reported by `SHOW TABLE STATUS`, which is incorrect on systems with more than 4TB space. (Bug #29097)
- If an `INSERT INTO ... SELECT` statement inserted into the same table that the `SELECT` retrieved from, and the `SELECT` included `ORDER BY` and `LIMIT` clauses, different data was inserted than the data produced by the `SELECT` executed by itself. (Bug #29095)
- Queries that performed a lookup into a `BINARY` index containing key values ending with spaces caused an assertion failure for debug builds and incorrect results for nondebug builds. (Bug #29087)
- The semantics of `BIGINT` depended on platform-specific characteristics. (Bug #29079)
- A byte-order issue in writing a spatial index to disk caused bad index files on some systems. (Bug #29070)
- Creation of a legal stored procedure could fail if no default database had been selected. (Bug #29050)
- `REPLACE`, `INSERT IGNORE`, and `UPDATE IGNORE` did not work for `FEDERATED` tables. (Bug #29019)
- Inserting into `InnoDB` tables and executing `RESET MASTER` in multiple threads cause assertion failure in debug server binaries. (Bug #28983)
- Updates to a `CSV` table could cause a server crash or update the table with incorrect values. (Bug #28971)
- For a `ucs2` column, `GROUP_CONCAT()` did not convert separators to the result character set before inserting them, producing a result containing a mixture of two different character sets. (Bug #28925)
- Dropping the definer of an active event caused the server to crash. (Bug #28924)
- For a join with `GROUP BY` or `ORDER BY` and a view reference in the `FROM` list, the query metadata erroneously showed empty table aliases and database names for the view columns. (Bug #28898)
- Creating an event using `ON SCHEDULE AT CURRENT_TIMESTAMP + INTERVAL ...` could in some cases cause `mysqld` to crash. (Bug #28881)
- Coercion of ASCII values to character sets that are a superset of ASCII sometimes was not done, resulting in `illegal mix of collations` errors. These cases now are resolved using repertoire, a new string expression attribute (see [String Repertoire](#)). (Bug #28875)
- Executing `ALTER EVENT` on an event whose definer's event creation privileges had been revoked cause the server to crash. (Bug #28873)
- `ALTER VIEW` is not supported as a prepared statement but was not being rejected. `ALTER VIEW` is now prohibited as a prepared statement or when called within stored routines. (Bug #28846)
- In strict SQL mode, errors silently stopped the SQL thread even for errors named using the `--slave-skip-errors` option. (Bug #28839)
- Fast `ALTER TABLE` (that works without rebuilding the table) acquired duplicate locks in the storage engine. In `MyISAM`, if `ALTER TABLE` was issued under `LOCK TABLE`, it caused all data inserted after `LOCK TABLE` to disappear. (Bug #28838)
- Runtime changes to the `log_queries_not_using_indexes` system variable were ignored. (Bug #28808)
- Selecting a column not present in the selected-from table caused an extra error to be produced by `SHOW ERRORS`. (Bug #28677)

- Creating an event to be executed at a time close to the end of the permitted range (2038-01-19 03:14:07 UTC) would cause the server to crash. (Bug #28641)
- For a statement of the form `CREATE t1 SELECT integer_constant`, the server created the column using the `DECIMAL` data type for large negative values that are within the range of `BIGINT`. (Bug #28625)
- Starting the server with an `innodb_force_recovery` value of 4 did not work. (Bug #28604)
- For `InnoDB` tables, MySQL unnecessarily sorted records in certain cases when the records were retrieved by `InnoDB` in the proper order already. (Bug #28591)
- `mysql_install_db` could fail to find script files that it needs. (Bug #28585)
- If a stored procedure was created and invoked prior to selecting a default database with `USE`, a `No database selected` error occurred. (Bug #28551)
- On Mac OS X, shared-library installation path names were incorrect. (Bug #28544)
- Using the `--skip-add-drop-table` option with `mysqldump` generated incorrect SQL if the database included any views. The recreation of views requires the creation and removal of temporary tables. This option suppressed the removal of those temporary tables. The same applied to `--compact` since this option also invokes `--skip-add-drop-table`. (Bug #28524)
- `mysqlbinlog --hexdump` generated incorrect output due to omission of the “#” comment character for some comment lines. (Bug #28293)
- `InnoDB` could crash if the server was shut down while `innodb_table_monitor` was running. (Bug #28254)
- A race condition in the interaction between `MyISAM` and the query cache code caused the query cache not to invalidate itself for concurrently inserted data. (Bug #28249)
- A duplicate-key error message could display an incorrect key value when not all columns of the key were used to select rows for update. (Bug #28158)
- Indexing column prefixes in `InnoDB` tables could cause table corruption. (Bug #28138)
- Index creation could fail due to truncation of key values to the maximum key length rather than to a multiple of the maximum character length. (Bug #28125)
- Instance Manager had a race condition when it received a shutdown request while a guarded `mysqld` instance was starting such that it could fail to stop the `mysqld` instance. (Bug #28030)
- `SELECT ... FOR UPDATE` with partitioned tables could cause a server crash. (Bug #28026)
- On Windows, Instance Manager would crash if an instance object failed to initialize during startup. This could happen if an incorrect `mysqld` path was supplied in the configuration file. (Bug #28012)
- The `LOCATE()` function returned `NULL` if any of its arguments evaluated to `NULL`. Likewise, the predicate, `LOCATE(str,NULL) IS NULL`, erroneously evaluated to `FALSE`. (Bug #27932)
- Dropping a user-defined function could cause a server crash if the function was still in use by another thread. (Bug #27564)
- For some event-creation problems, the server displayed messages that implied the problems were errors when they were only warnings. (Bug #27406)
- Unsafe aliasing in the source caused a client library crash when compiled with `gcc 4` at high optimization levels. (Bug #27383)

- Index-based range reads could fail for comparisons that involved contraction characters (such as `ch` in Czech or `ll` in Spanish). (Bug #27345)
- Aggregations in subqueries that refer to outer query columns were not always correctly referenced to the proper outer query. (Bug #27333)
- Error returns from the `time()` system call were ignored. (Bug #27198)
- Phantom reads could occur under `InnoDB SERIALIZABLE` isolation level. (Bug #27197)
- The `SUBSTRING()` function returned the entire string instead of an empty string when it was called from a stored procedure and when the length parameter was specified by a variable with the value “0”. (Bug #27130)
- Some functions when used in partitioning expressions could cause `mysqld` to crash. (Bug #27084)
- The server acquired a global mutex for temporary tables, although such tables are thread-specific. This affected performance by blocking other threads. (Bug #27062)
- `FEDERATED` tables had an artificially low maximum of key length. (Bug #26909)
- Updates to rows in a partitioned table could update the wrong column. (Bug #26827)
- Index creation could corrupt the table definition in the `.frm` file: 1) A table with the maximum number of key segments and maximum length key name would have a corrupted `.frm` file, due to incorrect calculation of the total key length. 2) `MyISAM` would reject a table with the maximum number of keys and the maximum number of key segments in all keys. (It would permit one less than this total maximum.) Now `MyISAM` accepts a table defined with the maximum. (Bug #26642)
- The Windows implementation of `pthread_join()` was incorrect and could cause crashes. (Bug #26564)
- After the first read of a `TEMPORARY` table, `CHECK TABLE` could report the table as being corrupt. (Bug #26325)
- If an operation had an `InnoDB` table, and two triggers, `AFTER UPDATE` and `AFTER INSERT`, competing for different resources (such as two distinct `MyISAM` tables), the triggers were unable to execute concurrently. In addition, `INSERT` and `UPDATE` statements for the `InnoDB` table were unable to run concurrently. (Bug #26141)
- A number of unsupported constructs—including prohibited constructs, the `UCASE()` function, and nested function calls—were permitted in partitioning expressions. (Bug #26082, Bug #18198, Bug #29308)
- `ALTER DATABASE` did not require at least one option. (Bug #25859)
- The index merge union access algorithm could produce incorrect results with `InnoDB` tables. The problem could also occur for queries that used `DISTINCT`. (Bug #25798)
- When using a `FEDERATED` table, the value of `LAST_INSERT_ID()` would not correctly update the C API interface, which would affect the autogenerated ID returned both through the C API and the MySQL protocol, affecting Connectors that used the protocol or C API. (Bug #25714)
- The server was blocked from opening other tables while the `FEDERATED` engine was attempting to open a remote table. Now the server does not check the correctness of a `FEDERATED` table at `CREATE TABLE` time, but waits until the table actually is accessed. (Bug #25679)
- Under ActiveState Perl, `mysql-test-run.pl` could kill itself when attempting to kill other processes. (Bug #25657)
- Several `InnoDB` assertion failures were corrected. (Bug #25645)

- A query with `DISTINCT` in the select list to which the loose-scan optimization for grouping queries was applied returned an incorrect result set when the query was used with the `SQL_BIG_RESULT` option. (Bug #25602)
- For a multiple-row insert into a `FEDERATED` table that refers to a remote transactional table, if the insert failed for a row due to constraint failure, the remote table would contain a partial commit (the rows preceding the failed one) instead of rolling back the statement completely. This occurred because the rows were treated as individual inserts.

Now `FEDERATED` performs bulk-insert handling such that multiple rows are sent to the remote table in a batch. This provides a performance improvement and enables the remote table to perform statement rollback properly should an error occur. This capability has the following limitations:

- The size of the insert cannot exceed the maximum packet size between servers. If the insert exceeds this size, it is broken into multiple packets and the rollback problem can occur.
- Bulk-insert handling does not occur for `INSERT ... ON DUPLICATE KEY UPDATE`.

(Bug #25513)

- The `FEDERATED` storage engine failed silently for `INSERT ... ON DUPLICATE KEY UPDATE` if a duplicate key violation occurred. `FEDERATED` does not support `ON DUPLICATE KEY UPDATE`, so now it correctly returns an `ER_DUP_KEY` error if a duplicate key violation occurs. (Bug #25511)
- In a stored function or trigger, when `InnoDB` detected deadlock, it attempted rollback and displayed an incorrect error message (`Explicit or implicit commit is not permitted in stored function or trigger`). Now `InnoDB` returns an error under these conditions and does not attempt rollback. Rollback is handled outside of `InnoDB` above the function/trigger level. (Bug #24989)
- Dropping a temporary `InnoDB` table that had been locked with `LOCK TABLES` caused a server crash. (Bug #24918)
- On Windows, executables did not include Vista manifests. (Bug #24732)

References: See also: Bug #22563.

- If MySQL/`InnoDB` crashed very quickly after starting up, it would not force a checkpoint. In this case, `InnoDB` would skip crash recovery at next startup, and the database would become corrupt. Now, if the redo log scan at `InnoDB` startup goes past the last checkpoint, crash recovery is forced. (Bug #23710)
- `SHOW INNODB STATUS` caused an assertion failure under high load. (Bug #22819)
- `SHOW BINLOG EVENTS` displayed incorrect values of `End_log_pos` for events associated with transactional storage engines. (Bug #22540)
- When determining which transaction to kill after deadlock has been detected, `InnoDB` now adds the number of locks to a transaction's weight, and avoids killing transactions that have modified nontransactional tables. This should reduce the likelihood of killing long-running transactions containing `SELECT ... FOR UPDATE` or `INSERT/REPLACE INTO ... SELECT` statements, and of causing partial updates if the target is a `MyISAM` table. (Bug #21293)
- `InnoDB` displayed an incorrect error message when a `CREATE TABLE` statement exceeded the `InnoDB` maximum permissible row size. (Bug #21101)
- Under heavy load with a large query cache, invalidating part of the cache could cause the server to freeze (that is, to be unable to service other operations until the invalidation was complete). (Bug #21074)

References: See also: Bug #39253.

- On Windows, the server used 10MB of memory for each connection thread, resulting in memory exhaustion. Now each thread uses 1MB. (Bug #20815)
- InnoDB produced an unnecessary (and harmless) warning: `InnoDB: Error: trying to declare trx to enter InnoDB, but InnoDB: it already is declared.` (Bug #20090)
- If a slave timed out while registering with the master to which it was connecting, auto-reconnect failed thereafter. (Bug #19328)
- If InnoDB reached its limit on the number of concurrent transactions (1023), it wrote a descriptive message to the error log but returned a misleading error message to the client, or an assertion failure occurred. (Bug #18828)

References: See also: Bug #46672.

- Under ActiveState Perl, `mysql-test-run.pl` would not run. (Bug #18415)
- The server crashed when the size of an ARCHIVE table grew larger than 2GB. (Bug #15787)
- `SQL_BIG_RESULT` had no effect for `CREATE TABLE ... SELECT SQL_BIG_RESULT ...` statements. (Bug #15130)
- On 64-bit Windows systems, the Config Wizard failed to complete the setup because 64-bit Windows does not resolve dynamic linking of the 64-bit `libmysql.dll` to a 32-bit application like the Config Wizard. (Bug #14649)
- `mysql_setpermission` tried to grant global-only privileges at the database level. (Bug #14618)
- For the general query log, logging of prepared statements executed using the C API differed from logging of prepared statements performed with `PREPARE` and `EXECUTE`. Logging for the latter was missing the `Prepare` and `Execute` lines. (Bug #13326)
- The `TABLE_COMMENT` column of `INFORMATION_SCHEMA.TABLES` and the `Comment` column in the output of `SHOW TABLE STATUS` displayed extraneous information for InnoDB and NDBCLUSTER tables. (Bug #11379)

References: See also: Bug #32440.

- The server returned data from `SHOW CREATE TABLE` statement or a `SELECT` statement on an `INFORMATION_SCHEMA` table using the `binary` character set. (Bug #10491)
- Backup software can cause `ERROR_SHARING_VIOLATION` or `ERROR_LOCK_VIOLATION` conditions during file operations. InnoDB now retries forever until the condition goes away. (Bug #9709)

Changes in MySQL 5.1.20 (2007-06-25)

This is a new Beta development release, fixing recently discovered bugs.



Note

This Beta 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.

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 Enterprise* (a commercial MySQL offering). For more details, please see (<http://www.mysql.com/products/enterprise>).

- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

Functionality Added or Changed

- **Incompatible Change:** It is no longer possible to partition the log tables. (Bug #27816)
- **Incompatible Change:** `mysqld_safe` now supports error logging to `syslog` on systems that support the `logger` command. The new `--syslog` and `--skip-syslog` options can be used instead of the `--log-error` option to control logging behavior, as described in [mysqld_safe — MySQL Server Startup Script](#). The default is to use `syslog`, which differs from the previous default behavior of writing an error log file.

Currently, logging to `syslog` may fail to operate correctly in some cases; if so, use `--skip-syslog` or `--log-error`. To maintain the older behavior if you were using no error-logging option, use `--skip-syslog`. If you were using `--log-error`, continue to use it.

Note: In 5.1.21, the default is changed to `--skip-syslog`, which is compatible with releases prior to 5.1.20. (Bug #4858)

- **Important Change; MySQL Cluster:** The `TimeBetweenWatchdogCheckInitial` configuration parameter was added to enable setting of a separate watchdog timeout for memory allocation during startup of the data nodes. (Bug #28899)
- **MySQL Cluster:** The cluster management client now stores command history between sessions. (Bug #29073)
- **MySQL Cluster:** `auto_increment_increment` and `auto_increment_offset` are now supported for `NDB` tables. (Bug #26342)
- **MySQL Cluster:** The server source tree now includes scripts to simplify building MySQL with SCI support. For more information about SCI interconnects and these build scripts, see [Configuring MySQL Cluster to use SCI Sockets](#). (Bug #25470)
- **MySQL Cluster:** A new configuration parameter `ODirect` causes `NDB` to attempt using `O_DIRECT` writes for LCP, backups, and redo logs, often lowering CPU usage.
- **Replication:** The `sql_mode`, `foreign_key_checks`, `unique_checks`, character set/collations, and `sql_auto_is_null` session variables are written to the binary log and honored during replication. See [The Binary Log](#).
- If a `MERGE` table cannot be opened or used because of a problem with an underlying table, `CHECK TABLE` now displays information about which table caused the problem. (Bug #26976)
- User variables and stored procedure variables are now supported for use in XPath expressions employed as arguments to the `ExtractValue()` and `UpdateXML()` functions.

This means that:

- XPath can now be used to load data from XML files using virtually any format, and so able to import data from most third party software which either has XML export functionality, or uses XML natively as a storage format.
- Various complex conditions can be put on rows and columns, so one can filter for desired rows (or skip unwanted rows) when loading XML.
- Various types of preprocessing using SQL functions are now possible when loading XML. For example, you can concatenate two XML tag or attribute values into a single column value using `CONCAT()`, or remove some parts of the data using `REPLACE()`.

See [XML Functions](#), for more information. (Bug #26518)

- Binary distributions for some platforms did not include shared libraries; now shared libraries are shipped for all platforms except AIX 5.2 64-bit. *Exception:* The library for the `libmysqld` embedded server is not shared except on Windows. (Bug #16520, Bug #26767, Bug #13450)
- Added a new `PAD_CHAR_TO_FULL_LENGTH` SQL mode. By default, trailing spaces are trimmed from `CHAR` column values on retrieval. If `PAD_CHAR_TO_FULL_LENGTH` is enabled, trimming does not occur and retrieved `CHAR` values are padded to their full length. This mode does not apply to `VARCHAR` columns, for which trailing spaces are retained on retrieval.
- XPath can now be used to load data from XML files using virtually any format, and so able to import data from most third party software which either has XML export functionality, or uses XML natively as a storage format.
- Various complex conditions can be put on rows and columns, so one can filter for desired rows (or skip unwanted rows) when loading XML.
- Various types of preprocessing using SQL functions are now possible when loading XML. For example, you can concatenate two XML tag or attribute values into a single column value using `CONCAT()`, or remove some parts of the data using `REPLACE()`.

Bugs Fixed

- **Security Fix:** A malformed password packet in the connection protocol could cause the server to crash. Thanks for Dormando for reporting this bug, and for providing details and a proof of concept. (Bug #28984, CVE-2007-3780)
- **Security Fix:** `CREATE TABLE LIKE` did not require any privileges on the source table. Now it requires the `SELECT` privilege.

In addition, `CREATE TABLE LIKE` was not isolated from alteration by other connections, which resulted in various errors and incorrect binary log order when trying to execute concurrently a `CREATE TABLE LIKE` statement and either DDL statements on the source table or DML or DDL statements on the target table. (Bug #23667, Bug #25578, CVE-2007-3781)

- **Incompatible Change:** Some error codes had error numbers in MySQL 5.1 different from the numbers in MySQL 5.0. In MySQL 5.1, error numbers have been changed to match the MySQL 5.0 values: Error codes with value of 1458 or higher have changed in MySQL 5.1 now. Client applications designed to work with MySQL 5.1 with hard-coded error code values (for example, in statements such as `if (mysql_errno(mysql) == 1463) { ... }`) need to be updated in the source code. All clients designed to work with MySQL 5.1 that test error codes (for example, in statements such as `if (mysql_errno(mysql) == ER_VIEW_RECURSIVE) { ... }`) should be recompiled. Existing 5.0 clients should now work, without changes or recompilation, against servers for MySQL 5.1.20 or higher. (Bug #29245)

- **Incompatible Change:** The names of stored functions referenced by views were not properly displayed by `SHOW CREATE VIEW`.

The fix corrects a problem introduced by Bug #23491. There is an incompatibility when upgrading from versions affected by that bug fix (MySQL 5.0.40 through 5.0.43, MySQL 5.1.18 through 5.1.19): If you use `mysqldump` before upgrading from an affected version and reload the data after upgrading to a higher version, you must drop and recreate your views. (Bug #28605)

References: This issue is a regression of: Bug #23491.

- **Incompatible Change:** When `mysqldump` was run with the `--delete-master-logs` option, binary log files were deleted before it was known that the dump had succeeded, not after. (The method for removing log files used `RESET MASTER` prior to the dump. This also reset the binary log sequence numbering to `.000001`.) Now `mysqldump` flushes the logs (which creates a new binary log number with the next sequence number), performs the dump, and then uses `PURGE BINARY LOGS` to remove the log files older than the new one. This also preserves log numbering because the new log with the next number is generated and only the preceding logs are removed. However, this may affect applications if they rely on the log numbering sequence being reset. (Bug #24733)
- **Incompatible Change:** The use of an `ORDER BY` or `DISTINCT` clause with a query containing a call to the `GROUP_CONCAT()` function caused results from previous queries to be redisplayed in the current result. The fix for this includes replacing a `BLOB` value used internally for sorting with a `VARCHAR`. This means that for long results (more than 65,535 bytes), it is possible for truncation to occur; if so, an appropriate warning is issued. (Bug #23856, Bug #28273)
- **MySQL Cluster; Replication:** (Replication): A replicated unique key permitted duplicate key inserts on the slave. (Bug #27044)
- **MySQL Cluster:** Memory corruption could occur due to a problem in the `DBTUP` kernel block. (Bug #29229)
- **MySQL Cluster:** A query having a large `IN(...)` or `NOT IN(...)` list in the `WHERE` condition on an `NDB` table could cause `mysqld` to crash. (Bug #29185)
- **MySQL Cluster:** In the event that two data nodes in the same node group and participating in a GCP crashed before they had written their respective `P0.sysfile` files, `QMGR` could refuse to start, issuing an invalid `Insufficient nodes for restart` error instead. (Bug #29167)
- **MySQL Cluster:** Attempting to restore a `NULL` row to a `VARBINARY` column caused `ndb_restore` to fail. (Bug #29103)
- **MySQL Cluster:** `ndb_error_reporter` now preserves timestamps on files. (Bug #29074)
- **MySQL Cluster:** It is now possible to set the maximum size of the allocation unit for table memory using the `MaxAllocate` configuration parameter. (Bug #29044)
- **MySQL Cluster:** When shutting down `mysqld`, the `NDB` binlog process was not shut down before log cleanup began. (Bug #28949)
- **MySQL Cluster:** `ndb_mgm` could hang when connecting to a nonexistent host. (Bug #28847)
- **MySQL Cluster:** A regression in the heartbeat monitoring code could lead to node failure under high load. This issue affected MySQL 5.1.19 and MySQL Cluster NDB 6.1.10 only. (Bug #28783)
- **MySQL Cluster:** A corrupt schema file could cause a `File already open` error. (Bug #28770)
- **MySQL Cluster:** Having large amounts of memory locked caused swapping to disk. (Bug #28751)

- **MySQL Cluster:** Setting `InitialNoOfOpenFiles` equal to `MaxNoOfOpenFiles` caused an error. This was due to the fact that the actual value of `MaxNoOfOpenFiles` as used by the cluster was offset by 1 from the value set in `config.ini`. (Bug #28749)
- **MySQL Cluster:** LCP files were not removed following an initial system restart. (Bug #28726)
- **MySQL Cluster:** `UPDATE IGNORE` statements involving the primary keys of multiple tables could result in data corruption. (Bug #28719)
- **MySQL Cluster:** A race condition could result when nonmaster nodes (in addition to the master node) tried to update active status due to a local checkpoint (that is, between `NODE_FAILREP` and `COPY_GCIREQ` events). Now only the master updates the active status. (Bug #28717)
- **MySQL Cluster:** A fast global checkpoint under high load with high usage of the redo buffer caused data nodes to fail. (Bug #28653)
- **MySQL Cluster:** The management client's response to `START BACKUP WAIT COMPLETED` did not include the backup ID. (Bug #27640)
- **Replication; Cluster Replication:** When replicating `MyISAM` or `InnoDB` tables to a MySQL Cluster, it was not possible to determine exactly what had been applied following a shutdown of the slave cluster or `mysqld` process. (Bug #26783)
- **Replication:** `DROP USER` statements that named multiple users, only some of which could be dropped, were replicated incorrectly. (Bug #29030)
- **Replication:** Using events in replication could cause the slave to crash. (Bug #28953)
- **Replication:** It was possible to set `SQL_SLAVE_SKIP_COUNTER` such that the slave would jump into the middle of an event group. (Bug #28618)

References: See also: Bug #12691.

- **Replication:** The result of executing of a prepared statement created with `PREPARE s FROM "SELECT 1 LIMIT ?"` was not replicated correctly. (Bug #28464)
- **Replication:** Recreating a view that already exists on the master would cause a replicating slave to terminate replication with a 'different error message on slave and master' error. (Bug #28244)
- **Replication:** Binary logging of prepared statements could produce syntactically incorrect queries in the binary log, replacing some parameters with variable names rather than variable values. This could lead to incorrect results on replication slaves. (Bug #26842, Bug #12826)
- **Replication:** Connections from one `mysqld` server to another failed on Mac OS X, affecting replication and `FEDERATED` tables. (Bug #26664)

References: See also: Bug #29083.

- **Replication:** When using transactions and replication, shutting down the master in the middle of a transaction would cause all slaves to stop replicating. (Bug #22725)
- **Replication:** Using `CREATE TABLE LIKE ...` would raise an assertion when replicated to a slave. (Bug #18950)
- **Disk Data:** When loading data into a cluster following a version upgrade, the data nodes could forcibly shut down due to page and buffer management failures (that is, `ndbrequire` failures in `PGMAN`). (Bug #28525)

- **Disk Data:** Repeated `INSERT` and `DELETE` operations on a Disk Data table having one or more large `VARCHAR` columns could cause data nodes to fail. (Bug #20612)
- **Cluster API:** The timeout set using the MGM API `ndb_mgm_set_timeout()` function was incorrectly interpreted as seconds rather than as milliseconds. (Bug #29063)
- **Cluster API:** An invalid error code could be set on transaction objects by `BLOB` handling code. (Bug #28724)
- The `TRUNCATE TABLE` statement was handled differently by the server when row-based logging was in effect, even though the binlogging format in effect does not effect the fact that `TRUNCATE TABLE` is always logged as a statement. (Bug #29130)
- If one of the queries in a `UNION` used the `SQL_CACHE` option and another query in the `UNION` contained a nondeterministic function, the result was still cached. For example, this query was incorrectly cached:

```
SELECT NOW() FROM t1 UNION SELECT SQL_CACHE 1 FROM t1;
```

(Bug #29053)

- Long path names for internal temporary tables could cause stack overflows. (Bug #29015)
- Using an `INTEGER` column from a table to `ROUND()` a number produced different results than using a constant with the same value as the `INTEGER` column. (Bug #28980)
- If a program binds a given number of parameters to a prepared statement handle and then somehow changes `stmt->param_count` to a different number, `mysql_stmt_execute()` could crash the client or server. (Bug #28934)
- Queries using UDFs or stored functions were cached. (Bug #28921)
- `INSERT .. ON DUPLICATE KEY UPDATE` could under some circumstances silently update rows when it should not have. (Bug #28904)
- Queries that used `UUID()` were incorrectly permitted into the query cache. (This should not happen because `UUID()` is nondeterministic.) (Bug #28897)
- Using a `VIEW` created with a nonexistent `DEFINER` could lead to incorrect results under some circumstances. (Bug #28895)
- For `InnoDB` tables that use the `utf8` character set, incorrect results could occur for DML statements such as `DELETE` or `UPDATE` that use an index on character-based columns. (Bug #28878)

References: See also: Bug #29449, Bug #30485, Bug #31395. This issue is a regression of: Bug #13195.

- Non-`utf8` characters could get mangled when stored in `CSV` tables. (Bug #28862)
- On Windows, `USE_TLS` was not defined for `mysqlclient.lib`. (Bug #28860)
- In MySQL 5.1.15, a new error code `ER_DUP_ENTRY_WITH_KEY_NAME` (1582) was introduced to replace `ER_DUP_ENTRY` (1062) so that the key name could be provided instead of the key number. This was unnecessary, so `ER_DUP_ENTRY` is used again and the key name is printed. The incompatibility introduced in 5.1.15 no longer applies. (Bug #28842)
- A subquery with `ORDER BY` and `LIMIT 1` could cause a server crash. (Bug #28811)
- Running `SHOW TABLE STATUS` while performing a high number of inserts on partitioned tables with a great many partitions could cause the server to crash. (Bug #28806)

- Using `BETWEEN` with nonindexed date columns and short formats of the date string could return incorrect results. (Bug #28778)
- Selecting `GEOMETRY` columns in a `UNION` caused a server crash. (Bug #28763)
- When constructing the path to the original `.frm` file, `ALTER ... RENAME` was unnecessarily (and incorrectly) lowercasing the entire path when not on a case-insensitive file system, causing the statement to fail. (Bug #28754)
- The `binlog_format` system variable value was empty if the server was started with binary logging disabled. Now it is set to `MIXED`. (Bug #28752)
- Searches on indexed and nonindexed `ENUM` columns could return different results for empty strings. (Bug #28729)
- Executing `EXPLAIN EXTENDED` on a query using a derived table over a grouping subselect could lead to a server crash. This occurred only when materialization of the derived tables required creation of an auxiliary temporary table, an example being when a grouping operation was carried out with usage of a temporary table. (Bug #28728)
- The result of evaluation for a view's `CHECK OPTION` option over an updated record and records of merged tables was arbitrary and dependent on the order of records in the merged tables during the execution of the `SELECT` statement. (Bug #28716)
- The “manager thread” of the LinuxThreads implementation was unintentionally started before `mysqld` had dropped privileges (to run as an unprivileged user). This caused signaling between threads in `mysqld` to fail when the privileges were finally dropped. (Bug #28690)
- Setting an interval of `EVERY 0 SECOND` for a scheduled event caused the server to crash. (Bug #28666)
- For debug builds, `ALTER TABLE` could trigger an assertion failure due to occurrence of a deadlock when committing changes. (Bug #28652)
- Attempting to create an index on a `BIT` column failed after modifying the column. (Bug #28631)
- Conversion of U+00A5 YEN SIGN and U+203E OVERLINE from `ucs2` to `ujis` produced incorrect results. (Bug #28600)
- Killing from one connection a long-running `EXPLAIN QUERY` started from another connection caused `mysqld` to crash. (Bug #28598)
- `SHOW GLOBAL VARIABLES` repeated some variable names. (Bug #28580)
- When one thread attempts to lock two (or more) tables and another thread executes a statement that aborts these locks (such as `REPAIR TABLE`, `OPTIMIZE TABLE`, or `CHECK TABLE`), the thread might get a table object with an incorrect lock type in the table cache. The result is table corruption or a server crash. (Bug #28574)
- Outer join queries with `ON` conditions over constant outer tables did not return `NULL`-complemented rows when conditions were evaluated to `FALSE`. (Bug #28571)
- An update on a multiple-table view with the `CHECK OPTION` clause and a subquery in the `WHERE` condition could cause an assertion failure. (Bug #28561)
- Calling the `UpdateXML()` function using invalid XPath syntax caused memory corruption possibly leading to a crash of the server. (Bug #28558)
- `PURGE MASTER LOGS BEFORE (subquery)` caused a server crash. Subqueries are forbidden in the `BEFORE` clause now. (Bug #28553)

- `mysqldump` calculated the required memory for a hex-blob string incorrectly causing a buffer overrun. This in turn caused `mysqldump` to crash silently and produce incomplete output. (Bug #28522)
- When upgrading from MySQL 5.1.17 to 5.1.18, `mysql_upgrade` and `mysql_fix_privilege_tables` did not upgrade the system tables relating to the Event Scheduler correctly. (Bug #28521)
- Passing a `DECIMAL` value as a parameter of a statement prepared with `PREPARE` resulted in an error. (Bug #28509)
- `mysql_affected_rows()` could return an incorrect result for `INSERT ... ON DUPLICATE KEY UPDATE` if the `CLIENT_FOUND_ROWS` flag was set. (Bug #28505)
- A query that grouped by the result of an expression returned a different result when the expression was assigned to a user variable. (Bug #28494)
- Subselects returning `LONG` values in MySQL versions later than 5.0.24a returned `LONGLONG` prior to this. The previous behavior was restored. (Bug #28492)

References: This issue is a regression of: Bug #19714.

- Performing `ALTER TABLE ... ADD PARTITION` or `ALTER TABLE DROP PARTITION` could result in inconsistent data, or cause the server to crash, if done concurrently with other accesses to the table. (Bug #28477, Bug #28488)
- Forcing the use of an index on a `SELECT` query when the index had been disabled would raise an error without running the query. The query now executes, with a warning generated noting that the use of a disabled index has been ignored. (Bug #28476)
- The query `SELECT '2007-01-01' + INTERVAL column_name DAY FROM table_name` caused `mysqld` to fail. (Bug #28450)
- A server crash could happen under rare conditions such that a temporary table outgrew heap memory reserved for it and the remaining disk space was not big enough to store the table as a `MyISAM` table. (Bug #28449)
- Using `ALTER TABLE` to move columns resulted only in the columns being renamed. The table contents were not changed. (Bug #28427)
- The test case for `mysqldump` failed with `bin-log` disabled. (Bug #28372)
- Attempting to `LOAD_FILE` from an empty floppy drive under Windows, caused the server to hang. For example, if you opened a connection to the server and then issued the command `SELECT LOAD_FILE('a:test');`, with no floppy in the drive, the server was inaccessible until the modal pop-up dialog box was dismissed. (Bug #28366)
- `mysqltest` used a too-large stack size on PowerPC/Debian Linux, causing thread-creation failure for tests that use many threads. (Bug #28333)
- When using a `MEMORY` table on Mac OS X, dropping a table and then creating a table with the same name could cause the information of the deleted table to remain accessible, leading to index errors. (Bug #28309)
- The `IS_UPDATABLE` column in the `INFORMATION_SCHEMA.VIEWS` table was not always set correctly. (Bug #28266)
- For `CAST()` of a `NULL` value with type `DECIMAL`, the return value was incorrectly initialized, producing a runtime error for binaries built using Visual C++ 2005. (Bug #28250)

- When the query cache was fully used, issuing `RENAME DATABASE` or `RENAME SCHEMA` could cause the server to hang, with 100% CPU usage. (Bug #28211)
- The `Bytes_received` and `Bytes_sent` status variables could hold only 32-bit values (not 64-bit values) on some platforms. (Bug #28149)
- Some valid identifiers were not parsed correctly. (Bug #28127)
- Storing a large number into a `FLOAT` or `DOUBLE` column with a fixed length could result in incorrect truncation of the number if the column's length was greater than 31. (Bug #28121)
- Sending debugging information from a dump of the Event Scheduler to `COM_DEBUG` could cause the server to crash. (Bug #28075)
- The `PARTITION_COMMENT` column of the `INFORMATION_SCHEMA.PARTITIONS` table had the wrong default value. (Bug #28007)
- `DECIMAL` values beginning with nine 9 digits could be incorrectly rounded. (Bug #27984)
- For attempts to open a nonexistent table, the server should report `ER_NO_SUCH_TABLE` but sometimes reported `ER_TABLE_NOT_LOCKED`. (Bug #27907)
- Following an invalid call to `UpdateXML()`, calling the function again (even if valid) crashed the server. (Bug #27898)
- A stored program that uses a variable name containing multibyte characters could fail to execute. (Bug #27876)
- The server made strong assumptions about the structure of the `general_log` and `slow_log` log tables: It supported only the table structure defined in the `mysql` database creation scripts. The server also permitted limited `ALTER TABLE` operations on the log tables, but adding an `AUTO_INCREMENT` column did not properly initialize the column, and subsequent inserts into the table could fail to generate correct sequence numbers. Now an `ALTER TABLE` statement that adds an `AUTO_INCREMENT` column populates the column correctly. In addition, when the server writes a log table row, it will set columns not present in the original table structure to their default values. (Bug #27857)
- `ON` conditions from `JOIN` expressions were ignored when checking the `CHECK OPTION` clause while updating a multiple-table view that included such a clause. (Bug #27827)
- On some systems, `udf_example.c` returned an incorrect result length. Also on some systems, `mysql-test-run.pl` could not find the shared object built from `udf_example.c`. (Bug #27741)
- The modification of a table by a partially completed multi-column update was not recorded in the binlog, rather than being marked by an event and a corresponding error code. (Bug #27716)
- `SHOW ENGINES` and queries on `INFORMATION_SCHEMA.ENGINES` did not use the same values for representing the same storage engine states. (Bug #27684)
- `HASH` indexes on `VARCHAR` columns with binary collations did not ignore trailing spaces from strings before comparisons. This could result in duplicate records being successfully inserted into a `MEMORY` table with unique key constraints. A consequence was that internal `MEMORY` tables used for `GROUP BY` calculation contained duplicate rows that resulted in duplicate-key errors when converting those temporary tables to `MyISAM`, and that error was incorrectly reported as a `table is full` error. (Bug #27643)
- An error occurred trying to connect to `mysqld-debug.exe`. (Bug #27597)
- A stack overrun could occur when storing `DATETIME` values using repeated prepared statements. (Bug #27592)

- If a stored function or trigger was killed, it aborted but no error was thrown, permitting the calling statement to continue without noticing the problem. This could lead to incorrect results. (Bug #27563)
- When `ALTER TABLE` was used to add a new `DATE` column with no explicit default value, `'0000-00-00'` was used as the default even if the SQL mode included the `NO_ZERO_DATE` mode to prohibit that value. A similar problem occurred for `DATETIME` columns. (Bug #27507)
- `ALTER TABLE ... ENABLE KEYS` could cause `mysqld` to crash when executed on a table containing on a `MyISAM` table containing billions of rows. (Bug #27029)
- Binary content `0x00` in a `BLOB` column sometimes became `0x5C 0x00` following a dump and reload, which could cause problems with data using multibyte character sets such as `GBK` (Chinese). This was due to a problem with `SELECT INTO OUTFILE` whereby `LOAD DATA` later incorrectly interpreted `0x5C` as the second byte of a multibyte sequence rather than as the `SOLIDUS` (`"\"`) character, used by MySQL as the escape character. (Bug #26711)
- The server crashed when attempting to open a table having a `#mysql50#` prefix in the database or table name. The server now will not open such tables. (This prefix is reserved by `mysql_upgrade` for accessing 5.0 tables that have names not yet encoded for 5.1.) (Bug #26402)
- A `FLUSH TABLES WITH READ LOCK` statement followed by a `FLUSH LOGS` statement caused a deadlock if the general log or the slow query log was enabled. (Bug #26380)
- The query `SELECT /*2*/ user, host, db, info FROM INFORMATION_SCHEMA.PROCESSLIST WHERE (command!='Daemon' || user='event_scheduler') AND (info IS NULL OR info NOT LIKE '%processlist%') ORDER BY INFO` yielded inconsistent results. (Bug #26338)
- For a given user variable `@v`, the statements `SELECT @v` and `CREATE TABLE ... AS SELECT @v` did not return the same data type. (Bug #26277)
- Statements within triggers ignored the value of the `low_priority_updates` system variable. (Bug #26162)

References: See also: Bug #29963.

- The embedded server library displayed error messages at startup if the `mysql.plugin` table was not present. This no longer occurs. (Bug #25800)
- On Windows, an application that called `mysql_thread_init()` but forgot to call `mysql_thread_end()` would get this error: `Error in my_thread_global_end()`. (Bug #25621)
- Embedded `/* ... */` comments were handled incorrectly within the definitions of stored programs and views, resulting in malformed definitions (the trailing `*/` was stripped). This also affected binary log contents. (Bug #25411, Bug #26302)
- Due to a race condition, executing `FLUSH PRIVILEGES` in one thread could cause brief table unavailability in other threads. (Bug #24988)
- In `SHOW SLAVE STATUS` output, `Last_Errno` and `Last_Error` were not set after `master_retry_count` errors had occurred. To provide additional information, the statement now displays four additional columns:
 - `Last_IO_Errno`: The number of the last error that caused the I/O thread to stop
 - `Last_IO_Error`: A description of the last error that caused the I/O thread to stop
 - `Last_SQL_Errno`: The number of the last error that caused the SQL thread to stop
 - `Last_SQL_Error`: A description of the last error that caused the SQL thread to stop

Also, `Last_Errno` and `Last_Error` now are aliases for `Last_SQL_Errno` and `Last_SQL_Error`. (Bug #24954)

- A too-long `shared-memory-base-name` value could cause a buffer overflow and crash the server or clients. (Bug #24924)
- When `mysqld` was run as a Windows service, shared memory objects were not created in the global namespace and could not be used by clients to connect. (Bug #24731)
- On some Linux distributions where LinuxThreads and NPTL `glibc` versions both are available, statically built binaries can crash because the linker defaults to LinuxThreads when linking statically, but calls to external libraries (such as `libnss`) are resolved to NPTL versions. This cannot be worked around in the code, so instead if a crash occurs on such a binary/OS combination, print an error message that provides advice about how to fix the problem. (Bug #24611)
- A number of `SHOW` statements caused `mysqld` to crash on recent versions of Solaris. This issue is believed to be present only in MySQL 5.1.12 and later. (Bug #23810)
- The server deducted some bytes from the `key_cache_block_size` option value and reduced it to the next lower 512 byte boundary. The resulting block size was not a power of two. Setting the `key_cache_block_size` system variable to a value that is not a power of two resulted in `MyISAM` table corruption. (Bug #23068, Bug #28478, Bug #25853)
- Conversion errors could occur when constructing the condition for an `IN` predicate. The predicate was treated as if the affected column contains `NULL`, but if the `IN` predicate is inside `NOT`, incorrect results could be returned. (Bug #22855)
- Linux binaries were unable to dump core after executing a `setuid()` call. (Bug #21723)
- Stack overflow caused server crashes. (Bug #21476)
- The server was ignoring the return value of the `parse()` function for full-text parser plugins. (Bug #18839)
- Granting access privileges to an individual table where the database or table name contained an underscore failed. (Bug #18660)
- 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)
- The `check-cpu` script failed to detect AMD64 Turion processors correctly. (Bug #17707)
- When using `mysqlbinlog` with `--read-from-remote-server` to load the data direct from a remote MySQL server would cause a core dump when dumping certain binary log events. (Bug #17654)
- Trying to shut down the server following a failed `LOAD DATA INFILE` caused `mysqld` to crash. (Bug #17233)
- 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)
- Using up-arrow for command-line recall in `mysql` could cause a segmentation fault. (Bug #10218)
- The result for `CAST()` when casting a value to `UNSIGNED` was limited to the maximum signed `BIGINT` value (9223372036854775808), rather than the maximum unsigned value (18446744073709551615). (Bug #8663)

- The internal functions for table preparation, creation, and alteration were not re-execution friendly, causing problems in code that: repeatedly altered a table; repeatedly created and dropped a table; opened and closed a cursor on a table, altered the table, and then reopened the cursor; used `ALTER TABLE` to change a table's current `AUTO_INCREMENT` value; created indexes on `utf8` columns.

Re-execution of `CREATE DATABASE`, `CREATE TABLE`, and `ALTER TABLE` statements in stored routines or as prepared statements also caused incorrect results or crashes. (Bug #4968, Bug #6895, Bug #19182, Bug #19733, Bug #22060, Bug #24879)

Changes in MySQL 5.1.19 (2007-05-25)

This is a new Beta development release, fixing recently discovered bugs.



Note

This Beta 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.

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 Enterprise* (a commercial MySQL offering). For more details, please see (<http://www.mysql.com/products/enterprise>).

- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

Functionality Added or Changed

- **Incompatible Change:** `INSERT DELAYED` is now downgraded to a normal `INSERT` if the statement uses functions that access tables or triggers, or that is called from a function or a trigger.

This was done to resolve the following interrelated issues:

- The server could abort or deadlock for `INSERT DELAYED` statements for which another insert was performed implicitly (for example, using a stored function that inserted a row).
- A trigger using an `INSERT DELAYED` caused the error `INSERT DELAYED can't be used with table ... because it is locked with LOCK TABLES` although the target table was not actually locked.
- `INSERT DELAYED` into a table with a `BEFORE INSERT` or `AFTER INSERT` trigger gave an incorrect `NEW` pseudocolumn value and caused the server to deadlock or abort.

(Bug #21483)

References: See also: Bug #20497, Bug #21714.

- **MySQL Cluster:** Formerly, restoring a cluster backup made on a MySQL 5.0 Cluster to a 5.1 cluster using a 5.1 version of `ndb_restore` did not resize `VARCHAR` columns as might be expected; now, the default behavior of `ndb_restore` in such cases is to resize the `VARCHAR` columns. This

changed default behavior can be overridden using the `--no-upgrade` (or `-u`) option when invoking `ndb_restore`. (Bug #22240)

- The `BLACKHOLE` storage engine now supports `INSERT DELAYED`. Previously, `INSERT DELAYED` statements for `BLACKHOLE` tables were not supported, and caused the server to crash. (Bug #27998)
- A new status variable, `Com_call_procedure`, indicates the number of calls to stored procedures. (Bug #27994)
- The `BLACKHOLE` storage engine now supports `LOCK TABLES` and `UNLOCK TABLES`. (Bug #26241)
- `GLOBAL_STATUS`
- `GLOBAL_VARIABLES`
- `SESSION_VARIABLES`
- The data type used for the `VARIABLE_VALUE` column of the following `INFORMATION_SCHEMA` tables has been changed to `VARCHAR`:
 - `GLOBAL_STATUS`
 - `SESSION_STATUS`
 - `GLOBAL_VARIABLES`
 - `SESSION_VARIABLES`

For more information, see [The INFORMATION_SCHEMA GLOBAL_STATUS and SESSION_STATUS Tables](#), and [The INFORMATION_SCHEMA GLOBAL_VARIABLES and SESSION_VARIABLES Tables](#).

References: See also: Bug #26994.

- `SESSION_STATUS`

Bugs Fixed

- **Security Fix:** UDFs are supposed to be loadable only from the plugin directory, but this restriction was not being enforced. (Bug #28341)
- **Security Fix:** Use of a view could enable a user to gain update privileges for tables in other databases. (Bug #27878, CVE-2007-3782)
- **MySQL Cluster:** When an API node sent more than 1024 signals in a single batch, `NDB` would process only the first 1024 of these, and then hang. (Bug #28443)
- **MySQL Cluster:** A delay in obtaining `AUTO_INCREMENT` IDs could lead to excess temporary errors. (Bug #28410)
- **MySQL Cluster:** Local checkpoint files relating to dropped `NDB` tables were not removed. (Bug #28348)
- **MySQL Cluster:** Multiple operations involving deletes followed by reads were not handled correctly.



Note

This issue could also affect MySQL Cluster Replication.

(Bug #28276)

- **MySQL Cluster:** Repeated insertion of data generated by `mysqldump` into `NDB` tables could eventually lead to failure of the cluster. (Bug #27437)
- **MySQL Cluster:** Restarting a data node caused SQL nodes to log repeatedly and unnecessarily the status of the event buffer, causing a memory leak of approximately 4 MB for each `mysqld` process each time this occurred.

(This issue was known to occur in MySQL 5.1.16 and later only.) (Bug #27292)
- **MySQL Cluster:** `ndb_mgmd` failed silently when the cluster configuration file contained invalid `[tcp]` entries. (Bug #27207)
- **MySQL Cluster:** `ndb_connectstring` did not appear in the output of `SHOW VARIABLES`. (Bug #26675)
- **MySQL Cluster:** A failure to release internal resources following an error could lead to problems with single user mode. (Bug #25818)
- **MySQL Cluster:** DDL operations were not supported on a partially started cluster. (Bug #24631)
- **Disk Data:** Extremely large inserts into Disk Data tables could lead to data node failure in some circumstances. (Bug #27942)
- **Cluster API:** In a multi-operation transaction, a delete operation followed by the insertion of an implicit `NULL` failed to overwrite an existing value. (Bug #20535)
- Some `ALTER TABLE` statements that worked in MySQL 5.0 did not work in 5.1. (Bug #28415)
- `mysql_upgrade` failed if certain SQL modes were set. Now it sets the mode itself to avoid this problem. (Bug #28401)
- A query with a `NOT IN` subquery predicate could cause a crash when the left operand of the predicate evaluated to `NULL`. (Bug #28375)
- A buffer overflow could occur when using `DECIMAL` columns on Windows operating systems. (Bug #28361)
- `libmysql.dll` could not be dynamically loaded on Windows. (Bug #28358)
- Grouping queries with correlated subqueries in `WHERE` conditions could produce incorrect results. (Bug #28337)
- `EXPLAIN` for a query on an empty table immediately after its creation could result in a server crash. (Bug #28272)
- Comparing a `DATETIME` column value with a user variable yielded incorrect results. (Bug #28261)
- Portability problems caused by use of `isinf()` were corrected. (Bug #28240)
- When dumping procedures, `mysqldump --compact` generated output that restored the session variable `sql_mode` without first capturing it. When dumping routines, `mysqldump --compact` neither set nor retrieved the value of `sql_mode`. (Bug #28223)
- Comparison of the string value of a date showed as unequal to `CURTIME()`. Similar behavior was exhibited for `DATETIME` values. (Bug #28208)
- For `InnoDB`, in some rare cases the optimizer preferred a more expensive `ref` access to a less expensive range access. (Bug #28189)

- Comparisons of `DATE` or `DATETIME` values for the `IN()` function could yield incorrect results. (Bug #28133)
- It was not possible to use the value `-9223372036854775808` (that is, `-MAXVALUE + 1`) when specifying a `LIST` partition. (Bug #28005)
- The server could hang for `INSERT IGNORE ... ON DUPLICATE KEY UPDATE` if an update failed. (Bug #28000)
- `CAST()` to `DECIMAL` did not check for overflow. (Bug #27957)
- The second execution of a prepared statement from a `UNION` query with `ORDER BY RAND()` caused the server to crash. This problem could also occur when invoking a stored procedure containing such a query. (Bug #27937)
- Views ignored precision for `CAST()` operations. (Bug #27921)
- Changes to some system variables should invalidate statements in the query cache, but invalidation did not happen. (Bug #27792)
- `LOAD DATA` did not use `CURRENT_TIMESTAMP` as the default value for a `TIMESTAMP` column for which no value was provided. (Bug #27670)
- Selecting `MIN()` on an indexed column that contained only `NULL` values caused `NULL` to be returned for other result columns. (Bug #27573)
- Using a `TEXT` local variable in a stored routine in an expression such as `SET var = SUBSTRING(var, 3)` produced an incorrect result. (Bug #27415)
- The error message for error number `137` did not report which database/table combination reported the problem. (Bug #27173)
- A large filesort could result in a division by zero error and a server crash. (Bug #27119)
- Some `InnoDB` variables were missing from the output of `mysqld --verbose --help`. (Bug #26987)
- Flow control optimization in stored routines could cause exception handlers to never return or execute incorrect logic. (Bug #26977)
- Some test suite files were missing from some MySQL-test packages. (Bug #26609)
- Running `CHECK TABLE` concurrently with a `SELECT`, `INSERT` or other statement on Windows could corrupt a MyISAM table. (Bug #25712)
- Concurrent execution of `CREATE TABLE ... SELECT` and other statements involving the target table suffered from various race conditions, some of which might have led to deadlocks. (Bug #24738)
- An attempt to execute `CREATE TABLE ... SELECT` when a temporary table with the same name already existed led to the insertion of data into the temporary table and creation of an empty nontemporary table. (Bug #24508)
- A statement of the form `CREATE TABLE IF NOT EXISTS t1 SELECT f1() AS i` failed with a deadlock error if the stored function `f1()` referred to a table with the same name as the to-be-created table. Now it correctly produces a message that the table already exists. (Bug #22427)
- Quoted labels in stored routines were mishandled, rendering the routines unusable. (Bug #21513)
- `CURDATE()` is less than `NOW()`, either when comparing `CURDATE()` directly (`CURDATE() < NOW()` is true) or when casting `CURDATE()` to `DATE` (`CAST(CURDATE() AS DATE) < NOW()` is true). However, storing `CURDATE()` in a `DATE` column and comparing `col_name < NOW()` incorrectly yielded false.

This is fixed by comparing a `DATE` column as `DATETIME` for comparisons to a `DATETIME` constant. (Bug #21103)

- `CREATE TABLE IF NOT EXISTS ... SELECT` caused a server crash if the target table already existed and had a `BEFORE INSERT` trigger. (Bug #20903)
- Deadlock occurred for attempts to execute `CREATE TABLE IF NOT EXISTS ... SELECT` when `LOCK TABLES` had been used to acquire a read lock on the target table. (Bug #20662, Bug #15522)
- For dates with 4-digit year parts less than 200, an incorrect implicit conversion to add a century was applied for date arithmetic performed with `DATE_ADD()`, `DATE_SUB()`, `+` `INTERVAL`, and `- INTERVAL`. (For example, `DATE_ADD('0050-01-01 00:00:00', INTERVAL 0 SECOND)` became `'2050-01-01 00:00:00'`.) (Bug #18997)
- Changing the size of a key buffer that is under heavy use could cause a server crash. The fix partially removes the limitation that `LOAD INDEX INTO CACHE` fails unless all indexes in a table have the same block size. Now the statement fails only if `IGNORE LEAVES` is specified. (Bug #17332)

Changes in MySQL 5.1.18 (2007-05-08)

This is a new Beta development release, fixing recently discovered bugs.



Note

This Beta 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.

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 Enterprise* (a commercial MySQL offering). For more details, please see (<http://www.mysql.com/products/enterprise>).

- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

Functionality Added or Changed

- **Incompatible Change; MySQL Cluster:** The internal specifications for columns in `NDB` tables has changed to enable compatibility with future MySQL Cluster releases that are expected to permit online adding and dropping of columns. This change is not backward compatible with earlier versions of MySQL Cluster.

See [Internal column specification changes](#), for important information prior to upgrading a MySQL Cluster to MySQL 5.1.18 or later from MySQL 5.1.17 or earlier.

References: See also: Bug #28205.

- **Incompatible Change; Replication:** The `INFORMATION_SCHEMA.EVENTS` and `mysql.event` tables have been changed to facilitate replication of events. When upgrading to MySQL 5.1.18, you must run `mysql_upgrade` prior to working with events. Until you have done so, any statement relating to the Event Scheduler or these tables (including `SHOW EVENTS`) will fail with the errors `Expected field status at position 12 to have type enum`

('ENABLED', 'SLAVESIDE_DISABLED', 'DISABLED'), found enum('ENABLED', 'DISABLED') and Table mysql.event is damaged. Can not open.

These changes were made as part of fixes for the following bugs:

- The effects of scheduled events were not replicated (that is, binary logging of scheduled events did not work).
- Effects of scheduled events on a replication master were both replicated and executed on the slave, causing double execution of events.
- `CREATE FUNCTION` statements and their effects were not replicated correctly.

For more information, see [Replication of Invoked Features](#). (Bug #17857, Bug #16421, Bug #20384, Bug #17671)

- **Incompatible Change; Cluster Replication:** The definition of the `mysql.ndb_apply_status` table has changed such that an online upgrade is not possible from MySQL 5.1.17 or earlier for a replication slave cluster; you must shut down all SQL nodes as part of the upgrade procedure. See [Upgrading and Downgrading MySQL Cluster](#), before upgrading.

For more information about the changes to `mysql.ndb_apply_status` see [MySQL Cluster Replication Schema and Tables](#).

- **Incompatible Change:** Prior to this release, when `DATE` values were compared with `DATETIME` values, the time portion of the `DATETIME` value was ignored, or the comparison could be performed as a string compare. Now a `DATE` value is coerced to the `DATETIME` type by adding the time portion as `00:00:00`. To mimic the old behavior, use the `CAST()` function as shown in this example: `SELECT date_col = CAST(NOW() AS DATE) FROM table;` (Bug #28929)
- **Incompatible Change:** The plugin interface and its handling of system variables was changed. Command-line options such as `--skip-innodb` now cause an error if `InnoDB` is not built-in or plugin-loaded. You should use `--loose-skip-innodb` if you do not want any error even if `InnoDB` is not available. The `--loose` prefix modifier should be used for all command-line options where you are uncertain whether the plugin exists and when you want the operation to proceed even if the option is necessarily ignored due to the absence of the plugin. (For a description of how `--loose` works, see [Using Options on the Command Line](#).)
- **Important Change:** When upgrading to MySQL 5.1.18 or later from a previous MySQL version and scheduled events have been used, the upgrade utilities do not accommodate changes in event-related system tables. As a workaround, you can dump events before the upgrade, then restore them from the dump afterward. This issue was fixed in MySQL 5.1.20.

References: See also: Bug #28521.

- **MySQL Cluster:** The behavior of the `ndb_restore` utility has been changed as follows:
 - It is now possible to restore selected databases or tables using `ndb_restore`.
 - Several options have been added for use with `ndb_restore --print_data` to facilitate the creation of structured data dump files. These options can be used to make dumps made using `ndb_restore` more like those produced by `mysqldump`.

For details of these changes, see `ndb_restore` — [Restore a MySQL Cluster Backup](#). (Bug #26899, Bug #26900)

- **MySQL Cluster:** The following changes were made in the `ndb_size.pl` utility:

- When `ndb_size.pl` calculates a value for a given configuration parameter that is less than the default value, it now suggests the default value instead.
- The dependency on `HTML::Template` was removed, with the result that the file `ndb_size.tmpl` is no longer needed or included.

(Bug #24227, Bug #24228)

- **Replication; Cluster Replication:** Some circular replication setups are now supported for MySQL Cluster. See [Known Issues in MySQL Cluster Replication](#), for detailed information. (Bug #17095, Bug #25688)
- **Cluster API:** The MGM API now supports explicit setting of network timeouts using the `ndb_mgm_set_timeout()` function. A utility function `ndb_mgm_number_of_mgmd_in_connect_string()` is also implemented to facilitate calculation of timeouts based on the number of management servers in the cluster.
- `mysqld_multi` now understands the `--no-defaults`, `--defaults-file`, and `--defaults-extra-file` options. The `--config-file` option is deprecated; if given, it is treated like `--defaults-extra-file`. (Bug #27390)
- If a set function `S` with an outer reference `S(outer_ref)` cannot be aggregated in the outer query against which the outer reference has been resolved, MySQL interprets `S(outer_ref)` the same way that it would interpret `S(const)`. However, standard SQL requires throwing an error in this situation. An error now is thrown for such queries if the `ANSI` SQL mode is enabled. (Bug #27348)
- Several additional data types are supported for columns in `INFORMATION_SCHEMA` tables: `DATE`, `TIME`, `BLOB`, `FLOAT`, and all integer types. (Bug #27047)
- The output of `mysql --xml` and `mysqldump --xml` now includes a valid XML namespace. (Bug #25946)
- If you use SSL for a client connection, you can tell the client not to authenticate the server certificate by specifying neither `--ssl-ca` nor `--ssl-capath`. The server still verifies the client according to any applicable requirements established using `GRANT` statements for the client, and it still uses any `--ssl-ca/--ssl-capath` values that were passed to server at startup time. (Bug #25309)
- Added a `MASTER_SSL_VERIFY_SERVER_CERT` option for the `CHANGE MASTER TO` statement, and a `Master_SSL_Verify_Server_Cert` output column to the `SHOW SLAVE STATUS` statement. The option value also is written to the `master.info` file. (Bug #19991)
- The `innodb_log_archive` system variable has been removed. The impact of this change should be low because the variable was unused, anyway.
- Added the `--auto-generate-sql-add-autoincrement`, `--auto-generate-sql-execute-number`, `--auto-generate-sql-guid-primary`, `--auto-generate-sql-secondary-indexes`, `--auto-generate-sql-unique-query-number`, `--auto-generate-sql-unique-write-number`, `--post-query`, and `--pre-query`, options for `mysqlslap`. Removed the `--lock-directory`, `--slave`, and `--use-threads` options.
- Added `--write-binlog` option for `mysqlcheck`. This option is enabled by default, but can be given as `--skip-write-binlog` to cause `ANALYZE TABLE`, `OPTIMIZE TABLE`, and `REPAIR TABLE` statements generated by `mysqlcheck` not to be written to the binary log. (Bug #26262)
- New command-line options: To alleviate ambiguities in variable names, all variables related to plugins can be specified using a `plugin` part in the name. For example, every time where we used to have `innodb` in the command-line options, you can now write `plugin-innodb`:

```
--skip-plugin-innodb
--plugin-innodb-buffer-pool-size=#
```

Furthermore, this is the preferred syntax. It helps to avoid ambiguities when a plugin, say, `wait`, has an option called `timeout`. `--wait-timeout` will still set a system variable, but `--plugin-wait-timeout` will set the plugin variable. Also, there is a new command-line option `--plugin-load` to install or load plugins at initialization time without using the `mysql.plugin` table.

- Storage engine plugins may now be uninstalled at run time. However, a plugin is not actually uninstalled until after its reference count drops to zero. The `default_storage_engine` system variable consumes a reference count, so uninstalling will not complete until said reference is removed.
- The `mysql_create_system_tables` script was removed because `mysql_install_db` no longer uses it.
- Renamed the `old_mode` system variable to `old`.

Bugs Fixed

- **Security Fix:** The requirement of the `DROP` privilege for `RENAME TABLE` was not enforced. (Bug #27515, CVE-2007-2691)
- **Security Fix:** If a stored routine was declared using `SQL SECURITY INVOKER`, a user who invoked the routine could gain privileges. (Bug #27337, CVE-2007-2692)
- **Security Fix:** A user with only the `ALTER` privilege on a partitioned table could obtain information about the table that should require the `SELECT` privilege. (Bug #23675, CVE-2007-2693)
- **MySQL Cluster; Replication:** (Replication): An `UPDATE` on the master became a `DELETE` on slaves. (Bug #27378)
- **MySQL Cluster:** The cluster waited 30 seconds instead of 30 milliseconds before reading table statistics. (Bug #28093)
- **MySQL Cluster:** Under certain rare circumstances, `ndbd` could get caught in an infinite loop when one transaction took a read lock and then a second transaction attempted to obtain a write lock on the same tuple in the lock queue. (Bug #28073)
- **MySQL Cluster:** Under some circumstances, a node restart could fail to update the Global Checkpoint Index (GCI). (Bug #28023)
- **MySQL Cluster:** `INSERT IGNORE` wrongly ignored `NULL` values in unique indexes. (Bug #27980)
- **MySQL Cluster:** The name of the month “March” was given incorrectly in the cluster error log. (Bug #27926)
- **MySQL Cluster:** `NDB` tables having `MEDIUMINT AUTO_INCREMENT` columns were not restored correctly by `ndb_restore`, causing spurious duplicate key errors. This issue did not affect `TINYINT`, `INT`, or `BIGINT` columns with `AUTO_INCREMENT`. (Bug #27775)
- **MySQL Cluster:** `NDB` tables with indexes whose names contained space characters were not restored correctly by `ndb_restore` (the index names were truncated). (Bug #27758)
- **MySQL Cluster:** An `INSERT` followed by a delete `DELETE` on the same `NDB` table caused a memory leak. (Bug #27756)

References: This issue is a regression of: Bug #20612.

- **MySQL Cluster:** It was not possible to add a unique index to an `NDB` table while in single user mode. (Bug #27710)
- **MySQL Cluster:** Under certain rare circumstances performing a `DROP TABLE` or `TRUNCATE TABLE` on an `NDB` table could cause a node failure or forced cluster shutdown. (Bug #27581)
- **MySQL Cluster:** Memory usage of a `mysqld` process grew even while idle. (Bug #27560)
- **MySQL Cluster:** Using more than 16GB for `DataMemory` caused problems with variable-size columns. (Bug #27512)
- **MySQL Cluster:** A data node failing while another data node was restarting could leave the cluster in an inconsistent state. In certain rare cases, this could lead to a race condition and the eventual forced shutdown of the cluster. (Bug #27466)
- **MySQL Cluster:** When using the `MemReportFrequency` configuration parameter to generate periodic reports of memory usage in the cluster log, `DataMemory` usage was not always reported for all data nodes. (Bug #27444)
- **MySQL Cluster:** When trying to create an `NDB` table after the server was started with `--ndbcluster` but without `--ndb-connectstring`, `mysqld` produced a memory allocation error. (Bug #27359)
- **MySQL Cluster:** Performing a delete followed by an insert during a local checkpoint could cause a `Rowid already allocated` error. (Bug #27205)
- **MySQL Cluster:** In an `NDB` table having a `TIMESTAMP` column using `DEFAULT CURRENT_TIMESTAMP`, that column would assume a random value when another column in the same row was updated. (Bug #27127)
- **MySQL Cluster:** Error messages displayed when running in single user mode were inconsistent. (Bug #27021)
- **MySQL Cluster:** On Solaris, the value of an `NDB` table column declared as `BIT(33)` was always displayed as `0`. (Bug #26986)
- **MySQL Cluster:** Performing `ALTER TABLE ... ENGINE=MERGE` on an `NDB` table caused `mysqld` to crash. (Bug #26898)
- **MySQL Cluster:** The `NDBCLUSTER` table handler did not set bits in null bytes correctly. (Bug #26591)
- **MySQL Cluster:** In some cases, `AFTER UPDATE` and `AFTER DELETE` triggers on `NDB` tables that referenced subject table did not see the results of operation which caused invocation of the trigger, but rather saw the row as it was prior to the update or delete operation.

This was most noticeable when an update operation used a subquery to obtain the rows to be updated. An example would be `UPDATE tbl1 SET col2 = val1 WHERE tbl1.col1 IN (SELECT col3 FROM tbl2 WHERE c4 = val2)` where there was an `AFTER UPDATE` trigger on table `tbl1`. In such cases, the trigger failed to execute.

The problem occurred because the actual update or delete operations were deferred to be able to perform them later as one batch. The fix for this bug solves the problem by disabling this optimization for a given update or delete if the table has an `AFTER` trigger defined for this operation. (Bug #26242)

- **MySQL Cluster:** Joins on multiple tables containing `BLOB` columns could cause data nodes run out of memory, and to crash with the error `NdbObjectIdMap::expand unable to expand`. (Bug #26176)
- **MySQL Cluster:** `START BACKUP NOWAIT` caused a spurious `Out of backup record` error in the management client (`START BACKUP` and `START BACKUP WAIT STARTED` performed normally). (Bug #25446)

- **MySQL Cluster:** Adding of indexes online failed for `NDB` tables having `BLOB` or `TEXT` columns. (Bug #25431)
- **MySQL Cluster:** When a cluster data node suffered a “hard” failure (such as a power failure or loss of a network connection) TCP sockets to the missing node were maintained indefinitely. Now socket-based transporters check for a response and terminate the socket if there is no activity on the socket after 2 hours. (Bug #24793)
- **MySQL Cluster:** The `ndb_resize.pl` utility did not calculate memory usage for indexes correctly. (Bug #24229)
- **MySQL Cluster:** While a data node was stopped, dropping a table then creating an index on a different table caused that node to fail during restart. This was due to the re-use of the dropped table's internal ID for the index without verifying that the index now referred to a different database object. (Bug #21755)
- **MySQL Cluster:** When trying to create tables on an SQL node not connected to the cluster, a misleading error message `Table 'tbl_name' already exists` was generated. The error now generated is `Could not connect to storage engine`. (Bug #11217, Bug #18676)
- **Replication; Cluster Replication:** An SQL node acting as a replication master server could be a single point of failure; that is, if it failed, the replication slave had no way of knowing this, which could result in a mismatch of data between the master and the slave. (Bug #21494)
- **Replication:** Out-of-memory errors were not reported. Now they are written to the error log. (Bug #26844)
- **Replication:** Improved out-of-memory detection when sending logs from a master server to slaves, and log a message when allocation fails. (Bug #26837)
- **Replication:** Aborting a statement on the master that applied to a nontransactional statement broke replication. The statement was written to the binary log but not completely executed on the master. Slaves receiving the statement executed it completely, resulting in loss of data synchrony. Now an error code is written to the error log so that the slaves stop without executing the aborted statement. (That is, replication stops, but synchrony to the point of the stop is preserved and you can investigate the problem.) (Bug #26551)
- **Replication:** When `RAND()` was called multiple times inside a stored procedure, the server did not write the correct random seed values to the binary log, resulting in incorrect replication. (Bug #25543)
- **Replication:** `GRANT` statements were not replicated if the server was started with the `--replicate-ignore-table` or `--replicate-wild-ignore-table` option. (Bug #25482)
- **Replication:** Restoration of the default database after stored routine or trigger execution on a slave could cause replication to stop if the database no longer existed. (Bug #25082)
- **Replication:** If a rotate event occurred in the middle of a nontransaction group, the group position would be updated by the rotate event indicating an illegal group start position that was effectively inside a group. This can happen if, for example, a rotate occurs between an `Intvar` event and the associated `Query` event, or between the table map events and the rows events when using row-based replication. (Bug #23171)
- **Replication:** Row-based replication of `MyISAM` to non-`MyISAM` tables did not work correctly for `BIT` columns. This has been corrected, but the fix introduces an incompatibility into the binary log format. (The incompatibility is corrected by the fix for Bug #27779.) (Bug #22583)
- **Disk Data; Cluster Replication:** An issue with replication of Disk Data tables could in some cases lead to node failure. (Bug #28161)

- **Disk Data:** Changes to a Disk Data table made as part of a transaction could not be seen by the client performing the changes until the transaction had been committed. (Bug #27757)
- **Disk Data:** When in single user mode, it was possible to create log file groups and tablespaces from any SQL node connected to the cluster. (Bug #27712)
- **Disk Data:** `CREATE TABLE ... LIKE disk_data_table` created an in-memory NDB table. (Bug #25875)
- **Disk Data:** When restarting a data node following the creation of a large number of Disk Data objects (approximately 200 such objects), the cluster could not assign a node ID to the restarting node. (Bug #25741)
- **Disk Data:** Creating an excessive number of Disk Data tables (1000 or more) could cause data nodes to fail. (Bug #24951)
- **Disk Data:** Changing a column specification or issuing a `TRUNCATE TABLE` statement on a Disk Data table caused the table to become an in-memory table.

This fix supersedes an incomplete fix that was made for this issue in MySQL 5.1.15. (Bug #24667, Bug #25296)

- **Disk Data:** Setting the value of the `UNDO BUFFER SIZE` to 64K or less in a `CREATE LOGFILE GROUP` statement led to failure of cluster data nodes. (Bug #24560)
- **Disk Data:** Creating an excessive number of data files for a single tablespace caused data nodes to crash. (Bug #24521)
- **Disk Data:** It was possible to drop the last remaining datafile in a tablespace using `ALTER TABLESPACE`, even when there was still an empty table using the tablespace.

**Note**

The datafile could be not dropped if the table still contained any rows, so this bug involved no loss of data.

(Bug #21699)

- **Cluster Replication:** Some queries that updated multiple tables were not backed up correctly. (Bug #27748)
- **Cluster Replication:** It was possible for API nodes to begin interacting with the cluster subscription manager before they were fully connected to the cluster. (Bug #27728)
- **Cluster Replication:** Under very high loads, checkpoints could be read or written with checkpoint indexes out of order. (Bug #27651)
- **Cluster Replication:** Trying to replicate a large number of frequent updates with a relatively small relay log (`max-relay-log-size` set to 1M or less) could cause the slave to crash. (Bug #27529)
- **Cluster Replication:** Setting `sql_log_bin` to zero did not disable binary logging.

This issue affected only the NDB storage engine. (Bug #27076)

- **Cluster API:** For BLOB reads on operations with lock mode `LM_CommittedRead`, the lock mode was not upgraded to `LM_Read` before the state of the BLOB had already been calculated. The NDB API methods affected by this problem included the following:
 - `NdbOperation::readTuple()`

- `NdbScanOperation::readTuples()`
 - `NdbIndexScanOperation::readTuples()`
- (Bug #27320)
- **Cluster API:** Using `NdbBlob::writeData()` to write data in the middle of an existing blob value (that is, updating the value) could overwrite some data past the end of the data to be changed. (Bug #27018)
 - A performance degradation was observed for outer join queries to which a not-exists optimization was applied. (Bug #28188)
 - `SELECT * INTO OUTFILE ... FROM INFORMATION_SCHEMA.SCHEMATA` failed with an `Access denied` error, even for a user who had the `FILE` privilege. (Bug #28181)
 - Early `NULL`-filtering optimization did not work for `eq_ref` table access. (Bug #27939)
 - Nongrouped columns were permitted by `*` in `ONLY_FULL_GROUP_BY` SQL mode. (Bug #27874)
 - Some equi-joins containing a `WHERE` clause that included a `NOT IN` subquery caused a server crash. (Bug #27870)
 - An error message suggested the use of `mysql_fix_privilege_tables` after an upgrade, but the recommended program is now `mysql_upgrade`. (Bug #27818)
 - Debug builds on Windows generated false alarms about uninitialized variables with some Visual Studio runtime libraries. (Bug #27811)
 - Certain queries that used uncorrelated scalar subqueries caused `EXPLAIN` to crash. (Bug #27807)
 - Performing a `UNION` on two views that had `ORDER BY` clauses resulted in an `Unknown column` error. (Bug #27786)
 - `mysql_install_db` is supposed to detect existing system tables and create only those that do not exist. Instead, it was exiting with an error if tables already existed. (Bug #27783)
 - The `LEAST()` and `GREATEST()` functions compared `DATE` and `DATETIME` values as strings, which in some cases could lead to an incorrect result. (Bug #27759)
 - A memory leak in the event scheduler was uncovered by Valgrind. (Bug #27733)
 - `mysqld` did not check the length of option values and could crash with a buffer overflow for long values. (Bug #27715)
 - Comparisons using row constructors could fail for rows containing `NULL` values. (Bug #27704)
 - `mysqldump` could not connect using SSL. (Bug #27669)
 - `SELECT DISTINCT` could return incorrect results if the select list contained duplicated columns. (Bug #27659)
 - On Linux, the server could not create temporary tables if `lower_case_table_names` was set to 1 and the value of `tmpdir` was a directory name containing any uppercase letters. (Bug #27653)
 - For `InnoDB` tables, a multiple-row `INSERT` of the form `INSERT INTO t (id...) VALUES (NULL...) ON DUPLICATE KEY UPDATE id=VALUES(id)`, where `id` is an `AUTO_INCREMENT` column, could cause `ERROR 1062 (23000): Duplicate entry...` errors or lost rows. (Bug #27650)

- When MySQL logged slow query information to a `CSV` table, it used an incorrect formula to calculate the `query_time` and `lock_time` values. (Bug #27638)
- The XML output representing an empty result was an empty string rather than an empty `<resultset/>` element. (Bug #27608)
- Comparison of a `DATE` with a `DATETIME` did not treat the `DATE` as having a time part of `00:00:00`. (Bug #27590)

References: See also: Bug #32198.

- With `NO_AUTO_VALUE_ON_ZERO` SQL mode enabled, `LOAD DATA` operations could assign incorrect `AUTO_INCREMENT` values. (Bug #27586)
- Group relay log rotation updated only the log position and not the name, causing the slave to stop. (Bug #27583)
- Incorrect results could be returned for some queries that contained a select list expression with `IN` or `BETWEEN` together with an `ORDER BY` or `GROUP BY` on the same expression using `NOT IN` or `NOT BETWEEN`. (Bug #27532)
- The fix for Bug #17212 provided correct sort order for misordered output of certain queries, but caused significant overall query performance degradation. (Results were correct (good), but returned much more slowly (bad).) The fix also affected performance of queries for which results were correct. The performance degradation has been addressed. (Bug #27531)

References: This issue is a regression of: Bug #17212.

- The `CRC32()` function returns an unsigned integer, but the metadata was signed, which could cause certain queries to return incorrect results. (For example, queries that selected a `CRC32()` value and used that value in the `GROUP BY` clause.) (Bug #27530)
- An interaction between `SHOW TABLE STATUS` and other concurrent statements that modify the table could result in a divide-by-zero error and a server crash. (Bug #27516)
- Evaluation of an `IN()` predicate containing a decimal-valued argument caused a server crash. (Bug #27513, Bug #27362, CVE-2007-2583)
- A race condition between `DROP TABLE` and `SHOW TABLE STATUS` could cause the latter to display incorrect information. (Bug #27499)
- In out-of-memory conditions, the server might crash or otherwise not report an error to the Windows event log. (Bug #27490)
- Passing nested row expressions with different structures to an `IN` predicate caused a server crash. (Bug #27484)
- The `decimal.h` header file was incorrectly omitted from binary distributions. (Bug #27456)
- With `innodb_file_per_table` enabled, attempting to rename an `InnoDB` table to a nonexistent database caused the server to exit. (Bug #27381)
- Nested aggregate functions could be improperly evaluated. (Bug #27363)
- A stored function invocation in the `WHERE` clause was treated as a constant. (Bug #27354)
- For the `INFORMATION_SCHEMA.SESSION_STATUS` and `GLOBAL_STATUS` tables, some status values were incorrectly converted to the data type of the `VARIABLE_VALUE` column. (Bug #27327)

- Failure to allocate memory associated with `transaction_prealloc_size` could cause a server crash. (Bug #27322)
- A subquery could get incorrect values for references to outer query columns when it contained aggregate functions that were aggregated in outer context. (Bug #27321)
- The server did not shut down cleanly. (Bug #27310)
- In a view, a column that was defined using a `GEOMETRY` function was treated as having the `LONGBLOB` data type rather than the `GEOMETRY` type. (Bug #27300)
- `mysqldump` crashed if it got no data from `SHOW CREATE PROCEDURE` (for example, when trying to dump a routine defined by a different user and for which the current user had no privileges). Now it prints a comment to indicate the problem. It also returns an error, or continues if the `--force` option is given. (Bug #27293)
- Queries containing subqueries with `COUNT(*)` aggregated in an outer context returned incorrect results. This happened only if the subquery did not contain any references to outer columns. (Bug #27257)
- Use of an aggregate function from an outer context as an argument to `GROUP_CONCAT()` caused a server crash. (Bug #27229)
- String truncation upon insertion into an integer or year column did not generate a warning (or an error in strict mode). (Bug #27176, Bug #26359)
- `mysqlbinlog` produced different output with the `-R` option than without it. (Bug #27171)
- Storing `NULL` values in spatial fields caused excessive memory allocation and crashes on some systems. (Bug #27164)
- Row equalities in `WHERE` clauses could cause memory corruption. (Bug #27154)
- `ON DUPLICATE KEY UPDATE` failed for a table partitioned by `KEY` on a primary key `VARCHAR` column. (Bug #27123)
- `GROUP BY` on a `ucs2` column caused a server crash when there was at least one empty string in the column. (Bug #27079)
- Duplicate members in `SET` or `ENUM` definitions were not detected. Now they result in a warning; if strict SQL mode is enabled, an error occurs instead. (Bug #27069)
- For `FEDERATED` tables, `SHOW CREATE TABLE` could fail when the table name was longer than the connection name. (Bug #27036)
- `mysql_install_db` could terminate with an error after failing to determine that a system table already existed. (Bug #27022)
- In a `MEMORY` table, using a `BTREE` index to scan for updatable rows could lead to an infinite loop. (Bug #26996)
- `make_win_bin_dist` neglected to copy some required `MyISAM` table files. (Bug #26922)
- For `InnoDB` tables having a clustered index that began with a `CHAR` or `VARCHAR` column, deleting a record and then inserting another before the deleted record was purged could result in table corruption. (Bug #26835)
- `mysqldump` would not dump a view for which the `DEFINER` no longer exists. (Bug #26817)
- Duplicates were not properly identified among (potentially) long strings used as arguments for `GROUP_CONCAT(DISTINCT)`. (Bug #26815)

- `ALTER VIEW` requires the `CREATE VIEW` and `DROP` privileges for the view. However, if the view was created by another user, the server erroneously required the `SUPER` privilege. (Bug #26813)
- If the name of a table given to `myisamchk -rq` was a packed table and the name included the `.MYI` extension, `myisamchk` incorrectly created a file with a `.MYI.MYI` extension. (Bug #26782)
- Creating a temporary table with `InnoDB` when using the one-file-per-table setting, and when the host file system for temporary tables was `tmpfs`, would cause an assertion within `mysqld`. This was due to the use of `O_DIRECT` when opening the temporary table file. (Bug #26662)
- `mysql_upgrade` did not detect failure of external commands that it runs. (Bug #26639)
- The range optimizer could cause the server to run out of memory. (Bug #26625)
- The range optimizer could consume a combinatorial amount of memory for certain classes of `WHERE` clauses. (Bug #26624)
- `mysqldump` could crash or exhibit incorrect behavior when some options were given very long values, such as `--fields-terminated-by="some very long string"`. The code has been cleaned up to remove a number of fixed-sized buffers and to be more careful about error conditions in memory allocation. (Bug #26346)
- A possible buffer overflow in `SHOW PROCEDURE CODE` was removed. (Bug #26303)
- The `FEDERATED` engine did not permit the local and remote tables to have different names. (Bug #26257)
- The temporary file-creation code was cleaned up on Windows to improve server stability. (Bug #26233)
- For `MyISAM` tables, `COUNT(*)` could return an incorrect value if the `WHERE` clause compared an indexed `TEXT` column to the empty string (`' '`). This happened if the column contained empty strings and also strings starting with control characters such as tab or newline. (Bug #26231)
- For `INSERT INTO ... SELECT` where index searches used column prefixes, insert errors could occur when key value type conversion was done. (Bug #26207)
- `mysqlbinlog --base64-output` produced invalid SQL. (Bug #26194)
- For `DELETE FROM tbl_name ORDER BY col_name` (with no `WHERE` or `LIMIT` clause), the server did not check whether `col_name` was a valid column in the table. (Bug #26186)
- Executing an `INSERT ... SELECT ... FROM INFORMATION_SCHEMA.GLOBAL_STATUS` statement from within an event caused a server crash. (Bug #26174)
- `mysqldump` could not dump log tables. (Bug #26121)
- On Windows, trying to use backslash (`\`) characters in paths for `DATA DIRECTORY` and `INDEX DIRECTORY` when creating partitioned tables caused MySQL to crash.

(You must use `/` characters when specifying paths for these options, regardless of platform. See [Overview of Partitioning in MySQL](#), for an example using absolute paths for `DATA DIRECTORY` and `INDEX DIRECTORY` when creating a partitioned table on Windows.) (Bug #26074, Bug #25141)
- `mysqldump` crashed for `MERGE` tables if the `--complete-insert (-c)` option was given. (Bug #25993)
- Index hints (`USE INDEX`, `IGNORE INDEX`, `FORCE INDEX`) cannot be used with `FULLTEXT` indexes, but were not being ignored. (Bug #25951)
- Setting a column to `NOT NULL` with an `ON DELETE SET NULL` clause foreign key crashes the server. (Bug #25927)

- Corrupted `MyISAM` tables that have different definitions in the `.frm` and `.MYI` tables might cause a server crash. (Bug #25908)
- If `CREATE TABLE t1 LIKE t2` failed due to a full disk, an empty `t2.frm` file could be created but not removed. This file then caused subsequent attempts to create a table named `t2` to fail. This is easily corrected at the file system level by removing the `t2.frm` file manually, but now the server removes the file if the create operation does not complete successfully. (Bug #25761)
- In certain situations, `MATCH ... AGAINST` returned false hits for `NULL` values produced by `LEFT JOIN` when no full-text index was available. (Bug #25729)
- Concurrent `CREATE SERVER` and `ALTER SERVER` statements could cause a deadlock. (Bug #25721)
- `CREATE SERVER`, `DROP SERVER`, and `ALTER SERVER` did not require any privileges. Now these statements require the `SUPER` privilege. (Bug #25671)
- On Windows, connection handlers did not properly decrement the server's thread count when exiting. (Bug #25621)
- `OPTIMIZE TABLE` might fail on Windows when it attempts to rename a temporary file to the original name if the original file had been opened, resulting in loss of the `.MYD` file. (Bug #25521)
- For `SHOW ENGINE INNODB STATUS`, the `LATEST DEADLOCK INFORMATION` was not always cleared properly. (Bug #25494)
- `mysql_stmt_fetch()` did an invalid memory deallocation when used with the embedded server. (Bug #25492)
- `mysql_upgrade` did not pass a password to `mysqlcheck` if one was given. (Bug #25452)
- On Windows, `mysql_upgrade` was sensitive to lettercase of the names of some required components. (Bug #25405)
- During a call to `mysql_change_user()`, when authentication fails or the database to change to is unknown, a subsequent call to any function that does network communication leads to packets out of order. This problem was introduced in MySQL 5.1.14. (Bug #25371)
- Difficult repair or optimization operations could cause an assertion failure, resulting in a server crash. (Bug #25289)
- For storage engines that permit the current auto-increment value to be set, using `ALTER TABLE ... ENGINE` to convert a table from one such storage engine to another caused loss of the current value. (For storage engines that do not support setting the value, it cannot be retained anyway when changing the storage engine.) (Bug #25262)
- Duplicate entries were not assessed correctly in a `MEMORY` table with a `BTREE` primary key on a `utf8 ENUM` column. (Bug #24985)
- Several math functions produced incorrect results for large unsigned values. `ROUND()` produced incorrect results or a crash for a large number-of-decimals argument. (Bug #24912)
- The result set of a query that used `WITH ROLLUP` and `DISTINCT` could lack some rollup rows (rows with `NULL` values for grouping attributes) if the `GROUP BY` list contained constant expressions. (Bug #24856)
- Selecting the result of `AVG()` within a `UNION` could produce incorrect values. (Bug #24791)
- For queries that used `ORDER BY` with `InnoDB` tables, if the optimizer chose an index for accessing the table but found a covering index that enabled the `ORDER BY` to be skipped, no results were returned. (Bug #24778)

- The `NO_DIR_IN_CREATE` server SQL mode was not enforced for partitioned tables. (Bug #24633)
- `MBRDisjoint()`, `MBRequal()`, `MBRIntersects()`, `MBROverlaps()`, `MBRTouches()`, and `MBRWithin()` were inadvertently omitted from recent versions of MySQL (5.1.14 to 5.1.17). (Bug #24588)
- Access through `my_pread()` or `my_pwrite()` to table files larger than 2GB could fail on some systems. (Bug #24566)
- `MBROverlaps()` returned incorrect values in some cases. (Bug #24563)
- A problem in handling of aggregate functions in subqueries caused predicates containing aggregate functions to be ignored during query execution. (Bug #24484)
- The `MERGE` storage engine could return incorrect results when several index values that compare equality were present in an index (for example, `'gross'` and `'gross '`, which are considered equal but have different lengths). (Bug #24342)
- Some upgrade problems are detected and better error messages suggesting that `mysql_upgrade` be run are produced. (Bug #24248)
- The test for the `MYSQL_OPT_SSL_VERIFY_SERVER_CERT` option for `mysql_options()` was performed incorrectly. Also changed as a result of this bug fix: The `arg` option for the `mysql_options()` C API function was changed from `char *` to `void *`. (Bug #24121)
- Some views could not be created even when the user had the requisite privileges. (Bug #24040)
- The values displayed for the `Innodb_row_lock_time`, `Innodb_row_lock_time_avg`, and `Innodb_row_lock_time_max` status variables were incorrect. (Bug #23666)
- Using `CAST()` to convert `DATETIME` values to numeric values did not work. (Bug #23656)
- A damaged or missing `mysql.event` table caused `SHOW VARIABLES` to fail. (Bug #23631)
- `SHOW CREATE VIEW` qualified references to stored functions in the view definition with the function's database name, even when the database was the default database. This affected `mysqldump` (which uses `SHOW CREATE VIEW` to dump views) because the resulting dump file could not be used to reload the database into a different database. `SHOW CREATE VIEW` now suppresses the database name for references to stored functions in the default database. (Bug #23491)
- An `INTO OUTFILE` clause is permitted only for the final `SELECT` of a `UNION`, but this restriction was not being enforced correctly. (Bug #23345)
- The `AUTO_INCREMENT` value would not be correctly reported for `InnoDB` tables when using `SHOW CREATE TABLE` statement or `mysqldump` command. (Bug #23313)
- With the `NO_AUTO_VALUE_ON_ZERO` SQL mode enabled, `LAST_INSERT_ID()` could return 0 after `INSERT ... ON DUPLICATE KEY UPDATE`. Additionally, the next rows inserted (by the same `INSERT`, or the following `INSERT` with or without `ON DUPLICATE KEY UPDATE`), would insert 0 for the auto-generated value if the value for the `AUTO_INCREMENT` column was `NULL` or missing. (Bug #23233)
- Implicit conversion of `9912101` to `DATE` did not match `CAST(9912101 AS DATE)`. (Bug #23093)
- `SELECT COUNT(*)` from a table containing a `DATETIME NOT NULL` column could produce spurious warnings with the `NO_ZERO_DATE` SQL mode enabled. (Bug #22824)
- Using `SET GLOBAL` to change the `lc_time_names` system variable had no effect on new connections. (Bug #22648)

- `SOUNDEX()` returned an invalid string for international characters in multibyte character sets. (Bug #22638)
- A multiple-table `UPDATE` could return an incorrect rows-matched value if, during insertion of rows into a temporary table, the table had to be converted from a `MEMORY` table to a `MyISAM` table. (Bug #22364)
- `COUNT(decimal_expr)` sometimes generated a spurious truncation warning. (Bug #21976)
- `yaSSL` crashed on pre-Pentium Intel CPUs. (Bug #21765)
- A slave that used `--master-ssl-cipher` could not connect to the master. (Bug #21611)
- Database and table names have a maximum length of 64 characters (even if they contain multibyte characters), but were truncated to 64 bytes.

**Note**

This improves on a previous fix made for this bug in MySQL 5.1.12.

(Bug #21432)

- `InnoDB`: The first read statement, if served from the query cache, was not consistent with the `READ COMMITTED` isolation level. (Bug #21409)
- On Windows, if the server was installed as a service, it did not auto-detect the location of the data directory. (Bug #20376)
- Changing a `utf8` column in an `InnoDB` table to a shorter length did not shorten the data values. (Bug #20095)
- In some cases, the optimizer preferred a range or full index scan access method over lookup access methods when the latter were much cheaper. (Bug #19372)
- Conversion of `DATETIME` values in numeric contexts sometimes did not produce a double (`YYYYMMDDHHMMSS.uuuuuu`) value. (Bug #16546)
- `INSERT...ON DUPLICATE KEY UPDATE` could cause `Error 1032: Can't find record in ...` for inserts into an `InnoDB` table unique index using key column prefixes with an underlying `utf8` string column. (Bug #13191)
- Having the `EXECUTE` privilege for a routine in a database should make it possible to `USE` that database, but the server returned an error instead. This has been corrected. As a result of the change, `SHOW TABLES` for a database in which you have only the `EXECUTE` privilege returns an empty set rather than an error. (Bug #9504)

Changes in MySQL 5.1.17 (2007-04-04)

This is a new Beta development release, fixing recently discovered bugs.

**Note**

This Beta 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.

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 Enterprise* (a commercial MySQL offering). For more details, please see (<http://www.mysql.com/products/enterprise>).

- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

Functionality Added or Changed

- **Incompatible Change:** Scheduled events now use the session time zone that is current when a `CREATE EVENT` or `ALTER EVENT` statement executes is used to interpret times specified in the event definition (rather than UTC as in previous releases). The session time zone becomes the event time zone; that is, the time zone that is used for event scheduling and is in effect within the event as it executes. Because of this change, scheduled event metadata now includes time zone information, which can be seen in the `TIME_ZONE` column of the `INFORMATION_SCHEMA.EVENTS` table and the `Time zone` column in the output of the `SHOW EVENTS` statement. These columns have been added in this release, along with a `time_zone` column in the `mysql.event` table. Due to these changes, events created in previous versions of MySQL cannot be created, viewed, or used until `mysql.event` has been upgraded.

For retrievals from `INFORMATION_SCHEMA.EVENTS` or `SHOW EVENTS`, times previously displayed using UTC now use the event time zone. (Bug #16420)

- **Important Change; Replication:** The following options for controlling replication master configuration on a slave are now deprecated.

- `--master-host`
- `--master-user`
- `--master-password`
- `--master-port`
- `--master-connect-retry`
- `--master-ssl`
- `--master-ssl-ca`
- `--master-ssl-capath`
- `--master-ssl-cert`
- `--master-ssl-cipher`
- `--master-ssl-key`

To change the master configuration on a slave you should use the `CHANGE MASTER TO` statement.

References: See also: Bug #21490.

- **Important Change:** The `CREATE EVENT` and `ALTER EVENT` statements now support a `DEFINER` clause, similar to that used in the `CREATE TRIGGER` statement.

See [CREATE EVENT Syntax](#), for detailed information. (Bug #16425)

- **MySQL Cluster:** Added the `--skip-table-check` option (short form `-s`) for `ndb_restore`, which causes the restoration process to ignore any changes that may have occurred in table schemas after the backup was made. Previously, this was the default behavior.

See [ndb_restore — Restore a MySQL Cluster Backup](#), for more information. (Bug #24363)

- **MySQL Cluster:** For `ALTER TABLE` statements, `ADD INDEX` and `DROP INDEX` operations for dynamic (variable-width) columns on NDB tables are now performed as online operations (no table copying). This is also true for `CREATE INDEX` and `DROP INDEX`. Renaming of NDB and MyISAM tables and of columns in such tables is now performed in place without copying or locking the tables. As a result, these operations are now performed much more quickly than previously.

For more information, see [ALTER TABLE Online Operations in MySQL Cluster](#), [CREATE INDEX Syntax](#), and [DROP INDEX Syntax](#).

- Added a `--no-beep` option to `mysqladmin`. It suppresses the warning beep that is emitted by default for errors such as a failure to connect to the server. (Bug #26964)
- Added the `--service-startup-timeout` option for `mysql.server` to specify how long to wait for the server to start. If the server does not start within the timeout period, `mysql.server` exits with an error. (Bug #26952)
- Prefix lengths for columns in `SPATIAL` indexes can no longer be specified. For tables created in older versions of MySQL that have `SPATIAL` indexes containing prefixed columns, dumping and reloading the table causes the indexes to be created with no prefixes. (The full column width of each column is indexed.) (Bug #26794)
- Added the `innodb_stats_on_metadata` system variable to enable control over whether InnoDB performs statistics gathering when metadata statements are executed. See [InnoDB Startup Options and System Variables](#). (Bug #26598)
- Statements that affect `mysql` database tables now are written to the binary log using the following rules:
 - Data manipulation statements such as `INSERT` that change data in `mysql` database tables directly are logged according to the settings of the `binlog_format` system variable.
 - Statements such as `GRANT` that change the `mysql` database indirectly are logged as statements regardless of the value of `binlog_format`.

For more details, see [Logging Format for Changes to mysql Database Tables](#). (Bug #25091)

- The server now includes a timestamp in error messages that are logged as a result of unhandled signals (such as `mysqld got signal 11` messages). (Bug #24878)
- The syntax for index hints has been extended to enable more fine-grained control over the optimizer's selection of an execution plan for various phases of query processing. See [Index Hints](#). (Bug #21174)
- Added the `--secure-file-priv` option for `mysqld`, which limits the effect of the `LOAD_FILE()` function and the `LOAD DATA` and `SELECT ... INTO OUTFILE` statements to work only with files in a given directory. (Bug #18628)
- Added the `thread_handling` system variable to control whether the server use a single thread or one thread per connection. The `--one-thread` option now is deprecated; use `--thread_handling=one-thread` instead.

- Statements such as `GRANT` that change the `mysql` database indirectly are logged as statements regardless of the value of `binlog_format`.
- Added the read-only `hostname` system variable, which the server sets at startup to the server host name.
- Data manipulation statements such as `INSERT` that change data in `mysql` database tables directly are logged according to the settings of the `binlog_format` system variable.
- Added the `old_mode` system variable to cause the server to revert to certain behaviors present in older versions. Currently, this variable affects handling of index hints. See [Index Hints](#).
- Prepared statements now use the query cache under the conditions described in [How the Query Cache Operates](#). (Bug #735)

Bugs Fixed

- **Incompatible Change:** `INSERT DELAYED` statements are not supported for `MERGE` tables, but the `MERGE` storage engine was not rejecting such statements, resulting in table corruption. Applications previously using `INSERT DELAYED` into `MERGE` table will break when upgrading to versions with this fix. To avoid the problem, remove `DELAYED` from such statements. (Bug #26464)

- **Important Note:** The parser accepted invalid code in SQL condition handlers, leading to server crashes or unexpected execution behavior in stored programs. Specifically, the parser permitted a condition handler to refer to labels for blocks that enclose the handler declaration. This was incorrect because block label scope does not include the code for handlers declared within the labeled block.

The parser now rejects this invalid construct, but if you perform a binary upgrade (without dumping and reloading your databases), existing handlers that contain the construct are still invalid and should be rewritten *even if they appear to function as you expect*.

To find affected handlers, use `mysqldump` to dump all stored procedures and functions, triggers, and events. Then attempt to reload them into an upgraded server. Handlers that contain illegal label references will be rejected.

For more information about condition handlers and writing them to avoid invalid jumps, see [DECLARE ... HANDLER Syntax](#). (Bug #26503)

- **MySQL Cluster:** It was not possible to set `LockPagesInMainMemory` equal to 0. (Bug #27291)
- **MySQL Cluster:** A race condition could sometimes occur if the node acting as master failed while node IDs were still being allocated during startup. (Bug #27286)
- **MySQL Cluster:** When a data node was taking over as the master node, a race condition could sometimes occur as the node was assuming responsibility for handling of global checkpoints. (Bug #27283)
- **MySQL Cluster:** After putting the cluster in single user mode from one MySQL server, trying to drop an `NDB` table from a second MySQL server also connected to the cluster would cause the second MySQL server to hang. (Bug #27254)
- **MySQL Cluster:** `mysqld` could crash shortly after a data node failure following certain DML operations. (Bug #27169)
- **MySQL Cluster:** (Disk Data): Under some circumstances, a data node could fail during restart while flushing Disk Data UNDO logs. (Bug #27102)
- **MySQL Cluster:** The same failed request from an API node could be handled by the cluster multiple times, resulting in reduced performance. (Bug #27087)

- **MySQL Cluster:** The failure of a data node while restarting could cause other data nodes to hang or crash. (Bug #27003)
- **MySQL Cluster:** Creating a table on one SQL node while in single user mode caused other SQL nodes to crash. (Bug #26997)
- **MySQL Cluster:** `mysqld` processes would sometimes crash under high load. (Bug #26825)
- **MySQL Cluster:** Using only the `--print_data` option (and no other options) with `ndb_restore` caused `ndb_restore` to fail. (Bug #26741)

References: This issue is a regression of: Bug #14612.

- **MySQL Cluster:** The output from `ndb_restore --print_data` was incorrect for a backup made of a database containing tables with `TINYINT` or `SMALLINT` columns. (Bug #26740)
- **MySQL Cluster:** An infinite loop in an internal logging function could cause trace logs to fill up with `Unknown Signal type` error messages and thus grow to unreasonable sizes. (Bug #26720)
- **MySQL Cluster:** An invalid pointer was returned following a `FSCLOSECONF` signal when accessing the REDO logs during a node restart or system restart. (Bug #26515)
- **MySQL Cluster:** The management client command `node_id STATUS` displayed the message `Node node_id: not connected` when `node_id` was not the node ID of a data node.

**Note**

The `ALL STATUS` command in the cluster management client still displays status information for data nodes only. This is by design. See [Commands in the MySQL Cluster Management Client](#), for more information.

(Bug #21715)

- **MySQL Cluster:** When performing an upgrade or downgrade, no specific error information was made available when trying to upgrade data nodes or SQL nodes before upgrading management nodes. (Bug #21296)
- **MySQL Cluster:** Some values of `MaxNoOfTables` caused the error `Job buffer congestion` to occur. (Bug #19378)
- **Replication:** A multiple-row delayed insert with an auto-increment column could cause duplicate entries to be created on the slave in a replication environment. (Bug #26116, Bug #25507)
- **Replication:** Duplicating the usage of a user variable in a stored procedure or trigger would not be replicated correctly to the slave. (Bug #25167)
- **Replication:** `DROP TRIGGER` statements would not be filtered on the slave when using the `replication-wild-do-table` option. (Bug #24478)
- **Replication:** For `INSERT ... ON DUPLICATE KEY UPDATE` statements where some `AUTO_INCREMENT` values were generated automatically for inserts and some rows were updated, one auto-generated value was lost per updated row, leading to faster exhaustion of the range of the `AUTO_INCREMENT` column.

Because the original problem can affect replication (different values on master and slave), it is recommended that the master and its slaves be upgraded to the current version. (Bug #24432)

- **Replication:** Replication between master and slave would infinitely retry binary log transmission where the `max_allowed_packet` on the master was larger than that on the slave if the size of the transfer was between these two values. (Bug #23775)
- **Replication:** Loading data using `LOAD DATA INFILE` may not replicate correctly (due to character set incompatibilities) if the `character_set_database` variable is set before the data is loaded. (Bug #15126)
- **Replication:** User defined variables used within stored procedures and triggers are not replicated correctly when operating in statement-based replication mode. (Bug #14914, Bug #20141)
- **Disk Data:** A memory overflow could occur with tables having a large amount of data stored on disk, or with queries using a very high degree of parallelism on Disk Data tables. (Bug #26514)
- **Disk Data:** Use of a tablespace whose `INITIAL_SIZE` was greater than 1 GB could cause the cluster to crash. (Bug #26487)
- **Disk Data:** Creating multiple Disk Data tables using different tablespaces could sometimes cause the cluster to fail. (Bug #25992)
- **Disk Data:** `ALTER TABLE ... ADD COLUMN ...` on a Disk Data table moved data for existing nonindexed columns from the tablespace into memory. (Bug #25880)
- **Disk Data:** `DROP INDEX` on a Disk Data table did not always move data from memory into the tablespace. (Bug #25877)
- **Disk Data:** When creating a log file group, setting `INITIAL_SIZE` to less than `UNDO_BUFFER_SIZE` caused data nodes to crash. (Bug #25743)
- **Cluster Replication:** The simultaneous failure of a data node and an SQL node could cause replication to fail. (Bug #27005)
- **Cluster API:** A delete operation using a scan followed by an insert using a scan could cause a data node to fail. (Bug #27203)
- **Cluster API:** `NAND` and `NOR` operations with `NdbScanFilter` did not perform correctly. (Bug #24568)
- **Cluster API:** You can now use the `ndb_mgm_check_connection()` function to determine whether a management server is running.
- **MyISAM** tables converted to `ARCHIVE` were excessively large. (Bug #27533)
- `SELECT ... INTO OUTFILE` with a long `FIELDS ENCLOSED BY` value could crash the server. (Bug #27231)
- An `INSERT ... ON DUPLICATE KEY UPDATE` statement might modify values in a table but not flush affected data from the query cache, causing subsequent selects to return stale results. This made the combination of query cache plus `ON DUPLICATE KEY UPDATE` very unreliable. (Bug #27210)

References: See also: Bug #27006, Bug #27033. This issue is a regression of: Bug #19978.

- For `INSERT ... ON DUPLICATE KEY UPDATE` statements on tables containing `AUTO_INCREMENT` columns, `LAST_INSERT_ID()` was reset to 0 if no rows were successfully inserted or changed. "Not changed" includes the case where a row was updated to its current values, but in that case, `LAST_INSERT_ID()` should not be reset to 0. Now `LAST_INSERT_ID()` is reset to 0 only if no rows were successfully inserted or touched, whether or not touched rows were changed. (Bug #27033)

References: See also: Bug #27210, Bug #27006. This issue is a regression of: Bug #19978.

- Invalid optimization of pushdown conditions for queries where an outer join was guaranteed to read only one row from the outer table led to results with too few rows. (Bug #26963)
- For `MERGE` tables defined on underlying tables that contained a short `VARCHAR` column (shorter than four characters), using `ALTER TABLE` on at least one but not all of the underlying tables caused the table definitions to be considered different from that of the `MERGE` table, even if the `ALTER TABLE` did not change the definition. (Bug #26881)
- Use of a subquery containing `GROUP BY` and `WITH ROLLUP` caused a server crash. (Bug #26830)
- Setting `event_scheduler = 1` or `event_scheduler = ON` caused the server to crash if the server had been started with `--skip-grant-tables`. Starting the server with `--skip-grant-tables` now causes `event_scheduler` to be set to `DISABLED` automatically, overriding any other value that may have been set. (Bug #26807)
- Added support for `--debugger=dbx` for `mysql-test-run.pl` and added support for `--debugger=devenv`, `--debugger=DevEnv`, and `--debugger=/path/to/devenv`. (Bug #26792)
- A result set column formed by concatenation of string literals was incomplete when the column was produced by a subquery in the `FROM` clause. (Bug #26738)
- SSL connections failed on Windows. (Bug #26678)
- When using the result of `SEC_TO_TIME()` for time value greater than 24 hours in an `ORDER BY` clause, either directly or through a column alias, the rows were sorted incorrectly as strings. (Bug #26672)
- Use of a subquery containing a `UNION` with an invalid `ORDER BY` clause caused a server crash. (Bug #26661)
- In some error messages, inconsistent format specifiers were used for the translations in different languages. `comp_err` (the error message compiler) now checks for mismatches. (Bug #26571)
- Views that used a scalar correlated subquery returned incorrect results. (Bug #26560)
- `UNHEX() IS NULL` comparisons failed when `UNHEX()` returned `NULL`. (Bug #26537)
- On 64-bit Windows, large timestamp values could be handled incorrectly. (Bug #26536)
- `SHOW CREATE EVENT` failed to display the `STARTS` and `ENDS` clauses for an event defined with `STARTS NOW()`, `ENDS NOW()`, or both. (Bug #26429)
- If the server was started with `--skip-grant-tables`, selecting from `INFORMATION_SCHEMA` tables caused a server crash. (Bug #26285)
- For some values of the position argument, the `INSERT()` function could insert a NUL byte into the result. (Bug #26281)
- For an `INSERT` statement that should fail due to a column with no default value not being assigned a value, the statement succeeded with no error if the column was assigned a value in an `ON DUPLICATE KEY UPDATE` clause, even if that clause was not used. (Bug #26261)
- `INSERT DELAYED` statements inserted incorrect values into `BIT` columns. (Bug #26238)
- A query of type `index_merge`, and with a `WHERE` clause having the form `WHERE indexed_column_1=value_1 OR indexed_column_2=value_2` on a partitioned table caused the server to crash. (Bug #26117)
- `BENCHMARK()` did not work correctly for expressions that produced a `DECIMAL` result. (Bug #26093)

- For `MEMORY` tables, extending the length of a `VARCHAR` column with `ALTER TABLE` might result in an unusable table. (Bug #26080)
- The server could hang during binary log rotation. (Bug #26079)
- `LOAD DATA INFILE` sent an okay to the client before writing the binary log and committing the changes to the table had finished, thus violating ACID requirements. (Bug #26050)
- `X() IS NULL` and `Y() IS NULL` comparisons failed when `X()` and `Y()` returned `NULL`. (Bug #26038)
- Indexes on `TEXT` columns were ignored when `ref` accesses were evaluated. (Bug #25971)
- If a thread previously serviced a connection that was killed, excessive memory and CPU use by the thread occurred if it later serviced a connection that had to wait for a table lock. (Bug #25966)
- `VIEW` restrictions were applied to `SELECT` statements after a `CREATE VIEW` statement failed, as though the `CREATE` had succeeded. (Bug #25897)
- Several deficiencies in resolution of column names for `INSERT ... SELECT` statements were corrected. (Bug #25831)
- Inserting `utf8` data into a `TEXT` column that used a single-byte character set could result in spurious warnings about truncated data. (Bug #25815)
- On Windows, debug builds of `mysqld` could fail with heap assertions. (Bug #25765)
- In certain cases it could happen that deleting a row corrupted an `RTREE` index. This affected indexes on spatial columns. (Bug #25673)
- Using `mysqlbinlog` on a binary log would crash if there were a large number of row-based events related to a single statement. (Bug #25628)
- Expressions involving `SUM()`, when used in an `ORDER BY` clause, could lead to out-of-order results. (Bug #25376)
- Use of a `GROUP BY` clause that referred to a stored function result together with `WITH ROLLUP` caused incorrect results. (Bug #25373)
- A stored procedure that made use of cursors failed when the procedure was invoked from a stored function. (Bug #25345)
- On Windows, the server exhibited a file-handle leak after reaching the limit on the number of open file descriptors. (Bug #25222)
- The `REPEAT()` function did not permit a column name as the `count` parameter. (Bug #25197)
- A reference to a nonexistent column in the `ORDER BY` clause of an `UPDATE ... ORDER BY` statement could cause a server crash. (Bug #25126)
- A view on a join is insertable for `INSERT` statements that store values into only one table of the join. However, inserts were being rejected if the inserted-into table was used in a self-join because MySQL incorrectly was considering the insert to modify multiple tables of the view. (Bug #25122)
- Creating a table with latin characters in the name caused the output of `SHOW FULL TABLES` to have `ERROR` for the table type. (Bug #25081)
- MySQL would not compile when configured using `--without-query-cache`. (Bug #25075)
- It was not possible to use XPath keywords as tag names for expressions used in the `ExtractValue()` function. (Bug #24747)

- Increasing the width of a `DECIMAL` column could cause column values to be changed. (Bug #24558)
- `IF(expr, unsigned_expr, unsigned_expr)` was evaluated to a signed result, not unsigned. This has been corrected. The fix also affects constructs of the form `IS [NOT] {TRUE|FALSE}`, which were transformed internally into `IF()` expressions that evaluated to a signed result.

For existing views that were defined using `IS [NOT] {TRUE|FALSE}` constructs, there is a related implication. The definitions of such views were stored using the `IF()` expression, not the original construct. This is manifest in that `SHOW CREATE VIEW` shows the transformed `IF()` expression, not the original one. Existing views will evaluate correctly after the fix, but if you want `SHOW CREATE VIEW` to display the original construct, you must drop the view and re-create it using its original definition. New views will retain the construct in their definition. (Bug #24532)

- `SHOW ENGINE MUTEX STATUS` failed to produce an `Unknown table engine` error.

See [SHOW ENGINE Syntax](#). (Bug #24392)

- A user-defined variable could be assigned an incorrect value if a temporary table was employed in obtaining the result of the query used to determine its value. (Bug #24010)
- `mysqlimport` used a variable of the wrong type for the `--use-threads` option, which could cause a crash on some architectures. (Bug #23814)
- Queries that used a temporary table for the outer query when evaluating a correlated subquery could return incorrect results. (Bug #23800)
- On Windows, debug builds of `mysqlbinlog` could fail with a memory error. (Bug #23736)
- When using certain server SQL modes, the `mysql.proc` table was not created by `mysql_install_db`. (Bug #23669)
- `DOUBLE` values such as `20070202191048.000000` were being treated as illegal arguments by `WEEK()`. (Bug #23616)
- The server could crash if two or more threads initiated query cache resize operation at moments very close in time. (Bug #23527)
- `NOW()` returned the wrong value in statements executed at server startup with the `--init-file` option. (Bug #23240)
- Setting the `slow_query_log_file` system variable caused log output to go to the general log, not the slow query log. (Bug #23225)
- When nesting stored procedures within a trigger on a table, a false dependency error was thrown when one of the nested procedures contained a `DROP TABLE` statement. (Bug #22580)
- Instance Manager did not remove the angel PID file on a clean shutdown. (Bug #22511)
- `EXPLAIN EXTENDED` did not show `WHERE` conditions that were optimized away. (Bug #22331)
- `IN ((subquery))`, `IN (((subquery)))`, and so forth, are equivalent to `IN (subquery)`, which is always interpreted as a table subquery (so that it is permitted to return more than one row). MySQL was treating the “over-parenthesized” subquery as a single-row subquery and rejecting it if it returned more than one row. This bug primarily affected automatically generated code (such as queries generated by Hibernate), because humans rarely write the over-parenthesized forms. (Bug #21904)
- An `INSERT` trigger invoking a stored routine that inserted into a table other than the one on which the trigger was defined failed with a `Table '...' doesn't exist` referring to the second table when attempting to delete records from the first table. (Bug #21825)

- `CURDATE()` is less than `NOW()`, either when comparing `CURDATE()` directly (`CURDATE() < NOW()` is true) or when casting `CURDATE()` to `DATE` (`CAST(CURDATE() AS DATE) < NOW()` is true). However, storing `CURDATE()` in a `DATE` column and comparing `col_name < NOW()` incorrectly yielded false. This is fixed by comparing a `DATE` column as `DATETIME` for comparisons to a `DATETIME` constant. (Bug #21103)
- When a stored routine attempted to execute a statement accessing a nonexistent table, the error was not caught by the routine's exception handler. (Bug #20713, Bug #8407)
- For a stored procedure containing a `SELECT` statement that used a complicated join with an `ON` expression, the expression could be ignored during re-execution of the procedure, yielding an incorrect result. (Bug #20492)
- The conditions checked by the optimizer to permit use of indexes in `IN` predicate calculations were unnecessarily tight and were relaxed. (Bug #20420)
- When a `TIME_FORMAT()` expression was used as a column in a `GROUP BY` clause, the expression result was truncated. (Bug #20293)
- The creation of MySQL system tables was not checked for by `mysql-test-run.pl`. (Bug #20166)
- For index reads, the `BLACKHOLE` engine did not return end-of-file (which it must because `BLACKHOLE` tables contain no rows), causing some queries to crash. (Bug #19717)
- For `expr IN(value_list)`, the result could be incorrect if `BIGINT UNSIGNED` values were used for `expr` or in the value list. (Bug #19342)
- When attempting to call a stored procedure creating a table from a trigger on a table `tbl` in a database `db`, the trigger failed with `ERROR 1146 (42S02): Table 'db.tbl' doesn't exist`. However, the actual reason that such a trigger fails is due to the fact that `CREATE TABLE` causes an implicit `COMMIT`, and so a trigger cannot invoke a stored routine containing this statement. A trigger which does so now fails with `ERROR 1422 (HY000): Explicit or implicit commit is not permitted in stored function or trigger`, which makes clear the reason for the trigger's failure. (Bug #18914)
- While preparing prepared statements, the server acquired unnecessary table write locks. (Bug #18326)
- The update columns for `INSERT ... SELECT ... ON DUPLICATE KEY UPDATE` could be assigned incorrect values if a temporary table was used to evaluate the `SELECT`. (Bug #16630)
- For `SUBSTRING()` evaluation using a temporary table, when `SUBSTRING()` was used on a `LONGTEXT` column, the `max_length` metadata value of the result was incorrectly calculated and set to 0. Consequently, an empty string was returned instead of the correct result. (Bug #15757)
- Local variables in stored routines or triggers, when declared as the `BIT` type, were interpreted as strings. (Bug #12976)
- For some operations, system tables in the `mysql` database must be accessed. For example, the `HELP` statement requires the contents of the server-side help tables, and `CONVERT_TZ()` might need to read the time zone tables. However, to perform such operations while a `LOCK TABLES` statement is in effect, the server required you to also lock the requisite system tables explicitly or a lock error occurred:

```
mysql> LOCK TABLE t1 READ;
Query OK, 0 rows affected (0.02 sec)
mysql> HELP HELP;
ERROR 1100 (HY000) at line 4: Table 'help_topic' was not
locked with LOCK TABLES
```

Now, the server implicitly locks the system tables for reading as necessary so that you need not lock them explicitly. These tables are treated as just described:

```
mysql.help_category
mysql.help_keyword
mysql.help_relation
mysql.help_topic
mysql.proc
mysql.time_zone
mysql.time_zone_leap_second
mysql.time_zone_name
mysql.time_zone_transition
mysql.time_zone_transition_type
```

If you want to explicitly place a `WRITE` lock on any of those tables with a `LOCK TABLES` statement, the table must be the only one locked; no other table can be locked with the same statement. (Bug #9953)

Changes in MySQL 5.1.16 (2007-02-26)

This is a new Beta development release, fixing recently discovered bugs.



Note

This Beta 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.



Note

After release, a trigger failure problem was found to have been introduced. (Bug #27006) Users affected by this issue should upgrade to MySQL 5.1.17, which corrects the problem.

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 Enterprise* (a commercial MySQL offering). For more details, please see (<http://www.mysql.com/products/enterprise>).

- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

Functionality Added or Changed

- **Incompatible Change; Cluster API:** The `AbortOption` type is now a member of the `NdbOperation` class; its values and behavior have also changed. `NdbTransaction::AbortOption` can no longer be used, and applications written against the NDB API may need to be rewritten and recompiled to accommodate these changes. For more information about this change, see [The NdbOperation::AbortOption Type](#).

This also affects the behavior of the `NdbTransaction::execute()` method, which now reports failure only if the transaction was actually aborted.

- **MySQL Cluster:** Previously, when a data node failed to start more than 8 times in succession, this caused a forced shutdown of the cluster. Now, when a data node fails to start 7 consecutive times, the node does not start again until it is started with the `--initial` option, and a warning to this effect is written to the error log. (Bug #25984)
- **MySQL Cluster:** In the event that all cluster management and API nodes are configured with `ArbitrationRank = 0`, `ndb_mgmd` now issues the following warning when starting: `Cluster configuration warning: Neither MGM nor API nodes are configured with arbitrator, may cause complete cluster shutdown in case of host failure.` (Bug #23546)
- **MySQL Cluster:** A number of new and more descriptive error messages covering transporter errors were added. (Bug #22025)
- **MySQL Cluster:** A new configuration parameter `MemReportFrequency` enables additional control of data node memory usage. Previously, only warnings at predetermined percentages of memory allocation were given; setting this parameter enables that behavior to be overridden.
- **Cluster API:** A new `ndb_mgm_get_clusterlog_loglevel()` function was added to the MGM API.
- The `localhost` anonymous user account created during MySQL installation on Windows now has no global privileges. Formerly this account had all global privileges. For operations that require global privileges, the `root` account can be used instead. (Bug #24496)
- In the `INFORMATION_SCHEMA REFERENTIAL_CONSTRAINTS` table, the `UNIQUE_CONSTRAINT_NAME` column incorrectly named the referenced table. Now it names the referenced constraint, and a new column, `REFERENCED_TABLE_NAME`, names the referenced table. (Bug #21713)
- `RAND()` now permits nonconstant initializers (such as a column name) as its argument. In this case, the seed is initialized with the value for each invocation of `RAND()`. (One implication of this is that for equal argument values, `RAND()` will return the same value each time.) (Bug #6172)
- Added the `--auto-generate-sql-load-type` and `--auto-generate-sql-write-number` options for `mysqslap`.
- The bundled yaSSL library was upgraded to version 1.5.8.

Bugs Fixed

- **Security Fix:** Using an `INFORMATION_SCHEMA` table with `ORDER BY` in a subquery could cause a server crash.

We would like to thank Oren Isacson of Flowgate Security Consulting and Stefan Streichsbier of SEC Consult for informing us of this problem. (Bug #24630, Bug #26556, CVE-2007-1420)

- **MySQL Cluster; Partitioning:** A query with an `IN` clause against an `NDB` table employing explicit user-defined partitioning did not always return all matching rows. (Bug #25821)
- **MySQL Cluster; Replication:** (Replication): Under some circumstances, the binary log thread could shut down while the slave SQL thread was still using it. (Bug #26015, Bug #26019)
- **MySQL Cluster; Replication:** (Replication): The error message `Last_Errno: 4294967295, Error in Write_rows event` now supplies a valid error code. (Bug #19896)
- **MySQL Cluster:** An inadvertent use of unaligned data caused `ndb_restore` to fail on some 64-bit platforms, including Sparc and Itanium-2. (Bug #26739)

- **MySQL Cluster:** The `InvalidUndoBufferSize` error used the same error code (763) as the `IncompatibleVersions` error. `InvalidUndoBufferSize` now uses its own error code (779). (Bug #26490)
- **MySQL Cluster:** The failure of a data node when restarting it with `--initial` could lead to failures of subsequent data node restarts. (Bug #26481)
- **MySQL Cluster:** Takeover for local checkpointing due to multiple failures of master nodes was sometimes incorrectly handled. (Bug #26457)
- **MySQL Cluster:** The `LockPagesInMainMemory` parameter was not read until after distributed communication had already started between cluster nodes. When the value of this parameter was 1, this could sometimes result in data node failure due to missed heartbeats. (Bug #26454)
- **MySQL Cluster:** Under some circumstances, following the restart of a management node, all data nodes would connect to it normally, but some of them subsequently failed to log any events to the management node. (Bug #26293)
- **MySQL Cluster:** Condition pushdown did not work with prepared statements. (Bug #26225)
- **MySQL Cluster:** A memory leak could cause problems during a node or cluster shutdown or failure. (Bug #25997)
- **MySQL Cluster:** No appropriate error message was provided when there was insufficient REDO log file space for the cluster to start. (Bug #25801)
- **MySQL Cluster:** An `UPDATE` using an `IN` clause on an `NDB` table on which there was a trigger caused `mysqld` to crash. (Bug #25522)
- **MySQL Cluster:** A memory allocation failure in `SUMA` (the cluster Subscription Manager) could cause the cluster to crash. (Bug #25239)
- **MySQL Cluster:** The `ndb_size.tmpl` file (necessary for using the `ndb_size.pl` script) was missing from binary distributions. (Bug #24191)
- **MySQL Cluster:** The message `Error 0 in readAutoIncrementValue(): no Error` was written to the error log whenever `SHOW TABLE STATUS` was performed on a Cluster table that did not have an `AUTO_INCREMENT` column.

**Note**

This improves on and supersedes an earlier fix that was made for this issue in MySQL 5.1.12.

(Bug #21033)

- **MySQL Cluster:** When a node failed due to there being insufficient disk space to perform a local checkpoint, there was no indication that this was the source of the problem. Such a condition now produces an appropriate error message. (Bug #20121)
- **MySQL Cluster:** In the event that cluster backup parameters such as `BackupWriteSize` were incorrectly set, no appropriate error was issued to indicate that this was the case. (Bug #19146)
- **Replication:** If a slave server closed its relay log (for example, due to an error during log rotation), the I/O thread did not recognize this and still tried to write to the log, causing a server crash. (Bug #10798)
- **Disk Data; Cluster API:** A delete and a read performed in the same operation could cause one or more data nodes to crash. This could occur when the operation affected more than 5 columns concurrently, or when one or more of the columns was of the `VARCHAR` type and was stored on disk. (Bug #25794)

- **Cluster API:** After defining a delete operation (using `NdbOperation::deleteTuple()`) on a nonexistent primary key of a table having a `BLOB` or `TEXT` column, invoking `NdbTransaction::execute()` caused the calling application to enter an endless loop rather than raising an error.

This issue also affected `ndb_restore`; when restoring tables containing `BLOB` or `TEXT` columns, this could cause it to consume all available memory and then crash. (Bug #24028)

References: See also: Bug #27308, Bug #30177.

- **Cluster API:** `libndbclient.so` was not versioned. (Bug #13522)
- Using `ORDER BY` or `GROUP BY` could yield different results when selecting from a view and selecting from the underlying table. (Bug #26209)
- `DISTINCT` queries that were executed using a loose scan for an `InnoDB` table that had been emptied caused a server crash. (Bug #26159)
- A `WHERE` clause that used `BETWEEN` for `DATETIME` values could be treated differently for a `SELECT` and a view defined as that `SELECT`. (Bug #26124)
- Collation for `LEFT JOIN` comparisons could be evaluated incorrectly, leading to improper query results. (Bug #26017)
- The `WITH CHECK OPTION` clause for views was ignored for updates of multiple-table views when the updates could not be performed on fly and the rows to update had to be put into temporary tables first. (Bug #25931)
- `LOAD DATA INFILE` did not work with pipes. (Bug #25807)
- The `SEC_TO_TIME()` and `QUARTER()` functions sometimes did not handle `NULL` values correctly. (Bug #25643)
- View definitions that used the `!` operator were treated as containing the `NOT` operator, which has a different precedence and can produce different results. . (Bug #25580)
- An error in the name resolution of nested `JOIN ... USING` constructs was corrected. (Bug #25575)
- `GROUP BY` and `DISTINCT` did not group `NULL` values for columns that have a `UNIQUE` index. . (Bug #25551)
- The `--with-readline` option for `configure` did not work for commercial source packages, but no error message was printed to that effect. Now a message is printed. (Bug #25530)
- A yaSSL program named `test` was installed, causing conflicts with the `test` system utility. It is no longer installed. (Bug #25417)
- For a `UNIQUE` index containing many `NULL` values, the optimizer would prefer the index for `col IS NULL` conditions over other more selective indexes. . (Bug #25407)
- An `AFTER UPDATE` trigger on an `InnoDB` table with a composite primary key caused the server to crash. (Bug #25398)
- Passing a `NULL` value to a user-defined function from within a stored procedure crashes the server. (Bug #25382)
- `perror` crashed on some platforms due to failure to handle a `NULL` pointer. (Bug #25344)
- `mysql.server stop` timed out too quickly (35 seconds) waiting for the server to exit. Now it waits up to 15 minutes, to ensure that the server exits. (Bug #25341)

- A query that contained an `EXISTS` subquery with a `UNION` over correlated and uncorrelated `SELECT` queries could cause the server to crash. (Bug #25219)
- `mysql_kill()` caused a server crash when used on an SSL connection. (Bug #25203)
- `yaSSL` was sensitive to the presence of whitespace at the ends of lines in PEM-encoded certificates, causing a server crash. (Bug #25189)
- A query with `ORDER BY` and `GROUP BY` clauses where the `ORDER BY` clause had more elements than the `GROUP BY` clause caused a memory overrun leading to a crash of the server. (Bug #25172)
- Use of `ON DUPLICATE KEY UPDATE` defeated the usual restriction against inserting into a join-based view unless only one of the underlying tables is used. (Bug #25123)
- `ALTER TABLE ... ENABLE KEYS` acquired a global lock, preventing concurrent execution of other statements that use tables. (Bug #25044)
- `OPTIMIZE TABLE` caused a race condition in the I/O cache. (Bug #25042)
- A return value of `-1` from user-defined handlers was not handled well and could result in conflicts with server code. (Bug #24987)
- Certain joins using `Range checked for each record` in the query execution plan could cause the server to crash. (Bug #24776)
- `ALTER TABLE` caused loss of `CASCADE` clauses for `InnoDB` tables. (Bug #24741)
- If an `ORDER BY` or `GROUP BY` list included a constant expression being optimized away and, at the same time, containing single-row subselects that returned more than one row, no error was reported. If a query required sorting by expressions containing single-row subselects that returned more than one row, execution of the query could cause a server crash. (Bug #24653)
- For `ALTER TABLE`, using `ORDER BY expression` could cause a server crash. Now the `ORDER BY` clause permits only column names to be specified as sort criteria (which was the only documented syntax, anyway). (Bug #24562)
- Within stored routines or prepared statements, inconsistent results occurred with multiple use of `INSERT ... SELECT ... ON DUPLICATE KEY UPDATE` when the `ON DUPLICATE KEY UPDATE` clause erroneously tried to assign a value to a column mentioned only in its `SELECT` part. (Bug #24491)
- Expressions of the form `(a, b) IN (SELECT a, MIN(b) FROM t GROUP BY a)` could produce incorrect results when column `a` of table `t` contained `NULL` values while column `b` did not. (Bug #24420)
- If a prepared statement accessed a view, access to the tables listed in the query after that view was checked in the security context of the view. (Bug #24404)
- A nested query on a partitioned table returned fewer records than on the corresponding nonpartitioned table, when the subquery affected more than one partition. (Bug #24186)
- Expressions of the form `(a, b) IN (SELECT c, d ...)` could produce incorrect results if `a`, `b`, or both were `NULL`. (Bug #24127)
- Queries that evaluate `NULL IN (SELECT ... UNION SELECT ...)` could produce an incorrect result (`FALSE` instead of `NULL`). (Bug #24085)
- Some `UPDATE` statements were slower than in previous versions when the search key could not be converted to a valid value for the type of the search column. (Bug #24035)
- `ISNULL(DATE(NULL))` and `ISNULL(CAST(NULL AS DATE))` erroneously returned false. (Bug #23938)

- Within a stored routine, accessing a declared routine variable with `PROCEDURE ANALYSE()` caused a server crash. (Bug #23782)
- For an `InnoDB` table with any `ON DELETE` trigger, `TRUNCATE TABLE` mapped to `DELETE` and activated triggers. Now a fast truncation occurs and triggers are not activated. .

**Important**

As a result of this fix, `TRUNCATE TABLE` now requires the `DROP` privilege rather than the `DELETE` privilege.

(Bug #23556)

- With `ONLY_FULL_GROUP_BY` enabled, the server was too strict: Some expressions involving only aggregate values were rejected as nonaggregate (for example, `MAX(a) - MIN(a)`). (Bug #23417)
- The arguments to the `ENCODE()` and the `DECODE()` functions were not printed correctly, causing problems in the output of `EXPLAIN EXTENDED` and in view definitions. (Bug #23409)
- Some queries against `INFORMATION_SCHEMA` that used subqueries failed. . (Bug #23299)
- `readline` detection did not work correctly on NetBSD. (Bug #23293)
- The number of `setsockopt()` calls performed for reads and writes to the network socket was reduced to decrease system call overhead. (Bug #22943)
- Storing values specified as hexadecimal values 64 or more bits long in `BIT(64)`, `BIGINT`, or `BIGINT UNSIGNED` columns did not raise any warning or error if the value was out of range. (Bug #22533)
- Type conversion errors during formation of index search conditions were not correctly checked, leading to incorrect query results. (Bug #22344)
- For the `IF()` and `COALESCE()` function and `CASE` expressions, large unsigned integer values could be mishandled and result in warnings. (Bug #22026)
- Inserting `DEFAULT` into a column with no default value could result in garbage in the column. Now the same result occurs as when inserting `NULL` into a `NOT NULL` column. (Bug #20691)
- Indexes disabled with `ALTER TABLE ... DISABLE KEYS` could in some cases be used by specifying `FORCE INDEX`. (Bug #20604)
- If a duplicate key value was present in the table, `INSERT ... ON DUPLICATE KEY UPDATE` reported a row count indicating that a record was updated, even when no record actually changed due to the old and new values being the same. Now it reports a row count of zero. (Bug #19978)

References: See also: Bug #27006, Bug #27033, Bug #27210.

- `ORDER BY` values of the `DOUBLE` or `DECIMAL` types could change the result returned by a query. (Bug #19690)
- The `readline` library wrote to uninitialized memory, causing `mysql` to crash. (Bug #19474)
- Use of already freed memory caused SSL connections to hang forever. (Bug #19209)
- The server might fail to use an appropriate index for `DELETE` when `ORDER BY`, `LIMIT`, and a nonrestricting `WHERE` are present. (Bug #17711)
- The optimizer used a filesort rather than a `const` table read in some cases when the latter was possible. (Bug #16590)

- To enable installation of MySQL RPMs on Linux systems running RHEL 4 (which includes SE-Linux) additional information was provided to specify some actions that are permitted to the MySQL binaries. (Bug #12676)
- `CONNECTION` is no longer treated as a reserved word. (Bug #12204)
- The presence of `ORDER BY` in a view definition prevented the `MERGE` algorithm from being used to resolve the view even if nothing else in the definition required the `TEMPTABLE` algorithm. (Bug #12122)

Changes in MySQL 5.1.15 (2007-01-25)

This is a new Beta development release, fixing recently discovered bugs.



Note

This Beta 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.

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 Enterprise* (a commercial MySQL offering). For more details, please see (<http://www.mysql.com/products/enterprise>).

- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

Functionality Added or Changed

- **Incompatible Change; MySQL Cluster:** The `LockPagesInMainMemory` configuration parameter has changed its type and possible values.



Important

The values `true` and `false` are no longer accepted for this parameter. If you were using this parameter and had it set to `false` in a previous release, you must change it to `0`. If you had this parameter set to `true`, you should instead use `1` to obtain the same behavior as previously, or `2` to take advantage of new functionality introduced with this release, as described in the section cited above.

(Bug #25686)

- **Incompatible Change:** `InnoDB` rolls back only the last statement on a transaction timeout. A new option, `--innodb_rollback_on_timeout`, causes `InnoDB` to abort and roll back the entire transaction if a transaction timeout occurs (the same behavior as in MySQL 5.0.13 and earlier). (Bug #24200)
- **Incompatible Change:** Previously, the `DATE_FORMAT()` function returned a binary string. Now it returns a string with a character set and collation given by `character_set_connection` and `collation_connection` so that it can return month and weekday names containing non-ASCII characters. (Bug #22646)
- **Incompatible Change:** The following conditions apply to enabling the `read_only` system variable:

- If you attempt to enable `read_only` while you have any explicit locks (acquired with `LOCK TABLES` or have a pending transaction, an error will occur.
- If other clients hold explicit table locks or have pending transactions, the attempt to enable `read_only` blocks until the locks are released and the transactions end. While the attempt to enable `read_only` is pending, requests by other clients for table locks or to begin transactions also block until `read_only` has been set.
- `read_only` can be enabled while you hold a global read lock (acquired with `FLUSH TABLES WITH READ LOCK`) because that does not involve table locks.

Previously, the attempt to enable `read_only` would return immediately even if explicit locks or transactions were pending, so some data changes could occur for statements executing in the server at the same time. (Bug #22009, Bug #11733)

- **Incompatible Change:** Previously, the `ARCHIVE` storage engine created a metadata file with an extension of `.ARM` for each table. The engine no longer creates this file.
- **Important Change:** When using a `MERGE` table, the definition of the table and the underlying `MyISAM` tables are checked each time the tables are opened for access (including any `SELECT` or `INSERT` statement). Each table is compared for column order, types, sizes, and associated indexes. If there is a difference in any one of the tables, the statement will fail.
- **Important Change:** Previously, duplicate-key errors were indicated by the `ER_DUP_ENTRY` error code (1062). This code is no longer used. Instead, the server returns `ER_DUP_ENTRY_WITH_KEY_NAME` (1582), and the error message indicates the name of the index for which the duplicate occurred. Applications that test for duplicate keys should look for both error codes if they need to be compatible with current and older servers.

References: See also: Bug #28842.

- **MySQL Cluster:** The `NDB` storage engine could leak memory during file operations. (Bug #21858)
- **Replication:** Calling a nondeterministic stored routine when using statement-based replication now throws an error. Formerly, defining such a stored routine would cause an error to be thrown. (Bug #16456)
- On Unix, when searching the standard locations for option files, MySQL programs now also look for `/etc/mysql/my.cnf` after checking for `/etc/my.cnf` and before checking the remaining locations. (Bug #25104)
- The default value of the `max_connections` variable has been increased to 151 so that Web sites running on Apache and using MySQL will not have more processes trying to access MySQL than the default number of connections available.

The maximum number of Apache processes is determined by the Apache `MaxClient` setting, which defaults to 256, but is usually set to 150 in the `httpd.conf` commonly distributed with Apache. For more information about `MaxClient`, see http://httpd.apache.org/docs/2.2/mod/mpm_common.html#maxclients. (Bug #23883)

- The `Com_create_user` status variable was added (for counting `CREATE USER` statements). (Bug #22958)
- The `--memlock` option relies on system calls that are unreliable on some operating systems. If a crash occurs, the server now checks whether `--memlock` was specified and if so issues some information about possible workarounds. (Bug #22860)

- The (undocumented) `UNIQUE_USERS()` and `GROUP_UNIQUE_USERS()` functions were removed. (Bug #22687)
- Partitioning of tables using the `FEDERATED` storage engine is no longer permitted. Attempting to create such a table or to modify an existing table so that it uses both partitioning and `FEDERATED` now fails with an error. (Bug #22451)
- The `--skip-thread-priority` option now is enabled by default for binary Mac OS X distributions. Use of thread priorities degrades performance on Mac OS X. (Bug #18526)
- The bundled yaSSL library was upgraded to version 1.5.0.
- Remote servers for use with the `FEDERATED` storage engine now can be managed with the new `CREATE/ALTER/DROP SERVER` syntax.
- Added the `--disable-grant-options` option to `configure`. If `configure` is run with this option, the `--bootstrap`, `--skip-grant-tables`, and `--init-file` options for `mysqld` are disabled and cannot be used. For Windows, the `configure.js` script recognizes the `DISABLE_GRANT_OPTIONS` flag, which has the same effect.

Bugs Fixed

- **Incompatible Change:** For `ENUM` columns that had enumeration values containing commas, the commas were mapped to `0xff` internally. However, this rendered the commas indistinguishable from true `0xff` characters in the values. This no longer occurs. However, the fix requires that you dump and reload any tables that have `ENUM` columns containing any true `0xff` values. Dump the tables using `mysqldump` with the current server before upgrading from a version of MySQL 5.1 older than 5.1.15 to version 5.1.15 or newer. (Bug #24660)
- **MySQL Cluster; Partitioning:** Non-32-bit, nonaligned columns were not handled correctly in explicitly partitioned `NDB` tables. (Bug #25587)
- **MySQL Cluster; Replication:** (Replication): Connecting a `mysqld` to a cluster where not all nodes were running, starting the remaining cluster nodes, and then disconnecting from the cluster caused the `mysqld` process to crash. (Bug #25387)
- **MySQL Cluster:** It was not possible to create an `NDB` table with a key on two `VARCHAR` columns where both columns had a storage length in excess of 256. (Bug #25746)
- **MySQL Cluster:** Hosts in clusters with large numbers of nodes could experience excessive CPU usage while obtaining configuration data. (Bug #25711)
- **MySQL Cluster:** In some circumstances, shutting down the cluster could cause connected `mysqld` processes to crash. (Bug #25668)
- **MySQL Cluster:** Some aggregate queries such as `SELECT COUNT(*)` performed a table scan on `NDB` tables rather than checking table statistics, causing such queries to perform much more slowly in MySQL Cluster 5.1 than in 5.0. (Bug #25567)
- **MySQL Cluster:** Memory allocations for `TEXT` columns were calculated incorrectly, resulting in space being wasted and other issues. (Bug #25562)
- **MySQL Cluster:** The failure of a master node during a node restart could lead to a resource leak, causing later node failures. (Bug #25554)
- **MySQL Cluster:** The failure of a node during a local checkpoint could lead to other node failures. (Bug #25468)
- **MySQL Cluster:** A node shutdown occurred if the master failed during a commit. (Bug #25364)

- **MySQL Cluster:** Creating a nonunique index with the `USING HASH` clause silently created an ordered index instead of issuing a warning. (Bug #24820)
- **MySQL Cluster:** `ndb_config` failed when trying to use 2 management servers and node IDs. (Bug #23887)
- **MySQL Cluster:** When a data node was shut down using the management client `STOP` command, a connection event (`NDB_LE_Connected`) was logged instead of a disconnection event (`NDB_LE_Disconnected`). (Bug #22773)
- **MySQL Cluster:** The management server did not handle logging of node shutdown events correctly in certain cases. (Bug #22013)
- **MySQL Cluster:** `SELECT` statements with a `BLOB` or `TEXT` column in the selected column list and a `WHERE` condition including a primary key lookup on a `VARCHAR` primary key produced empty result sets. (Bug #19956)
- **MySQL Cluster:** When stopping and restarting multiple data nodes, the last node to be restarted would sometimes hang in Phase 100. (Bug #19645)
- **Replication:** Using row-based replication to replicate to a table having at least one extra `BIT` column with a default value on the slave as compared to the master could cause the slave to fail. (Bug #24490)
- **Replication:** When `SET PASSWORD` was written to the binary log, double quotation marks were included in the statement. If the slave was running in with the server SQL mode set to `ANSI_QUOTES`, then the event failed, which halted the replication process. (Bug #24158)
- **Replication:** A stored procedure, executed from a connection using a binary character set, and which wrote multibyte data, would write incorrectly escaped entries to the binary log. This caused syntax errors, and caused replication to fail. (Bug #23619, Bug #24492)
- **Replication:** Using `CREATE TABLE ... SELECT` and rolling back the transaction would leave an empty table on the master, but the instructions would not be recorded in the binary log and therefore replicated to the slave. This would result in a difference between the master and slave databases. An implicit commit has been added to ensure consistency. (Bug #22865)
- **Replication:** Changes to the `lc_time_names` system variable were not replicated. (Bug #22645)
- **Replication:** For `SET`, `SELECT`, and `DO` statements that invoked a stored function from a database other than the default database, the function invocation could fail to be replicated. (Bug #19725)
- **Disk Data:** Following three or more missed local checkpoints by a cluster node, a restart of the node caused incorrect undo information to be used for Disk Data tables. (Bug #25636)
- **Disk Data:** `MEDIUMTEXT` columns of Disk Data tables were stored in memory rather than on disk, even if the columns were not indexed. (Bug #25001)
- **Disk Data:** Performing a node restart with a newly dropped Disk Data table could lead to failure of the node during the restart. (Bug #24917)
- **Disk Data:** Changing a column specification or issuing a `TRUNCATE TABLE` statement on a Disk Data table caused the table to become an in-memory table. (Bug #24667, Bug #25296)
- **Disk Data:** When restoring from backup a cluster containing any Disk Data tables with hidden primary keys, a node failure resulted which could lead to a crash of the cluster. (Bug #24166)
- **Disk Data:** Repeated `CREATE`, `DROP`, or `TRUNCATE TABLE` in various combinations with system restarts between these operations could lead to the eventual failure of a system restart. (Bug #21948)

- **Disk Data:** Extents that should have been available for re-use following a `DROP TABLE` operation were not actually made available again until after the cluster had performed a local checkpoint. (Bug #17605)
- **Cluster Replication:** Certain errors in replication setups could lead to unexpected node failures. (Bug #25755)
- **Cluster Replication:** Connecting an API node to the cluster during a node restart while performing database operations could cause the restarting node to fail. (Bug #25329)
- **Cluster Replication:** Following a restart of the master cluster, the latest GCI was set to 0 upon reconnection to the slave. (Bug #21806)
- **Cluster API:** Deletion of an `Ndb_cluster_connection` object took a very long time. (Bug #25487)
- **Cluster API:** Invoking the `NdbTransaction::execute()` method using execution type `Commit` and abort option `AO_IgnoreError` could lead to a crash of the transaction coordinator (DBTC). (Bug #25090)
- **Cluster API:** A unique index lookup on a nonexistent tuple could lead to a data node timeout (error 4012). (Bug #25059)
- **Cluster API:** When using the `NdbTransaction::execute()` method, a very long timeout (greater than 5 minutes) could result if the last data node being polled was disconnected from the cluster. (Bug #24949)
- **Cluster API:** Due to an error in the computation of table fragment arrays, some transactions were not executed from the correct starting point. (Bug #24914)
- `mysqltest_embedded` crashed at startup. (Bug #25890)
- Referencing an ambiguous column alias in an expression in the `ORDER BY` clause of a query caused the server to crash. (Bug #25427)
- A number of issues were uncovered by Valgrind. (Bug #25396)
- Using a view in combination with a `USING` clause caused column aliases to be ignored. (Bug #25106)
- A multiple-table `DELETE QUICK` could sometimes cause one of the affected tables to become corrupted. (Bug #25048)
- An assertion failed incorrectly for prepared statements that contained a single-row uncorrelated subquery that was used as an argument of the `IS NULL` predicate. (Bug #25027)
- In the `INFORMATION_SCHEMA.KEY_COLUMN_USAGE` table, the value displayed for the `REFERENCED_TABLE_NAME` column was the table name as encoded for disk storage, not the actual table name. (Bug #25026)
- The `REPEAT()` function could return `NULL` when passed a column for the count argument. (Bug #24947)
- `mysql_upgrade` failed if the `--password` (or `-p`) option was given. (Bug #24896)
- Accessing a fixed record format table with a crashed key definition results in server/`myisamchk` segmentation fault. (Bug #24855)
- `mysqld_multi` and `mysqlaccess` looked for option files in `/etc` even if the `--sysconfdir` option for `configure` had been given to specify a different directory. (Bug #24780)
- If there was insufficient memory available to `mysqld`, this could sometimes cause the server to hang during startup. (Bug #24751)

- Optimizations that are legal only for subqueries without tables and `WHERE` conditions were applied for any subquery without tables. (Bug #24670)
- Under certain rare circumstances, local checkpoints were not performed properly, leading to an inability to restart one or more data nodes. (Bug #24664)
- A workaround was implemented to avoid a race condition in the NPTL `pthread_exit()` implementation. (Bug #24507)
- Under some circumstances, a `REORGANIZE PARTITION` statement could crash `mysqld`. (Bug #24502)
- `mysqltest` crashed with a stack overflow. (Bug #24498)
- Attempts to access a `MyISAM` table with a corrupt column definition caused a server crash. (Bug #24401)
- `ALTER TABLE ENABLE KEYS` or `ALTER TABLE DISABLE KEYS` combined with another `ALTER TABLE` option other than `RENAME TO` did nothing. In addition, if `ALTER TABLE` was used on a table having disabled keys, the keys of the resulting table were enabled. (Bug #24395)
- When opening a corrupted `.frm` file during a query, the server crashes. (Bug #24358)
- The `--extern` option for `mysql-test-run.pl` did not function correctly. (Bug #24354)
- Some joins in which one of the joined tables was a view could return erroneous results or crash the server. (Bug #24345)
- The `mysql.server` script used the `source` command, which is less portable than the `.` command; it now uses `.` instead. (Bug #24294)
- A view was not handled correctly if the `SELECT` part contained “\z”. (Bug #24293)
- `mysql_install_db` did not create the `mysql.plugin` table if strict SQL mode was enabled. (Bug #24270)
- A query using `WHERE unsigned_column NOT IN ('negative_value')` could cause the server to crash. (Bug #24261)
- `ALTER TABLE` statements that performed both `RENAME TO` and `{ENABLE|DISABLE} KEYS` operations caused a server crash. (Bug #24219)
- A `FETCH` statement using a cursor on a table which was not in the table cache could sometimes cause the server to crash. (Bug #24117)
- Hebrew-to-Unicode conversion failed for some characters. Definitions for the following Hebrew characters (as specified by the ISO/IEC 8859-8:1999) were added: LEFT-TO-RIGHT MARK (LRM), RIGHT-TO-LEFT MARK (RLM) (Bug #24037)
- On HP-UX, `mysqltest` (nonthread-safe) crashed due to being linked against a thread-safe `libmysys` library. (Bug #23984)
- The server was built even when `configure` was run with the `--without-server` option. (Bug #23973)

References: See also: Bug #32898.

- The MySQL 5.1.12 binaries for Windows were missing the `FEDERATED`, `EXAMPLE`, and `BLACKHOLE` storage engines. (Bug #23900)
- `ROW_COUNT()` did not work properly as an argument to a stored procedure. (Bug #23760)

- When reading from the standard input on Windows, `mysqlbinlog` opened the input in text mode rather than binary mode and consequently misinterpreted some characters such as **Control+Z**. (Bug #23735)
- `OPTIMIZE TABLE` tried to sort R-tree indexes such as spatial indexes, although this is not possible (see [OPTIMIZE TABLE Syntax](#)). (Bug #23578)
- The row count for `MyISAM` tables was not updated properly, causing `SHOW TABLE STATUS` to report incorrect values. (Bug #23526)
- The Instance Manager `DROP INSTANCE` command did not work. (Bug #23476)
- User-defined variables could consume excess memory, leading to a crash caused by the exhaustion of resources available to the `MEMORY` storage engine, due to the fact that this engine is used by MySQL for variable storage and intermediate results of `GROUP BY` queries. Where `SET` had been used, such a condition could instead give rise to the misleading error message `You may only use constant expressions with SET, rather than Out of memory (Needed NNNNNN bytes)`. (Bug #23443)
- A table created with the `ROW_FORMAT = FIXED` table option lost that option if an index was added or dropped with `CREATE INDEX` or `DROP INDEX`. (Bug #23404)
- A deadlock could occur, with the server hanging on `Closing tables`, with a sufficient number of concurrent `INSERT DELAYED`, `FLUSH TABLES`, and `ALTER TABLE` operations. (Bug #23312)
- Accuracy was improved for comparisons between `DECIMAL` columns and numbers represented as strings. (Bug #23260)
- The Instance Manager `STOP INSTANCE` command took too much time and caused Instance Manager to be unresponsive. (Bug #23215)
- If there was insufficient memory to store or update a blob record in a `MyISAM` table then the table will be marked as crashed. (Bug #23196)
- A compressed `MyISAM` table that became corrupted could crash `myisamchk` and possibly the MySQL Server. (Bug #23139)
- `CREATE TABLE ... SELECT` statements were not rolled back correctly. As part of the fix, such a statement now causes an implicit commit before and after it is executed. However, it does not cause a commit when used to create a temporary table. (Bug #22864)
- `mysql_upgrade` failed when called with a `--basedir` path name containing spaces. (Bug #22801)
- Using `INSTALL PLUGIN` followed by a restart of the server caused an error due to memory not being properly initialized. (Bug #22694)
- `SET lc_time_names = value` permitted only exact literal values, not expression values. (Bug #22647)
- A partitioned table that used the `DATA DIRECTORY` option, where the data directory was the same as the directory in which the table definition file resided, became corrupted following `ALTER TABLE ENGINE=ARCHIVE`. This was actually due to an issue with the `ARCHIVE` storage engine, and not with partitioned tables in general. (Bug #22634)
- The `STDDEV()` function returned a positive value for data sets consisting of a single value. (Bug #22555)
- `SHOW COLUMNS` reported some `NOT NULL` columns as `NULL`. (Bug #22377)
- A server crash occurred when using `LOAD DATA` to load a table containing a `NOT NULL` spatial column, when the statement did not load the spatial column. Now a `NULL supplied to NOT NULL column` error occurs. (Bug #22372)

- An `ALTER TABLE` statement that used a `RENAME` clause in combination with a `MODIFY` or `CHANGE` that did not actually change the table (for example, when it changed a column's type from `INT` to `INT`). The behavior caused by this bug differed according to whether or not the storage engine used by the table was transactional or nontransactional. For transactional tables (such as those using the `InnoDB` storage engine), the statement simply failed; for nontransactional tables (such as those using the `MyISAM` storage engine), the `ALTER TABLE` statement succeeding renaming the table, but subsequent `SELECT` statements against the renamed table failed. (Bug #22369)
- The Instance Manager `STOP INSTANCE` command could not be applied to instances in the `Crashed`, `Failed`, or `Abandoned` state. (Bug #22306)
- `DATE_ADD()` requires complete dates with no “zero” parts, but sometimes did not return `NULL` when given such a date. (Bug #22229)
- Some small double precision numbers (such as `1.00000001e-300`) that should have been accepted were truncated to zero. (Bug #22129)
- Changing the value of `MI_KEY_BLOCK_LENGTH` in `myisam.h` and recompiling MySQL resulted in a `myisamchk` that saw existing `MyISAM` tables as corrupt. (Bug #22119)
- For a nonexistent table, `DROP TEMPORARY TABLE` failed with an incorrect error message if `read_only` was enabled. (Bug #22077)
- A crash of the MySQL Server could occur when unpacking a `BLOB` column from a row in a corrupted `MyISAM` table. This could happen when trying to repair a table using either `REPAIR TABLE` or `myisamchk`; it could also happen when trying to access such a “broken” row using statements like `SELECT` if the table was not marked as crashed. (Bug #22053)
- The code for generating `USE` statements for binary logging of `CREATE PROCEDURE` statements resulted in confusing output from `mysqlbinlog` for `DROP PROCEDURE` statements. (Bug #22043)
- `STR_TO_DATE()` returned `NULL` if the format string contained a space following a nonformat character. (Bug #22029)
- It was possible to use `DATETIME` values whose year, month, and day parts were all zeros but whose hour, minute, and second parts contained nonzero values, an example of such an illegal `DATETIME` being `'0000-00-00 11:23:45'`.

**Note**

This fix was reverted in MySQL 5.1.18.

(Bug #21789)

References: See also: Bug #25301.

- SSL connections could hang at connection shutdown. (Bug #21781, Bug #24148)
- `yaSSL` crashed on pre-Pentium Intel CPUs. (Bug #21765)
- Using `FLUSH TABLES` in one connection while another connection is using `HANDLER` statements caused a server crash.

**Note**

This fix was reverted in MySQL 5.1.22

(Bug #21587)

References: See also: Bug #29474.

- The `FEDERATED` storage engine did not support the `euckr` character set. (Bug #21556)
- `InnoDB` crashed while performing XA recovery of prepared transactions. (Bug #21468)
- It was possible to set the backslash character (“\”) as the delimiter character using `DELIMITER`, but not actually possible to use it as the delimiter. (Bug #21412)
- Using `ALTER TABLE` to convert a `CSV` table containing `NULL` values to `MyISAM` resulted in warnings. (Bug #21328)
- When updating a table that used a `JOIN` of the table itself (for example, when building trees) and the table was modified on one side of the expression, the table would either be reported as crashed or the wrong rows in the table would be updated. (Bug #21310)
- `mysqld_error.h` was not installed when only the client libraries were built. (Bug #21265)
- `InnoDB`: During a restart of the MySQL Server that followed the creation of a temporary table using the `InnoDB` storage engine, MySQL failed to clean up in such a way that `InnoDB` still attempted to find the files associated with such tables. (Bug #20867)
- Selecting into variables sometimes returned incorrect wrong results. (Bug #20836)
- Queries of the form `SELECT ... WHERE string = ANY(...)` failed when the server used a single-byte character set and the client used a multibyte character set. (Bug #20835)

References: See also: Bug #34760.

- `mysql_fix_privilege_tables.sql` altered the `table_privs.table_priv` column to contain too few privileges, causing loss of the `CREATE VIEW` and `SHOW VIEW` privileges. (Bug #20589)
- A stored routine containing semicolon in its body could not be reloaded from a dump of a binary log. (Bug #20396)
- `SELECT ... FOR UPDATE`, `SELECT ... LOCK IN SHARE MODE`, `DELETE`, and `UPDATE` statements executed using a full table scan were not releasing locks on rows that did not satisfy the `WHERE` condition. (Bug #20390)
- The `BUILD/check-cpu` script did not recognize Celeron processors. (Bug #20061)
- Unsigned `BIGINT` values treated as signed values by the `MOD()` function. (Bug #19955)
- Compiling PHP 5.1 with the MySQL static libraries failed on some versions of Linux. (Bug #19817)
- The `DELIMITER` statement did not work correctly when used in an SQL file run using the `SOURCE` statement. (Bug #19799)
- `mysqltest` incorrectly tried to retrieve result sets for some queries where no result set was available. (Bug #19410)
- `VARBINARY` column values inserted on a MySQL 4.1 server had trailing zeros following upgrade to MySQL 5.0 or later. (Bug #19371)
- Some `CASE` statements inside stored routines could lead to excessive resource usage or a crash of the server. (Bug #19194, Bug #24854)

- Instance Manager could crash during shutdown. (Bug #19044)
- `myisampack` wrote to unallocated memory, causing a crash. (Bug #17951)
- `FLUSH LOGS` or `mysqladmin flush-logs` caused a server crash if the binary log was not open. (Bug #17733)
- `mysql_fix_privilege_tables` did not accept a password containing embedded space or apostrophe characters. (Bug #17700)
- No warning was issued for use of the `DATA DIRECTORY` or `INDEX DIRECTORY` table options on a platform that does not support them. (Bug #17498)
- The `FEDERATED` storage engine did not support the `utf8` character set. (Bug #17044)
- The optimizer removes expressions from `GROUP BY` and `DISTINCT` clauses if they happen to participate in `expression = constant` predicates of the `WHERE` clause, the idea being that, if the expression is equal to a constant, then it cannot take on multiple values. However, for predicates where the expression and the constant item are of different result types (for example, when a string column is compared to 0), this is not valid, and can lead to invalid results in such cases. The optimizer now performs an additional check of the result types of the expression and the constant; if their types differ, then the expression is not removed from the `GROUP BY` list. (Bug #15881)
- When a prepared statement failed during the prepare operation, the error code was not cleared when it was reused, even if the subsequent use was successful. (Bug #15518)
- Dropping a user-defined function sometimes did not remove the UDF entry from the `mysql.proc` table. (Bug #15439)
- Inserting a row into a table without specifying a value for a `BINARY(N) NOT NULL` column caused the column to be set to spaces, not zeros. (Bug #14171)
- On Windows, the `SLEEP()` function could sleep too long, especially after a change to the system clock. (Bug #14094, Bug #24686, Bug #17635)
- `mysqldump --order-by-primary` failed if the primary key name was an identifier that required quoting. (Bug #13926)
- Subqueries of the form `NULL IN (SELECT ...)` returned invalid results. (Bug #8804, Bug #23485)

Changes in MySQL 5.1.14 (2006-12-05)

This is a new Beta development release, fixing recently discovered bugs.



Note

This Beta 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.

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 Enterprise* (a commercial MySQL offering). For more details, please see (<http://www.mysql.com/products/enterprise>).

- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

Functionality Added or Changed

- **Incompatible Change; Cluster Replication:** Two major changes have taken place with regard to the MySQL Cluster system tables. These are:
 1. The `cluster` database is no longer used. The tables formerly found in the `cluster` database are now in the `mysql` database, and have been renamed as `ndb_binlog_index`, `ndb_apply_status`, and `ndb_schema`.
 2. The `mysql.ndb_apply_status` and `mysql.ndb_schema` tables (formerly `cluster.apply_status` and `cluster.schema` are now created by `ndb_restore`, in the event that they do not already exist on the slave cluster.



Note

When upgrading from versions of MySQL previous to 5.1.14 to 5.1.14 or later, `mysql_fix_privilege_tables` merely creates a new `mysql.ndb_binlog_index` table, but does not remove the existing `cluster` database (or, if upgrading from MySQL 5.1.7 or earlier, the existing `cluster_replication` database), nor any of the tables in it.

For more information, see [MySQL Cluster Replication Schema and Tables](#). (Bug #14612)

- **Incompatible Change; Cluster Replication:** The `cluster` database is no longer used. The tables formerly found in the `cluster` database are now in the `mysql` database, and have been renamed as `ndb_binlog_index`, `ndb_apply_status`, and `ndb_schema`.
- **Incompatible Change:** The `prepared_stmt_count` system variable has been converted to the `Prepared_stmt_count` global status variable (viewable with the `SHOW GLOBAL STATUS` statement). (Bug #23159)
- **Incompatible Change:** Previously, you could create a user-defined function (UDF) or stored function with the same name as a built-in function, but could not invoke the UDF. Now an error occurs if you try to create such a UDF. The server also now generates a warning if you create a stored function with the same name as a built-in function. It is not considered an error to create a stored function with the same name as a built-in function because you can invoke the function using `db_name.func_name()` syntax. However, the server now generates a warning in this case.

See [Function Name Parsing and Resolution](#), for the rules describing how the server interprets references to different kinds of functions. (Bug #22619, Bug #18239)

- **Important Change; Disk Data:** The output of `mysqldump` now includes by default all tablespace and logfile group definitions used by any tables or databases that are dumped.

This fix also introduces the `--no-tablespaces` option (short form: `-y`) for `mysqldump`, which has the effect of suppressing all `CREATE LOGFILE GROUP` and `CREATE TABLESPACE` statements in the output.



Note

The working of the `--all-tablespaces` or `-Y` option for `mysqldump` remains unaffected by this change.

(Bug #20839)

- **MySQL Cluster:** Backup messages are now printed to the Cluster log. (Bug #24544)
- **MySQL Cluster:** Setting the configuration parameter `LockPagesInMainMemory` had no effect. (Bug #24461)
- **MySQL Cluster:** The error message `Management server closed connection`, when recorded in the MySQL error log, now includes a timestamp indicating when the error took place. (Bug #21519)
- **MySQL Cluster:** It is now possible to create a unique hashed index on a column that is not defined as `NOT NULL`.

**Note**

This change applies only to tables using the `NDB` storage engine.

Unique indexes on columns in `NDB` tables do not store null values because they are mapped to primary keys in an internal index table (and primary keys cannot contain nulls).

Normally, an additional ordered index is created when one creates unique indexes on `NDB` table columns; this can be used to search for `NULL` values. However, if `USING HASH` is specified when such an index is created, no ordered index is created.

The reason for permitting unique hash indexes with null values is that, in some cases, the user wants to save space if a large number of records are pre-allocated but not fully initialized. This also assumes that the user will *not* try to search for null values. Since MySQL does not support indexes that are not permitted to be searched in some cases, the `NDB` storage engine uses a full table scan with pushed conditions for the referenced index columns to return the correct result.

A warning is returned if one creates a unique nullable hash index, since the query optimizer should be provided a hint not to use it with `NULL` values if this can be avoided. (Bug #21507)

- `DROP TRIGGER` now supports an `IF EXISTS` clause. (Bug #23703)
- Direct and indirect usage of stored routines, user-defined functions, and table references is now prohibited in `CREATE EVENT` and `ALTER EVENT` statements.
See [CREATE EVENT Syntax](#), and [ALTER EVENT Syntax](#), for more specific information. (Bug #22830)
- The XPath operators `<` and `>`, as implemented in the `ExtractValue()` function, operated in reverse.
With this fix, all standard XPath comparison operators should now be supported correctly for use with the `ExtractValue()` and `UpdateXML()` functions. (Bug #22823)
- For the `mysql` client, display of result set metadata now is enabled with the `--column-type-info` option rather than with `--debug-info/-T`.
- `mysqladmin`, `mysqlcheck`, `mysqldump`, `mysqlimport`, and `mysqlshow` now accept the `--debug-info` option, which displays debugging information and memory and CPU usage statistics at program exit.

Bugs Fixed

- **Performance:** Evaluation of subqueries that require the filesort algorithm were allocating and freeing the `sort_buffer_size` buffer many times, resulting in slow performance. Now the buffer is allocated once and reused. (Bug #21727)
- **MySQL Cluster; Replication:** (Replication): If errors occurred during purging of the binary logs, extraneous rows could remain left in the `binlog_index` table. (Bug #15021)

- **MySQL Cluster:** The failure of a data node failure during a schema operation could lead to additional node failures. (Bug #24752)
- **MySQL Cluster:** A committed read could be attempted before a data node had time to connect, causing a timeout error. (Bug #24717)
- **MySQL Cluster:** The simultaneous shutdown of `mysqld` and `ndbd` processes caused unnecessary locking. (Bug #24655)
- **MySQL Cluster:** The failure of the master node in a node group during the allocation of node IDs could cause `ndb_mgmd` to hang. (Bug #24543)
- **MySQL Cluster:** In certain rare cases, a data node could crash due to a typographical error in the MySQL Cluster source code. (Bug #24476)
- **MySQL Cluster:** Creating a new tables containing a `BLOB` column when the server was short of memory could cause the server to crash. (Bug #24470)
- **MySQL Cluster:** Sudden disconnection of an SQL or data node could lead to shutdown of data nodes with the error `failed ndbrequire`. (Bug #24447)
- **MySQL Cluster:** Any statement following the execution of `CREATE TABLE ... LIKE ndb_table` (where `ndb_table` was a table using the `NDB` storage engine), would cause the `mysql` client to hang. (Bug #24301)
- **MySQL Cluster:** (Disk Data): Excessive fragmentation of Disk Data files (including log files and data files) could occur during the course of normal use. (Bug #24143)
- **MySQL Cluster:** When the management client command `ALL RESTART -i` was executed while one data node was not running, all data nodes in the cluster were shut down. (Bug #24105)
- **MySQL Cluster:** A query using an index scan followed by a delete operation, and then a rollback could cause one or more data nodes to crash. (Bug #24039)
- **MySQL Cluster:** (Disk Data): Under some circumstances, a `DELETE` from a Disk Data table could cause `mysqld` to crash. (Bug #23542)
- **MySQL Cluster:** It was possible for the sum of the `MaxNoOfTables`, `MaxNoOfOrderedIndexes`, and `MaxNoOfUniqueHashIndexes` configuration parameters, plus the number of system tables to exceed the maximum value for a `Uint32` number. In such a case, the cluster's data nodes failed to start, and no reason for this could easily be determined from the error messages provided. (Bug #22548)
- **MySQL Cluster:** A value equal to or greater than the permitted maximum for `LongMessageBuffer` caused all data nodes to crash. (Bug #22547)
- **MySQL Cluster:** Multiple occurrences of error conditions were logged with duplicat error messages rather than being reported with a single error message stating that the error was encountered `N` times. (Bug #22313)
- **MySQL Cluster:** Given a table `mytbl` in a database `mydb` on a MySQL Server acting as an SQL node in a MySQL Cluster, then, following multiple `ALTER TABLE mytbl ENGINE=engine` statements—first, to change the storage engine used for a table to `NDB`, and then again to change the table to use a non-`NDB` storage engine—a `DROP DATABASE mydb` statement executed on any SQL node in the cluster would cause `mydb` to be dropped on *all* SQL nodes in the cluster, even if `mydb` contained non-`NDB` tables. (Bug #21495)
- **MySQL Cluster:** An incorrect error message was displayed in the event that the value of the `MaxNoOfOrderedIndexes` parameter was set too low. (Bug #20065)

- **MySQL Cluster:** An incorrect error message was displayed in the event that the value of the `DataMemory` parameter was insufficient for the amount of data to be stored by the cluster. (Bug #19808)
- **MySQL Cluster:** Some values of `MaxNoOfTriggers` could cause the server to become inaccessible following startup of the data nodes. (Bug #19454)
- **MySQL Cluster:** If the value set for `MaxNoOfAttributes` is excessive, a suitable error message is now returned. (Bug #19352)
- **MySQL Cluster:** Different error messages were returned for similar cases involving failure to allocate memory for Cluster operations. (Bug #19203)
- **MySQL Cluster:** A unique constraint violation was not ignored by an `UPDATE IGNORE` statement when the constraint violation occurred on a nonprimary key. (Bug #18487, Bug #24303)
- **Replication:** With row-based binary logging, replicated multiple-statement transaction deadlocks did not return the correct error code, causing the slave SQL thread to stop rather than roll back and re-execute. (Bug #23831)
- **Replication:** Changes to character set variables prior to an action on a replication-ignored table were forgotten by slave servers. (Bug #22877)
- **Replication:** On slave servers, transactions that exceeded the lock wait timeout failed to roll back properly. (Bug #20697)
- **Replication:** SQL statements close to the size of `max_allowed_packet` could produce binary log events larger than `max_allowed_packet` that could not be read by slave servers. (Bug #19402)
- **Disk Data:** `ndb_restore` sometimes failed when attempting to restore Disk Data tables due to data node failure caused by accessing uninitialized memory. (Bug #24331)
- **Disk Data:** It was possible to execute a statement for creating a Disk Data table that referred to a nonexistent tablespace, in which case the table created was actually an in-memory `NDB` table. Such a statement now fails instead, with an appropriate error message. (Bug #23576)
- **Cluster API:** Using `BIT` values with any of the comparison methods of the `NdbScanFilter` class caused data nodes to fail. (Bug #24503)
- **Cluster API:** Some MGM API function calls could yield incorrect return values in certain cases where the cluster was operating under a very high load, or experienced timeouts in inter-node communications. (Bug #24011)
- In some cases, a function that should be parsed as a user-defined function was parsed as a stored function. (Bug #24736)
- Some unnecessary Valgrind warnings were removed from the server. (Bug #24488, Bug #24533)
- The server source code had multiple exportable definitions of the `field_in_record_is_null()` function. These are now all declared `static`. (Bug #24190)
- The loose index scan optimization for `GROUP BY` with `MIN` or `MAX` was not applied within other queries, such as `CREATE TABLE ... SELECT ...`, `INSERT ... SELECT ...`, or in the `FROM` clauses of subqueries. (Bug #24156)
- Subqueries for which a pushed-down condition did not produce exactly one key field could cause a server crash. (Bug #24056)
- The size of `MEMORY` tables and internal temporary tables was limited to 4GB on 64-bit Windows systems. (Bug #24052)

- `LAST_DAY('0000-00-00')` could cause a server crash. (Bug #23653)
- A trigger that invoked a stored function could cause a server crash when activated by different client connections. (Bug #23651)
- The stack size for NetWare binaries was increased to 128KB to prevent problems caused by insufficient stack size. (Bug #23504)
- If elements in a nontop-level `IN` subquery were accessed by an index and the subquery result set included a `NULL` value, the quantified predicate that contained the subquery was evaluated to `NULL` when it should return a non-`NULL` value. (Bug #23478)
- When applying the `group_concat_max_len` limit, `GROUP_CONCAT()` could truncate multibyte characters in the middle. (Bug #23451)
- `mysql_affected_rows()` could return values different from `mysql_stmt_affected_rows()` for the same sequence of statements. (Bug #23383)
- Calculation of `COUNT(DISTINCT)`, `AVG(DISTINCT)`, or `SUM(DISTINCT)` when they are referenced more than once in a single query with `GROUP BY` could cause a server crash. (Bug #23184)
- With row-based binary logging, for `CREATE TABLE IF NOT EXISTS LIKE temporary_table` statements, the `IF NOT EXISTS` clause was not logged. (Bug #22762)
- `BENCHMARK()`, `ENCODE()`, `DECODE()`, and `FORMAT()` could only accept a constant for some parameters, and could not be used in prepared statements. (Bug #22684)
- Queries using a column alias in an expression as part of an `ORDER BY` clause failed, an example of such a query being `SELECT mycol + 1 AS mynum FROM mytable ORDER BY 30 - mynum`. (Bug #22457)
- Using `EXPLAIN` caused a server crash for queries that selected from `INFORMATION_SCHEMA` in a subquery in the `FROM` clause. (Bug #22413)
- Instance Manager option-parsing code caused memory-allocation errors. (Bug #22242)
- Trailing spaces were not removed from Unicode `CHAR` column values when used in indexes. This resulted in excessive usage of storage space, and could affect the results of some `ORDER BY` queries that made use of such indexes.

**Note**

When upgrading, it is necessary to re-create any existing indexes on Unicode `CHAR` columns of each affected table to take advantage of the fix. See [Rebuilding or Repairing Tables or Indexes](#).

(Bug #22052)

- With row-based binary logging, `CREATE TABLE IF NOT EXISTS SELECT` statements were not logged properly. (Bug #22027)
- In some cases, the parser failed to distinguish a user-defined function from a stored function. (Bug #21809)
- Inserting a default or invalid value into a spatial column could fail with `Unknown error` rather than a more appropriate error. (Bug #21790)
- Through the C API, the member strings in `MYSQL_FIELD` for a query that contained expressions could return incorrect results. (Bug #21635)

- View columns were always handled as having implicit derivation, leading to `illegal mix of collation errors` for some views in `UNION` operations. Now view column derivation comes from the original expression given in the view definition. (Bug #21505)
- `INET_ATON()` returned a signed `BIGINT` value, not an unsigned value. (Bug #21466)
- For debug builds, `mysqladmin shutdown` displayed an extraneous `skipped 9 bytes from file: socket (3)` message. (Bug #21428)
- For renaming of views, encoding of table name to file names was not performed. (Bug #21370)
- `CREATE FUNCTION X()` and `CREATE FUNCTION Y()` failed with a syntax error instead of warning the user that these function names are already used (for GIS functions). (Bug #21025)
- `CONCURRENT` did not work correctly for `LOAD DATA INFILE`. (Bug #20637)
- With `lower_case_table_names` set to 1, `SHOW CREATE TABLE` printed incorrect output for table names containing Turkish I (LATIN CAPITAL LETTER I WITH DOT ABOVE). (Bug #20404)
- A query with a subquery that references columns of a view from the outer `SELECT` could return an incorrect result if used from a prepared statement. (Bug #20327)
- For queries that select from a view, the server returned `MYSQL_FIELD` metadata inconsistently for view names and table names. For view columns, the server now returns the view name in the `table` field and, if the column selects from an underlying table, the table name in the `org_table` field. (Bug #20191)
- Invalidating the query cache caused a server crash for `INSERT INTO ... SELECT` statements that selected from a view. (Bug #20045)
- For a cast of a `DATETIME` value containing microseconds to `DECIMAL`, the microseconds part was truncated without generating a warning. Now the microseconds part is preserved. (Bug #19491)
- The server could send incorrect column count information to the client for queries that produce a larger number of columns than can fit in a two-byte number. (Bug #19216)
- For some problems relating to character set conversion or incorrect string values for `INSERT` or `UPDATE`, the server reported truncation or length errors instead. (Bug #18908)
- Constant expressions and some numeric constants used as input parameters to user-defined functions were not treated as constants. (Bug #18761)
- Attempting to use a view containing `DEFINER` information for a nonexistent user resulted in an error message that revealed the definer account. Now the definer is revealed only to users that have the `SUPER` privilege. Other users receive only an `access denied` message. (Bug #17254)
- `IN()` and `CHAR()` can return `NULL`, but did not signal that to the query processor, causing incorrect results for `IS NULL` operations. (Bug #17047)
- Warnings were generated when explicitly casting a character to a number (for example, `CAST('x' AS SIGNED)`), but not for implicit conversions in simple arithmetic operations (such as `'x' + 0`). Now warnings are generated in all cases. (Bug #11927)
- Metadata for columns calculated from scalar subqueries was limited to integer, double, or string, even if the actual type of the column was different. (Bug #11032)

Changes in MySQL 5.1.13 (Not released)

This is a new Beta development release, fixing recently discovered bugs.

**Note**

This Beta 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.

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 Enterprise* (a commercial MySQL offering). For more details, please see (<http://www.mysql.com/products/enterprise>).

- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

Functionality Added or Changed

- **Incompatible Change:** The number of function names affected by `IGNORE_SPACE` was reduced significantly in MySQL 5.1.13, from about 200 to about 30. (For details about `IGNORE_SPACE`, see [Function Name Parsing and Resolution](#).) This change improves the consistency of parser operation. However, it also introduces the possibility of incompatibility for old SQL code that relies on the following conditions:

- `IGNORE_SPACE` is disabled.
- The presence or absence of whitespace following a function name is used to distinguish between a built-in function and stored function that have the same name (for example, `PI()` versus `PI ()`).

For functions that are no longer affected by `IGNORE_SPACE` as of MySQL 5.1.13, that strategy no longer works. Either of the following approaches can be used if you have code that is subject to the preceding incompatibility:

- If a stored function has a name that conflicts with a built-in function, refer to the stored function with a schema name qualifier, regardless of whether whitespace is present. For example, write `schema_name.PI()` or `schema_name.PI ()`.
- Alternatively, rename the stored function to use a nonconflicting name and change invocations of the function to use the new name.

(Bug #21114)

- **Incompatible Change:** The `innodb_buffer_pool_ave_mem_mb` system variable has been removed and should no longer be used.
- **MySQL Cluster:** A change in the interfaces for the `INFORMATION_SCHEMA.FILES` table has made the table accessible to storage engines other than `NDB`. (Bug #23013)
- Binary distributions of MySQL 5.1.12 were built without support for partitioning. This has been corrected except for NetWare. (Bug #23949)
- If the user specified the server options `--max-connections=N` or `--table-open-cache=M`, a warning would be given in some cases that some values were recalculated, with the result that `--table-open-cache` could be assigned greater value.

In such cases, both the warning and the increase in the `--table-open-cache` value were completely harmless. Note also that it is not possible for the MySQL Server to predict or to control limitations on the maximum number of open files, since this is determined by the operating system.

The value of `--table-open-cache` is no longer increased automatically, and a warning is now given only if some values had to be decreased due to operating system limits. (Bug #21915)

- For the `CALL` statement, stored procedures that take no arguments now can be invoked without parentheses. That is, `CALL p()` and `CALL p` are equivalent. (Bug #21462)
- `mysql_upgrade` now passes all the parameters specified on the command line to both `mysqlcheck` and `mysql` using the `upgrade_defaults` file. (Bug #20100)
- `mysqldump --single-transaction` now uses `START TRANSACTION /*!40100 WITH CONSISTENT SNAPSHOT */` rather than `BEGIN` to start a transaction, so that a consistent snapshot will be used on those servers that support it. (Bug #19660)

Bugs Fixed

- **Performance:** InnoDB showed substandard performance with multiple queries running concurrently. (Bug #15815)
- **Important Change:** When installing MySQL on AIX 5.3, you must upgrade AIX to technology level 7 (5300-07) to ensure the required thread libraries are available.
- **MySQL Cluster:** Backup of a cluster failed if there were any tables with 128 or more columns. (Bug #23502)
- **MySQL Cluster:** Cluster backups failed when there were more than 2048 schema objects in the cluster. (Bug #23499)
- **MySQL Cluster:** Restoring a cluster failed if there were any tables with 128 or more columns. (Bug #23494)
- **MySQL Cluster:** The management client command `ALL DUMP 1000` would cause the cluster to crash if data nodes were connected to the cluster but not yet fully started. (Bug #23203)
- **MySQL Cluster:** `INSERT ... ON DUPLICATE KEY UPDATE` on an NDB table could lead to deadlocks and memory leaks. (Bug #23200)
- **MySQL Cluster:** An NDB source file included a `memset()` call with reversed arguments. (Bug #23169)
- **MySQL Cluster:** If a node restart could not be performed from the REDO log, no node takeover took place. This could cause partitions to be left empty during a system restart. (Bug #22893)
- **MySQL Cluster:** Multiple node restarts in rapid succession could cause a system restart to fail, or induce a race condition. (Bug #22892, Bug #23210)
- **MySQL Cluster:** Attempting to create a unique constraint with `USING HASH` on an NDB table caused `mysqld` to crash. (Bug #21873)
- **MySQL Cluster:** When inserting a row into an NDB table with a duplicate value for a nonprimary unique key, the error issued would reference the wrong key. (Bug #21072)
- **MySQL Cluster:** Aborting a cluster backup too soon after starting it caused a forced shutdown of the data nodes. (Bug #19148)

- **Replication:** Column names were not quoted properly for replicated views. (Bug #19736)
- **Replication:** Transient errors in replication from master to slave may trigger multiple `Got fatal error 1236: 'binlog truncated in the middle of event'` errors on the slave. (Bug #4053)
- **Disk Data:** In the event of an aborted multiple update, the space in the Disk Data log buffer to be freed as a result was actually freed twice, which could eventually lead to a crash. (Bug #23430)
- **Cluster API:** When multiple processes or threads in parallel performed the same ordered scan with exclusive lock and updated the retrieved records, the scan could skip some records, which as a result were not updated. (Bug #20446)
- `FORMAT(X,D)` did not accept a nonconstant value for `D`. (Bug #48374)
- There was a race condition in the `InnoDB fil_flush_file_spaces()` function. (Bug #24098)
- yaSSL-related memory leaks were detected by Valgrind. (Bug #23981)
- MySQL 5.0.26 introduced an ABI incompatibility, which this release reverts. Programs compiled against 5.0.26 are not compatible with any other version and must be recompiled. (Bug #23427)
- `M % 0` returns `NULL`, but `(M % 0) IS NULL` evaluated to false. (Bug #23411)
- For not-yet-authenticated connections, the `Time` column in `SHOW PROCESSLIST` was a random value rather than `NULL`. (Bug #23379)
- `InnoDB` crashed when trying to display an error message about a foreign key constraint violation when the two tables are in different schemas. (Bug #23368)
- MySQL failed to build on Linux/Alpha. (Bug #23256)

References: This issue is a regression of: Bug #21250.

- If `COMPRESS()` returned `NULL`, subsequent invocations of `COMPRESS()` within a result set or within a trigger also returned `NULL`. (Bug #23254)
- Insufficient memory (`myisam_sort_buffer_size`) could cause a server crash for several operations on `MyISAM` tables: repair table, create index by sort, repair by sort, parallel repair, bulk insert. (Bug #23175)
- The column default value in the output from `SHOW COLUMNS` or `SELECT FROM INFORMATION_SCHEMA.COLUMNS` was truncated to 64 characters. (Bug #23037)
- `mysql` did not check for errors when fetching data during result set printing. (Bug #22913)
- The return value from `my_seek()` was ignored. (Bug #22828)
- Use of `SQL_BIG_RESULT` did not influence the sort plan for query execution. (Bug #22781)
- The optimizer failed to use equality propagation for `BETWEEN` and `IN` predicates with string arguments. (Bug #22753)
- The `Handler_rollback` status variable sometimes was incremented when no rollback had taken place. (Bug #22728)
- The `Host` column in `SHOW PROCESSLIST` output was blank when the server was started with the `--skip-grant-tables` option. (Bug #22723)

- If a table contains an `AUTO_INCREMENT` column, inserting into an insertable view on the table that does not include the `AUTO_INCREMENT` column should not change the value of `LAST_INSERT_ID()`, because the side effects of inserting default values into columns not part of the view should not be visible. MySQL was incorrectly setting `LAST_INSERT_ID()` to zero. (Bug #22584)
- The optimizer used the `ref` join type rather than `eq_ref` for a simple join on strings. (Bug #22367)
- Some queries that used `MAX()` and `GROUP BY` could incorrectly return an empty result. (Bug #22342)
- If an `init_connect` SQL statement produced an error, the connection was silently terminated with no error message. Now the server writes a warning to the error log. (Bug #22158)
- An unhandled `NULL` pointer caused a server crash. (Bug #22138)
- Incorrect warnings occurred for use of `CREATE TABLE ... LIKE` or `REPAIR TABLE` with the log tables. (Bug #21966)
- The optimizer sometimes mishandled R-tree indexes for `GEOMETRY` data types, resulting in a server crash. (Bug #21888)
- Use of a DES-encrypted SSL certificate file caused a server crash. (Bug #21868)
- Use of `PREPARE` with a `CREATE PROCEDURE` statement that contained a syntax error caused a server crash. (Bug #21856)
- Adding a day, month, or year interval to a `DATE` value produced a `DATE`, but adding a week interval produced a `DATETIME` value. Now all produce a `DATE` value. (Bug #21811)
- Use of a subquery that invoked a function in the column list of the outer query resulted in a memory leak. (Bug #21798)
- It was not possible to do an atomic rename of the log tables without the possibility of losing rows. Now you can do this:

```
USE mysql;  
CREATE TABLE IF NOT EXISTS general_log2 LIKE general_log;  
RENAME TABLE general_log TO general_log_backup, general_log2 TO general_log;
```

(Bug #21785, Bug #17544)

- Within a prepared statement, `SELECT (COUNT(*) = 1)` (or similar use of other aggregate functions) did not return the correct result for statement re-execution. (Bug #21354)
- Within a stored routine, a view definition cannot refer to routine parameters or local variables. However, an error did not occur until the routine was called. Now it occurs during parsing of the routine creation statement.

**Note**

A side effect of this fix is that if you have already created such routines, and error will occur if you execute `SHOW CREATE PROCEDURE` or `SHOW CREATE FUNCTION`. You should drop these routines because they are erroneous.

(Bug #20953)

- In `mysql`, invoking `connect` or `\r` with very long `db_name` or `host_name` parameters caused buffer overflow. (Bug #20894)

- `WITH ROLLUP` could group unequal values. (Bug #20825)
- Range searches on columns with an index prefix could miss records. (Bug #20732)
- The server did not allocate sufficient memory for some queries for which a `DISTINCT` to `GROUP BY` conversion is possible and an `ORDER BY` clause is present, resulting in a server crash. (Bug #20503)
- `LIKE` searches failed for indexed `utf8` character columns. (Bug #20471)
- With `sql_mode = TRADITIONAL`, MySQL incorrectly aborted on warnings within stored routines and triggers. (Bug #20028)
- `mysqldump --xml` produced invalid XML for `BLOB` data. (Bug #19745)
- The range analysis optimizer did not take into account predicates for which an index could be used after reading `const` tables. In some cases this resulted in nonoptimal execution plans. (Bug #19579)
- `FLUSH INSTANCES` in Instance Manager triggered an assertion failure. (Bug #19368)
- For a debug server, a reference to an undefined user variable in a prepared statement executed with `EXECUTE` caused an assertion failure. (Bug #19356)
- Within a trigger for a base table, selecting from a view on that base table failed. (Bug #19111)
- The value of the `warning_count` system variable was not being calculated correctly (also affecting `SHOW COUNT(*) WARNINGS`). (Bug #19024)
- `DELETE IGNORE` could hang for foreign key parent deletes. (Bug #18819)
- `InnoDB` used table locks (not row locks) within stored functions. (Bug #18077)
- `mysql` would lose its connection to the server if its standard output was not writable. (Bug #17583)
- At shutdown, Instance Manager told guarded server instances to stop, but did not wait until they actually stopped. (Bug #17486)
- `mysql-test-run` did not work correctly for RPM-based installations. (Bug #17194)
- A client library crash was caused by executing a statement such as `SELECT * FROM t1 PROCEDURE ANALYSE()` using a server side cursor on a table `t1` that does not have the same number of columns as the output from `PROCEDURE ANALYSE()`. (Bug #17039)
- The `WITH CHECK OPTION` for a view failed to prevent storing invalid column values for `UPDATE` statements. (Bug #16813)
- `ALTER TABLE` was not able to rename a view. (Bug #14959)
- Statements such as `DROP PROCEDURE` and `DROP VIEW` were written to the binary log too late due to a race condition. (Bug #14262)
- A literal string in a `GROUP BY` clause could be interpreted as a column name. (Bug #14019)
- Entries in the slow query log could have an incorrect `Rows_examined` value. (Bug #12240)
- Lack of validation for input and output `TIME` values resulted in several problems: `SEC_TO_TIME()` in some cases did not clip large values to the `TIME` range appropriately; `SEC_TO_TIME()` treated `BIGINT UNSIGNED` values as signed; only truncation warnings were produced when both truncation and out-of-range `TIME` values occurred. (Bug #11655, Bug #20927)

- Several string functions could return incorrect results when given very large length arguments. (Bug #10963)
- `FROM_UNIXTIME()` did not accept arguments up to `POWER(2,31)-1`, which it had previously. (Bug #9191)
- `OPTIMIZE TABLE` with `myisam_repair_threads > 1` could result in `MyISAM` table corruption. (Bug #8283)

Changes in MySQL 5.1.12 (2006-10-24)

This is a new Beta development release, fixing recently discovered bugs.

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 Enterprise* (a commercial MySQL offering). For more details, please see (<http://www.mysql.com/products/enterprise>).

- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

Functionality Added or Changed

- **Incompatible Change; MySQL Cluster:** MySQL Cluster node and system restarts formerly required that all fragments use the same local checkpoint (LCP); beginning with this version, it is now possible for different fragments to use different LCPs during restarts. This means that data node file systems must be rebuilt as part of any upgrade to this version by restarting all data nodes with the `--initial` option.

See [Upgrade and downgrade compatibility: MySQL 5.1](#), and related sections of the Manual before upgrading a MySQL Cluster to version 5.1.12 or later. (Bug #21478, Bug #21271)

- **Incompatible Change:** In the `INFORMATION_SCHEMA.EVENTS` table, the `EVENT_DEFINITION` column now contains the SQL executed by a scheduled event.

The `EVENT_BODY` column now contains the language used for the statement or statements shown in `EVENT_DEFINITION`. In MySQL 5.1, the value shown in `EVENT_BODY` is always `SQL`.

These changes were made to bring this table into line with the `INFORMATION_SCHEMA.ROUTINES` table, and that table's `ROUTINE_BODY` and `ROUTINE_DEFINITION` columns. (Bug #16992)

- **Incompatible Change:** For `GRANT` and `REVOKE`, `ON *` previously granted and revoked privileges for the default database if there was a default database and global privileges if there was none. Now `ON *` requires a default database and produces an error if there is none.
- **Incompatible Change:** Support for the BerkeleyDB (BDB) engine has been dropped from this release. Any existing tables that are in BDB format will not be readable from within MySQL from 5.1.12 or newer. You should convert your tables to another storage engine *before* upgrading to 5.1.12.

Because of this change, the `SHOW [BDB] LOGS` statement has been dropped.

- **Incompatible Change:** A number of MySQL constructs are now prohibited in partitioning expressions, beginning with this release. These include the following:
 - A number of MySQL functions. For a complete list of these, see [Partitioning Limitations Relating to Functions](#).

- The bit operators `|`, `&`, `^`, `<<`, `>>`, and `~`.
- Nested function calls.
- Calls to stored routines, UDFs, or plugins.
- Character-to-integer conversions involving non-8-bit character sets or any of the `latin1_german2_ci`, `latin2_czech_cs`, or `cp1250_czech_cs` collations.

These restrictions were added in part as a result of Bug #18198 and related bug reports.

For more information about these and other restrictions on partitioned tables in MySQL, see [Restrictions and Limitations on Partitioning](#).

References: This issue is a regression of: Bug #18198.

- **Incompatible Change:** The permitted values for and behavior of the `event_scheduler` system variable have changed. Permitted values are now `ON`, `OFF`, and `DISABLED`, with `OFF` being the default. It is not possible to change its value to or from `DISABLED` while the server is running.

For details, see [Event Scheduler Overview](#).

- **Incompatible Change:** The plugin interface has changed: The `st_mysql_plugin` structure has a new `license` member to indicate the license type. (The permissible values are defined in `mysql/plugin.h`.) This change is not backward compatible, so the API version (`MYSQL_PLUGIN_INTERFACE_VERSION`) has changed. For additional information, see [Writing Plugins](#).

- **Incompatible Change:** The full-text parser plugin interface has changed in two ways:

- The `MYSQL_FTPARSER_PARAM` structure has a new `flags` member. This is zero if there are no special flags, or `MYSQL_FTFLAGS_NEED_COPY`, which means that `mysql_add_word()` must save a copy of the word (that is, it cannot use a pointer to the word because the word is in a buffer that will be overwritten.)

This flag might be set or reset by MySQL before calling the parser plugin, by the parser plugin itself, or by the `mysql_parse()` function.

- The `mysql_parse()` and `mysql_add_word()` functions now take a `MYSQL_FTPARSER_PARAM` as their first argument, not a `MYSQL_FTPARSER_PARAM:mysql_ftparam` as before.

These changes are not backward compatible, so the API version (`MYSQL_FTPARSER_INTERFACE_VERSION`) has changed. For additional information, see [Writing Plugins](#).

- **Incompatible Change:** Storage engines can be pluggable at runtime, so the distinction between disabled and invalid storage engines no longer applies. This affects the `NO_ENGINE_SUBSTITUTION` SQL mode, as described in [Server SQL Modes](#).
- **Incompatible Change:** The namespace for scheduled events has changed, such that events are no longer unique to individual users. This also means that a user with the `EVENT` privilege on a given database can now view, alter, or drop any events defined on that database.

If you used scheduled events in an earlier MySQL 5.1 release, you should rename any of them having the same name and defined on the same database but belonging to different users—so that all events in a given database have unique names—*before* upgrading to 5.1.12 (or newer).

For additional information, see [The Event Scheduler and MySQL Privileges](#).

- **Important Change; MySQL Cluster; Partitioning:** It is no longer possible to create Cluster tables using any partitioning type other than `[LINEAR] KEY`. Attempting to do so now raises an error.
- **Important Change; MySQL Cluster:** `LOAD DATA INFILE` no longer causes an implicit commit for all storage engines. It now causes an implicit commit only for tables using the `NDB` storage engine. (Bug #11151)
- **Important Change; MySQL Cluster:** The status variables `Ndb_connected_host` and `Ndb_connected_port` were renamed to `Ndb_config_from_host` and `Ndb_config_from_port`, respectively.
- **Important Change; Replication:** The default value for the `--binlog-format` server option is now `MIXED`.
- **MySQL Cluster:** The `ndb_config` utility now accepts `-c` as a short form of the `--ndb-connectstring` option. (Bug #22295)
- **MySQL Cluster:** Added the `--bind-address` option for `ndbd`. This permits a data node process to be bound to a specific network interface. (Bug #22195)
- **MySQL Cluster:** The `Ndb_number_of_storage_nodes` system variable was renamed to `Ndb_number_of_data_nodes`. (Bug #20848)
- **MySQL Cluster:** The `HELP` command in the Cluster management client now provides command-specific help. For example, `HELP RESTART` in `ndb_mgm` provides detailed information about the `RESTART` command. (Bug #19620)
- **MySQL Cluster:** A number of erroneous, misleading, or missing error messages have been corrected. (Bug #17297, Bug #19543)
- **MySQL Cluster:** Backup messages are no longer printed to the cluster log.
- **MySQL Cluster:** Added the `--ndb-use-copying-alter-table` option to `mysqld` to provide a fallback in case of problems with online `ALTER TABLE` operations on `NDB` tables.
- **Replication:** The default binary log format (as used during replication) is now Mixed based, automatically using a combination of row-based and statement based log events as appropriate.
- **Cluster API:** Two new NDB API methods `aggregate()` and `validate()` were added to the `Table` class. This was done to rectify the following issues:
 - Under some conditions, the data distribution could become unbalanced in a MySQL Cluster with 2 or more node groups following the creation of a new table.
 - Data was stored unevenly between partitions due to all `BLOB` data being placed in partition 0.(Bug #21690)
- The number of `InnoDB` threads is no longer limited to 1,000 on Windows. (Bug #22268)
- The `STATE` column of the `INFORMATION_SCHEMA.PROCESSLIST` table was increased from 30 to 64 characters to accommodate longer state values. (Bug #21652)
- `mysqldump` now has a `--flush-privileges` option. It causes `mysqldump` to emit a `FLUSH PRIVILEGES` statement after dumping the `mysql` database. This option should be used any time the dump contains the `mysql` database and any other database that depends on the data in the `mysql` database for proper restoration. (Bug #21424)

- `mysqslap` threads now try to connect up to 10 times if the initial connect attempt fails. (Bug #21297)
- For `mysqldump`, the output generated by the server when using the `--xml` option has changed with regard to null values. It now matches the output from `mysqldump --xml`. That is, a column containing a `NULL` value is now reported as

```
<field name="column_name" xsi:nil="true" />
```

whereas a column containing the string value '`NULL`' is reported as

```
<field name="column_name">NULL</field>
```

and a column containing an empty string is reported as

```
<field name="column_name"></field>
```

(Bug #21263)

- The `mysqld` and `mysqlmanager` man pages have been reclassified from volume 1 to volume 8. (Bug #21220)
- `InnoDB` now honors `IGNORE INDEX`. Previously using `IGNORE INDEX` in cases where an index sort would be slower than a filesort had no effect when used with `InnoDB` tables. (Bug #21174)
- `TIMESTAMP` columns that are `NOT NULL` now are reported that way by `SHOW COLUMNS` and `INFORMATION_SCHEMA`. (Bug #20910)
- Memory consumption of the `InnoDB` data dictionary cache was roughly halved by cleaning up the data structures. (Bug #20877)
- The `BINARY` keyword now is forbidden as a data type attribute in stored routines (for example, `DECLARE v1 VARCHAR(25) BINARY`), because `DECLARE` does not support collations, and in this context `BINARY` specifies the binary collation of the variable's character set. (Bug #20701)
- The following statements now can be executed as prepared statements (using `PREPARE` plus `EXECUTE`):

```
CACHE INDEX
CHANGE MASTER
CHECKSUM {TABLE | TABLES}
{CREATE | RENAME | DROP} DATABASE
{CREATE | RENAME | DROP} USER
FLUSH {TABLE | TABLES | TABLES WITH READ LOCK | HOSTS | PRIVILEGES
      | LOGS | STATUS | MASTER | SLAVE | DES_KEY_FILE | USER_RESOURCES}
GRANT
REVOKE
KILL
LOAD INDEX INTO CACHE
RESET {MASTER | SLAVE | QUERY CACHE}
SHOW BINLOG EVENTS
SHOW CREATE {PROCEDURE | FUNCTION | EVENT | TABLE | VIEW}
SHOW {AUTHORS | CONTRIBUTORS | WARNINGS | ERRORS}
SHOW {MASTER | BINARY} LOGS
SHOW {MASTER | SLAVE} STATUS
SLAVE {START | STOP}
INSTALL PLUGIN
UNINSTALL PLUGIN
```

(Bug #20665)

- In the `INFORMATION_SCHEMA.ROUTINES` table the `ROUTINE_DEFINITION` column now is defined as `NULL` rather than `NOT NULL`. Also, `NULL` rather than the empty string is returned as the column value if the user does not have sufficient privileges to see the routine definition. (Bug #20230)
- 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)
- The MySQL distribution now compiles on UnixWare 7.13. (Bug #20190)
- `configure` now defines the symbol `DEBUG_ON` in `config.h` to indicate whether the source tree is configured to be compiled with debugging support. (Bug #19517)
- `TEXT` and `BLOB` columns do not support `DEFAULT` values. However, when a default of `' '` was specified, the specification was silently ignored. This now results in a warning, or an error in strict mode. (Bug #19498)
- For `mysqlshow`, if a database name argument contains wildcard characters (such as `"_"`) but matches a single database name exactly, treat the name as a literal name. This enables a command such as `mysqlshow information_schema` to work without having to escape the wildcard character. (Bug #19147)
- The source distribution has been updated so that the UDF example can be compiled under Windows with `CMake`. See [UDF Compiling and Installing](#). (Bug #19121)
- The default value of the `tmp_table_size` system variable was lowered from 32MB to 16MB because it is bounded by the value of `max_heap_table_size`, which has a default of 16MB. (Bug #18875)
- Log table changes: By default, the log tables use the `CSV` storage engine, as before. But now the log tables can be altered to use the `MyISAM` storage engine. You cannot use `ALTER TABLE` to alter a log table that is in use. The log must be disabled first. No engines other than `CSV` or `MyISAM` are legal for the log tables. The use of `DROP TABLE` for log tables is similarly restricted: It cannot be used to drop a log table that is in use. The log must be disabled first. (These changes also correct a deadlock that occurred for an attempt to drop an in-use log table.) (Bug #18559)
- Added the `--set-charset` option to `mysqlbinlog` to enable the character set to be specified for processing binary log files. (Bug #18351)
- The `ExtractValue()` function now produces an error when passed an XML fragment that is not well-formed.

(Previously, the function permitted invalid XML fragments to be used.) (Bug #18201)
- On Windows, typing **Control+C** while a query was running caused the `mysql` client to crash. Now it causes `mysql` to attempt to kill the current statement. If this cannot be done, or **Control+C** is typed again before the statement is killed, `mysql` exits. (In other words, `mysql`'s behavior with regard to **Control+C** is now the same as it is on Unix platforms.) (Bug #17926)

References: See also: Bug #1989.

- The bundled yaSSL library licensing has added a FLOSS exception similar to MySQL to resolve licensing incompatibilities with MySQL. (See the `extra/yassl/FLOSS-EXCEPTIONS` file in a MySQL source distribution for details.) (Bug #16755)
- `SHOW CREATE TABLE` now shows constraints for `InnoDB` tables. (Bug #16614)
- `EXPLAIN EXTENDED` now shows a `filtered` column that is an estimated percentage of the examined rows that will be joined with the previous tables. This was added while dealing with a problem of MySQL choosing the wrong index for some queries. (Bug #14940)

- The `mysql` client now permits `\1` in the `prompt` command argument, to insert the current delimiter into the prompt. (Bug #14448)
- The `mysql` client used the default character set if it automatically reconnected to the server, which is incorrect if the character set had been changed. To enable the character set to remain synchronized on the client and server, the `mysql` command `charset` (or `\C`) that changes the default character set and now also issues a `SET NAMES` statement. The changed character set is used for reconnects. (Bug #11972)
- The `LEFT()` and `RIGHT()` functions return `NULL` if any argument is `NULL`. (Bug #11728)
- If a `DROP VIEW` statement named multiple views, it stopped with an error if a nonexistent view was named and did not drop the remaining views. Now it continues on and reports an error at the end, similar to `DROP TABLE`. (Bug #11551)
- For a successful dump, `mysqldump` now writes a SQL comment to the end of the dump file in the following format:

```
-- Dump completed on YYYY-MM-DD hh:mm:ss
```

(Bug #10877)

- There were several issues regarding how `SHOW STATUS` affected some status variables and logging which could impact monitoring the MySQL Server. The behavior of this statement has been modified in two ways:
 - `SHOW STATUS` is no longer logged to the slow query log.
 - `SHOW STATUS` no longer updates any session status variables, except for `com_show_status`.

However, `SHOW STATUS` continues to update *global* status variables to enable monitoring of what the server is actually doing. This is because `SHOW STATUS` creates temporary tables that may affect performance if it is called excessively often. (Bug #10210)

References: See also: Bug #19764.

- 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)
- The `LOAD DATA FROM MASTER` and `LOAD TABLE FROM MASTER` statements are deprecated. See [LOAD DATA FROM MASTER Syntax](#), for recommended alternatives. (Bug #9125, Bug #20596, Bug #14399, Bug #12187, Bug #15025, Bug #18822)
- It is no longer possible to create partitioned tables using the `CSV` storage engine.
- Binary MySQL distributions no longer include a `mysqld-max` server. Instead, distributions contain a binary that includes the features previously included in the `mysqld-max` binary.
- `SHOW STATUS` is no longer logged to the slow query log.
- Program Database (PDB) files (with file name extension `.pdb`) are now included by default in Windows distributions. These can be used to help diagnose problems with `mysqld` and other tools. See [Debugging a MySQL Server](#).
- `INFORMATION_SCHEMA` contains new tables, `GLOBAL_STATUS`, `SESSION_STATUS`, `GLOBAL_VARIABLES`, and `SESSION_VARIABLES`, that correspond to the output from the `SHOW {GLOBAL|SESSION} STATUS` and `SHOW {GLOBAL|SESSION} VARIABLES` statements.

- `SHOW STATUS` no longer updates any session status variables, except for `com_show_status`.
- A new system variable, `lc_time_names`, specifies the locale that controls the language used to display day and month names and abbreviations. This variable affects the output from the `DATE_FORMAT()`, `DAYNAME()` and `MONTHNAME()` functions. See [MySQL Server Locale Support](#).
- Using `--with-debug` to configure MySQL with debugging support enables you to use the `--debug="d,parser_debug"` option when you start the server. This causes the Bison parser that is used to process SQL statements to dump a parser trace to the server's standard error output. Typically, this output is written to the error log.
- For the `mysql` client, typing **Control+C** causes `mysql` to attempt to kill the current statement. If this cannot be done, or **Control+C** is typed again before the statement is killed, `mysql` exits. Previously, **Control+C** caused `mysql` to exit in all cases. (Bug #1989)
- The bundled yaSSL library was upgraded to version 1.3.7.
- The Instance Manager `--passwd` option has been renamed to `--print-password-line`. Other options were added to enable management of the IM password file from the command line: `--add-user`, `--drop-user`, `--edit-user`, `--list-users`, `--check-password-file`, `--clean-password-file`, `--username`, and `--password`. The `--mysqld-safe-compatible` option was added to cause the Instance Manager to act similarly to `mysqld_safe`.
- Added the `SHOW CONTRIBUTORS` statement.
- The general query log and slow query logs now can be enabled or disabled at runtime with the `general_log` and `slow_query_log` system variables, and the name of the log files can be changed by setting the `general_log_file` and `slow_query_log_file` system variables. See [The General Query Log](#), and [The Slow Query Log](#).

Bugs Fixed

- **Security Fix:** A stored routine created by one user and then made accessible to a different user using `GRANT EXECUTE` could be executed by that user with the privileges of the routine's definer. (Bug #18630, CVE-2006-4227)
- **Security Fix:** On Linux, and possibly other platforms using case-sensitive file systems, 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. (Bug #17647, CVE-2006-4226)
- **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, CVE-2006-4031)
- **Incompatible Change:** For `utf8` columns, the full-text parser incorrectly considered several nonword punctuation and whitespace characters as word characters, causing some searches to return incorrect results.

The fix involves a change to the full-text parser, so any tables that have `FULLTEXT` indexes on `utf8` columns must be repaired with `REPAIR TABLE`:

```
REPAIR TABLE tbl_name QUICK;
```

(Bug #19580)

- **MySQL Cluster; Packaging:** 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)
- **MySQL Cluster; Replication:** A `DELETE FROM table` with no `WHERE` clause (deleting all rows) running concurrently with `INSERT` statements on a storage engine with row-level locking (such as `NDB`) could produce inconsistent results when using statement-based replication. (Bug #19066)
- **MySQL Cluster; Replication:** (Replication): A node failure could send duplicate events, causing a `mysqld` replicating tables containing `BLOBs` to crash.
- **MySQL Cluster:** (NDB API): Inactivity timeouts for scans were not correctly handled. (Bug #23107)
- **MySQL Cluster:** Inserting into an `NDB` table failed when the table had no primary key but had a unique key added after table was created on one or more `NOT NULL` columns. This occurred when the unique key had been adding using either `ALTER TABLE` or `CREATE UNIQUE KEY`. (Bug #22838)
- **MySQL Cluster:** (NDB API): Attempting to read a nonexistent tuple using `Commit` mode for `NdbTransaction::execute()` caused node failures. (Bug #22672)
- **MySQL Cluster:** The `--help` output from `NDB` binaries did not include file-related options. (Bug #21994)
- **MySQL 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)
- **MySQL Cluster:** A scan timeout returned Error 4028 (Node failure caused abort of transaction) instead of Error 4008 (Node failure caused abort of transaction...). (Bug #21799)
- **MySQL Cluster:** The node recovery algorithm was missing a version check for tables in the `ALTER_TABLE_COMMITTED` state (as opposed to the `TABLE_ADD_COMMITTED` state, which has the version check). This could cause inconsistent schemas across nodes following node recovery. (Bug #21756)
- **MySQL Cluster:** A memory leak occurred when running `ndb_mgm -e "SHOW"`. (Bug #21670)
- **MySQL Cluster:** The server provided a nondescriptive error message when encountering a fatally corrupted REDO log. (Bug #21615)
- **MySQL Cluster:** The output for the `--help` option used with `NDB` executable programs (such as `ndbd`, `ndb_mgm`, `ndb_restore`, `ndb_config`, and others mentioned in [MySQL Cluster Programs](#)) referred to the `Ndb.cfg` file, instead of to `my.cnf`. (Bug #21585)
- **MySQL Cluster:** A partial rollback could lead to node restart failures. (Bug #21536)
- **MySQL Cluster:** Partition distribution keys were updated only for the primary and starting replicas during node recovery. This could lead to node failure recovery for clusters having an odd number of replicas.

**Note**

For best results, use values for `NumberOfReplicas` that are even powers of 2.

(Bug #21535)

- **MySQL Cluster:** The `ndb_mgm` management client did not set the exit status on errors, always returning 0 instead. (Bug #21530)

- **MySQL 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)
- **MySQL Cluster:** Attempting to create an `NDB` table on a MySQL server with an existing non-Cluster table with the same name in the same database could result in data loss or corruption. Now, if such a table is encountered during autodiscovery, a warning is written to the error log of the affected `mysqld`, and the local table is overwritten. (Bug #21378)
- **MySQL Cluster:** Cluster logs were not rotated following the first rotation cycle. (Bug #21345)
- **MySQL 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 the same node group to shut down. (Bug #21213)
- **MySQL Cluster:** The `ndb_size.pl` script did not account for `TEXT` and `BLOB` column values correctly. (Bug #21204)
- **MySQL Cluster:** Some queries involving joins on very large `NDB` tables could crash the MySQL server. (Bug #21059)
- **MySQL Cluster:** Condition pushdown did not work correctly with `DATETIME` columns. (Bug #21056)
- **MySQL Cluster:** Responses to the `ALL DUMP 1000` management client command were printed multiple times in the cluster log for each cluster node. (Bug #21044)
- **MySQL Cluster:** The message `Error 0 in readAutoIncrementValue(): no Error` was written to the error log whenever `SHOW TABLE STATUS` was performed on a Cluster table that did not have an `AUTO_INCREMENT` column. (Bug #21033)
- **MySQL Cluster:** Restarting a data node while DDL operations were in progress on the cluster could cause other data nodes to fail. This could also lead to `mysqld` hanging or crashing under some circumstances. (Bug #21017, Bug #21050)
- **MySQL 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)
- **MySQL Cluster:** A race condition could in some circumstances following a `DROP TABLE`. (Bug #20897)
- **MySQL Cluster:** Under some circumstances, local checkpointing would hang, keeping any unstarted nodes from being started. (Bug #20895)
- **MySQL 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)
- **MySQL Cluster:** A vague error message was returned when reading both schema files during a restart of the cluster. (Bug #20860)
- **MySQL 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)
- **MySQL Cluster:** When attempting to restart the cluster following a data import, the cluster failed during Phase 4 of the restart with `Error 2334: Job buffer congestion`. (Bug #20774)
- **MySQL Cluster:** `REPLACE` statements did not work correctly on an `NDB` table having both a primary key and a unique key. In such cases, proper values were not set for columns which were not explicitly referenced in the statement. (Bug #20728)
- **MySQL Cluster:** The server did not honor the value set for `ndb_cache_check_time` in the `my.cnf` file. (Bug #20708)

- **MySQL Cluster:** Truncating a table on one `mysqld` caused other `mysqld` processes connected to the cluster to return `ERROR 1412 (HY000): Table definition has changed, please retry transaction` on subsequent queries. (Bug #20705)
- **MySQL Cluster:** Using an invalid node ID with the management client `STOP` command could cause `ndb_mgm` to hang. (Bug #20575)
- **MySQL Cluster:** Renaming of table columns was not supported as fast as a `ALTER TABLE` for NDB tables. (Bug #20456)
- **MySQL Cluster:** `ndb_size.pl` and `ndb_error_reporter` were missing from RPM packages. (Bug #20426)
- **MySQL Cluster:** Running `ndbd --nowait-nodes=id` where `id` was the node ID of a node that was already running failed with an invalid error message. (Bug #20419)
- **MySQL Cluster:** Data nodes added while the cluster was running in single user mode were all assigned node ID 0, which could later cause multiple node failures. Adding nodes while in single user mode is no longer possible. (Bug #20395)
- **MySQL Cluster:** The `ndb_mgm` client command `ALL CLUSTERLOG STATISTICS=15` had no effect. (Bug #20336)
- **MySQL Cluster:** A node failure during a scan could sometime cause the node to crash when restarting too quickly following the failure. (Bug #20197)
- **MySQL 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)
- **MySQL Cluster:** `SHOW ENGINE NDB STATUS` could sometimes return an incorrect value of 0 for the latest epoch, which could cause problems with synchronizing the binlog. (Bug #20142)
- **MySQL Cluster:** An internal formatting error caused some management client error messages to be unreadable. (Bug #20016)
- **MySQL Cluster:** Creating tables with variable-size columns caused `DataMemory` to be used but not freed when the tables were dropped. (Bug #20007)
- **MySQL 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)
- **MySQL Cluster:** Running management client commands while `mgmd` was in the process of disconnecting could cause the management server to fail. (Bug #19932)
- **MySQL Cluster:** Under certain conditions, a starting node could miss transactions, leading to inconsistencies between the primary and backup replicas. (Bug #19929)
- **MySQL Cluster:** An uncommitted row could sometimes be checkpointed and thus incorrectly included in a backup. (Bug #19928)
- **MySQL Cluster:** In some cases where `SELECT COUNT(*)` from an NDB table should have yielded an error, `MAX_INT` was returned instead. (Bug #19914)
- **MySQL Cluster:** `TEXT` columns in Cluster tables having both an explicit primary key and a unique key were not correctly updated by `REPLACE` statements. (Bug #19906)
- **MySQL Cluster:** The cluster's data nodes failed while trying to load data when `NoOfFrangmentLogFiles` was set equal to 1. (Bug #19894)

- **MySQL Cluster:** Following the restart of a management node, the Cluster management client did not automatically reconnect. (Bug #19873)
- **MySQL Cluster:** Restoring a backup with `ndb_restore` failed when the backup had been taken from a cluster whose `DataMemory` had been completely used up. (Bug #19852)
- **MySQL Cluster:** Error messages given when trying to make online changes to parameters such as `NoOfReplicas` that can only be changed using a complete shutdown and restart of the cluster did not indicate the true nature of the problem. (Bug #19787)
- **MySQL Cluster:** Under some circumstances, repeated DDL operations on one `mysqld` could cause failure of a second `mysqld` attached to the same cluster. (Bug #19770)
- **MySQL Cluster:** `ndb_restore` did not always make clear that it had recovered successfully from temporary errors while restoring a cluster backup. (Bug #19651)
- **MySQL 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)
- **MySQL Cluster:** `LOAD DATA LOCAL` failed to ignore duplicate keys in Cluster tables. (Bug #19496)
- **MySQL Cluster:** For `ndb_mgmd`, Valgrind revealed problems with a memory leak and a dependency on an uninitialized variable. (Bug #19318, Bug #20333)
- **MySQL Cluster:** A `DELETE` of many rows immediately followed by an `INSERT` on the same table could cause the `ndbd` process on the backup replica to crash. (Bug #19293)
- **MySQL Cluster:** An excessive number of `ALTER TABLE` operations could cause the cluster to fail with `NDB` error code 773 (`Out of string memory, please modify StringMemory`). (Bug #19275)
- **MySQL 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)
- **MySQL Cluster:** In rare situations with resource shortages, a crash could result from an insufficient number of `IndexScanOperation` objects. (Bug #19198)
- **MySQL Cluster:** Running out of `DataMemory` could sometimes crash `ndbd` and `mysqld` processes. (Bug #19185)
- **MySQL Cluster:** It was possible to use port numbers greater than 65535 for `ServerPort` in the `config.ini` file. (Bug #19164)
- **MySQL Cluster:** `ndb_mgm -e show | head` would hang after displaying the first 10 lines of output. (Bug #19047)
- **MySQL Cluster:** The error returned by the cluster when too many nodes were defined did not make clear the nature of the problem. (Bug #19045)
- **MySQL Cluster:** The management client `ALL STOP` command shut down `mgmd` processes (as well as `ndbd` processes). (Bug #18966)
- **MySQL Cluster:** `TRUNCATE TABLE` failed to reset the `AUTO_INCREMENT` counter. (Bug #18864)
- **MySQL Cluster:** Restarting a failed node could sometimes crash the cluster. (Bug #18782)
- **MySQL Cluster:** Trying to create or drop a table while a node was restarting caused the node to crash. This is now handled by raising an error. (Bug #18781)

- **MySQL Cluster:** Repeated [CREATE](#) - [INSERT](#) - [DROP](#) operations on 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)
- **MySQL Cluster:** A [CREATE TABLE](#) statement involving foreign key constraints raised an error rather than being silently ignored (see [CREATE TABLE Syntax](#)).

This bug affected Cluster in MySQL 5.1 only. (Bug #18483)
- **MySQL Cluster:** The server failed with a nondescriptive error message when out of data memory. (Bug #18475)
- **MySQL Cluster:** For [NDB](#) and possibly [InnoDB](#) tables, a [BEFORE UPDATE](#) trigger could insert incorrect values. (Bug #18437)
- **MySQL Cluster:** The [DATA_LENGTH](#) and [AVG_ROW_LENGTH](#) columns of the [INFORMATION_SCHEMA.TABLES](#) table did not report the size of variable-width column values correctly.

See [The INFORMATION_SCHEMA TABLES Table](#), for more information. (Bug #18413)
- **MySQL Cluster:** [SELECT ... FOR UPDATE](#) failed to lock the selected rows. (Bug #18184)
- **MySQL Cluster:** (Disk Data): Deletes from Disk Data tables used a nonoptimal scan to find the rows to be deleted, resulting in poor performance. The fix causes disk order rather than memory order to be used, and can improve performance of Disk Data deletes by up to ~300% in some cases. (Bug #17929)
- **MySQL Cluster:** [perror](#) did not properly report [NDB](#) error codes. (Bug #16561)
- **MySQL Cluster:** A problem with takeover during a system restart caused ordered indexes to be rebuilt incorrectly. This also adversely affected MySQL Cluster Replication. (Bug #15303)
- **MySQL Cluster:** A cluster data node could crash when an ordered index became full before the table containing the index was full. (Bug #14935)
- **MySQL Cluster:** The management client [ALL STATUS](#) command could sometimes report the status of some data nodes incorrectly. (Bug #13985)
- **MySQL Cluster:** New `mysqld` processes were permitted to connect without a restart of the cluster, causing the cluster to crash. (Bug #13266)
- **MySQL Cluster:** Cluster system status variables were not updated properly. (Bug #11459)
- **MySQL Cluster:** (NDBAPI): Update operations on blobs were not checked for illegal operations.

**Note**

Read locks with blob update operations are now upgraded from read committed to read shared.

- **MySQL Cluster:** The loss of one or more data nodes could sometimes cause `ndb_mgmd` to use a high amount of CPU (15 percent or more, as opposed to 1 to 2 percent normally).
- **Partitioning:** Old partition and subpartition files were not always removed following [ALTER TABLE ... REORGANIZE PARTITION](#) statements. (Bug #20770)
- **Replication; Cluster Replication:** In some cases, a large number of MySQL servers sending requests to the cluster simultaneously could cause the cluster to crash. This could also be triggered by many [NDB](#) API clients making simultaneous event subscriptions or unsubscriptions. (Bug #20683)

- **Replication; Cluster Replication:** Data definition and data manipulation statements on different tables were not serialized correctly in the binary log. For example, there was no guarantee that a `CREATE TABLE` statement and an update on a different table would occur in the same order in the binary log as they did on the cluster being replicated. (Bug #18947)
- **Replication:** `BIT` columns were not replicated properly under row-based replication. (Bug #22550)
- **Replication:** For row-based replication, log rotation could occur at an improper time. (Bug #21474)
- **Replication:** In mixed-format binary logging mode, stored functions, triggers, and views that use functions in their body that require row-based logging did not replicate reliably because the logging did not switch from statement-based to row-based format. For example, `INSERT INTO t SELECT FROM v`, where `v` is a view that selects `UUID()` could cause problems. This limitation has been removed. (Bug #20930)
- **Replication:** A race condition during slave server shutdown caused an assert failure. (Bug #20850)
- **Replication:** With mixed-format binary logging, `INSERT DELAYED` statements were logged using statement-based logging, and they did not replicate properly for statements that used values such as `UUID()`, `RAND()`, or user-defined variables that require row-based logging. To correct this, the `DELAYED` handler thread now switches to row-based logging if the logging format is mixed. (Bug #20633, Bug #20649)
- **Replication:** With the `auto_increment_increment` system variable set larger than 1, if the next generated `AUTO_INCREMENT` value would be larger than the column's maximum value, the value would be clipped down to that maximum value and inserted, even if the resulting value would not be in the generated sequence. This could cause problems for master-master replication. Now the server clips the value down to the previous value in the sequence, which correctly produces a duplicate-key error if that value already exists in the column. (Bug #20524)
- **Replication:** In mixed binary logging mode, a temporary switch from statement-based logging to row-based logging occurs when storing a row that uses a function such as `UUID()` into a temporary table. However, temporary table changes are not written to the binary log under row-based logging, so the row does not exist on the slave. A subsequent select from the temporary table to a nontemporary table using statement-based logging works correctly on the master, but not on the slave where the row does not exist. Replication no longer switches back from row-based logging to statement-based logging until there are no temporary tables for the session. (Bug #20499)
- **Replication:** `CREATE PROCEDURE`, `CREATE FUNCTION`, `CREATE TRIGGER`, and `CREATE VIEW` statements containing multi-line comments (`/* ... */`) could not be replicated. (Bug #20438)
- **Replication:** A stored procedure that used `LAST_INSERT_ID()` did not replicate properly using statement-based binary logging. (Bug #20339)
- **Replication:** When using row based replication, a `CREATE TABLE ... SELECT` statement was replicated, even if the table creation failed on the master (for example, due to a duplicate key failure). (Bug #20265)
- **Replication:** If a table on a slave server had a higher `AUTO_INCREMENT` counter than the corresponding master table (even though all rows of the two tables were identical), in some cases `REPLACE` or `INSERT ... ON DUPLICATE KEY UPDATE` would not replicate properly using statement-based logging. (Different values would be inserted on the master and slave.) (Bug #20188)
- **Replication:** Shutting down a slave in a replication scenario where temporary tables are in use would cause the slave to produce a core dump. (Bug #19881)
- **Replication:** The effect of a stored function or trigger that caused `AUTO_INCREMENT` values to be generated for multiple tables was not logged properly if statement-based logging was used. Only the

first table's value was logged, causing replication to fail. Under mixed logging format, this is dealt with by switching to row-based logging for the function or trigger. For statement-based logging, this remains a problem. (Bug #19630)

- **Replication:** For row-based replication, the `BINLOG` statement did not lock tables properly, causing a crash for some table types. (Bug #19459)
- **Replication:** Column names supplied for a view created on a master server could be lost on a slave server. (Bug #19419)
- **Replication:** The dropping of a temporary table whose name contained a backtick (``) character was not correctly written to the binary log, which also caused it not to be replicated correctly. (Bug #19188)
- **Replication:** With row-based replication, replicating a statement to a slave where the table had additional columns relative to the master table did not work. (Bug #19069)
- **Replication:** Valgrind revealed an issue with `mysqld` that was corrected: memory corruption in replication slaves when switching databases. (Bug #19022)
- **Replication:** A redundant table map event could be generated in the binary log when there were no actual changes to a table being replicated. In addition, a slave failed to stop when attempting to replicate a table that did not exist on the slave. (Bug #18948)
- **Replication:** Row-based replication failed when the query cache was enabled on the slave. (Bug #17620)
- **Replication:** Compilation on Windows failed if row based replication was disabled using `--without-row-based-replication`. (Bug #16837)
- **Replication:** An invalid `GRANT` statement for which `Ok` was returned on a replication master caused an error on the slave and replication to fail. (Bug #6774)
- **Disk Data:** On some platforms, `ndbd` compiled with `gcc 4` would crash when attempting to run `CREATE LOGFILE GROUP`. (Bug #21981)
- **Disk Data:** Trying to create a Disk Data table using a nonexistent tablespace or to drop a nonexistent data file from a tablespace produced an uninformative error message. (Bug #21751)
- **Disk Data:** Errors could occur when dropping a data file during a node local checkpoint. (Bug #21710)
- **Disk Data:** Creating a tablespace and log file group, then attempting to restart the cluster without using the `--initial` option and without having created any Disk Data tables could cause a forced shutdown of the cluster and raise a configuration error. (Bug #21172)
- **Disk Data:** `mysqldump` did not back up tablespace or log file group information for Disk Data tables correctly.

Specifically, `UNDO_BUFFER_SIZE` and `INITIAL_SIZE` values were misreported. This meant that trying to restore from such a backup would produce error 1296: `Got error 1504 'Out of logbuffer memory' from NDB`. (Bug #20809)

- **Disk Data:** Running a large number of scans on Disk Data could cause subsequent scans to perform poorly. (Bug #20334)
- **Disk Data:** `INFORMATION_SCHEMA.FILES` records for UNDO files showed incorrect values in the `EXTENT_SIZE`, `FREE_EXTENTS`, and `TOTAL_EXTENTS` columns. (Bug #20073)
- **Disk Data:** A data file created for one tablespace could be dropped using `ALTER TABLESPACE ... DROP DATAFILE` using a different tablespace. (Bug #20053)

- **Disk Data:** Trying to create Disk Data tables when running the cluster in diskless mode caused cluster data nodes to crash.

**Note**

Disk Data tables are now disabled when running in diskless mode.

(Bug #20008)

- **Disk Data:** An issue with disk allocation could sometimes cause a forced shutdown of the cluster when running a mix of memory and Disk Data tables. (Bug #18780)
- **Disk Data:** The failure of a `CREATE TABLESPACE` or `CREATE LOGFILE GROUP` statement did not revert all changes made prior to the point of failure. (Bug #16341)
- **Cluster Replication:** One or more of the `mysqld` processes could fail when subjecting a Cluster replication setup with multiple `mysqld` processes on both the master and slave clusters to high loads. (Bug #19768)
- **Cluster API:** The `storage/ndb` directory was missing from the server binary distribution, making it impossible to compile NDB API and MGM API applications. This directory can be found as `/usr/include/storage/ndb` after installing that distribution. (Bug #21955)
- **Cluster API:** Invoking the MGM API function `ndb_mgm_listen_event()` caused a memory leak. (Bug #21671)
- **Cluster API:** The inclusion of `my_config.h` in `NdbApi.h` required anyone wishing to write NDB API applications against MySQL 5.1 to have a complete copy of the 5.1 sources. (Bug #21253)
- **Cluster API:** The MGM API function `ndb_logevent_get_fd()` was not implemented. (Bug #21129)
- **Cluster API:** The `NdbOperation::getBlobHandle()` method, when called with the name of a nonexistent column, caused a segmentation fault. (Bug #21036)
- **Cluster API:** `NdbScanOperation::readTuples()` and `NdbIndexScanOperation::readTuples()` ignored the `batch` parameter. (Bug #20252)
- `ALTER EVENT` statements including only a `COMMENT` clause failed with a syntax error on two platforms: Linux for S/390, and OS X 10.4 for 64-bit PowerPC. (Bug #23423)
- When `event_scheduler` was set to `DISABLED`, its value was not displayed correctly by `SHOW VARIABLES` or `SELECT @@global.event_scheduler`. (Bug #22662)
- `ALTER EVENT` in the body of a stored procedure led to a crash when the procedure was called. This affected only those `ALTER EVENT` statements which changed the interval of the event. (Bug #22397)
- The optimizer could make an incorrect index choice for indexes with a skewed key distribution. (Bug #22393)
- Deleting entries from a large `MyISAM` index could cause index corruption when it needed to shrink. Deletes from an index can happen when a record is deleted, when a key changes and must be moved, and when a key must be un-inserted because of a duplicate key. This can also happen in `REPAIR TABLE` when a duplicate key is found and in `myisamchk` when sorting the records by an index. (Bug #22384)
- Instance Manager had a race condition involving `mysqld` PID file removal. (Bug #22379)
- yaSSL had a conflicting definition for `socklen_t` on hurd-i386 systems. (Bug #22326)

- Conversion of values inserted into a `BIT` column could affect adjacent columns. (Bug #22271)
- 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)
- `mysql_com.h` unnecessarily referred to the `ulong` type. (Bug #22227)
- The source distribution would not build on Windows due to a spurious dependency on `ib_config.h`. (Bug #22224)
- Execution of a prepared statement that uses an `IN` subquery with aggregate functions in the `HAVING` clause could cause a server crash. (Bug #22085)
- The `CSV` storage engine failed to detect some table corruption. (Bug #22080)
- Using `GROUP_CONCAT()` on the result of a subquery in the `FROM` clause that itself used `GROUP_CONCAT()` could cause a server crash. (Bug #22015)
- Running `SHOW MASTER LOGS` at the same time as binary log files were being switched would cause `mysqld` to hang. (Bug #21965)
- `libmysqlclient` defined a symbol `BN_bin2bn` which belongs to OpenSSL. This could break applications that also linked against OpenSSL's `libcrypto` library. The fix required correcting an error in a build script that was failing to add rename macros for some functions. (Bug #21930)
- `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)
- A misleading error message was displayed when attempting to define a unique key that was not valid for a partitioned table. (Bug #21862)
- A query that used `GROUP BY` and an `ALL` or `ANY` quantified subquery in a `HAVING` clause could trigger an assertion failure. (Bug #21853)
- An `InnoDB` mutex was not acquired and released under the same condition, leading to deadlock in some rare situations involving XA transactions. (Bug #21833)
- A `NUL` byte within a prepared statement string caused the rest of the string not to be written to the query log, permitting logging to be bypassed. (Bug #21813)
- `COUNT(*)` queries with `ORDER BY` and `LIMIT` could return the wrong result.

**Note**

This problem was introduced by the fix for Bug #9676, which limited the rows stored in a temporary table to the `LIMIT` clause. This optimization is not applicable to nongroup queries with aggregate functions. The current fix disables the optimization in such cases.

(Bug #21787)

References: This issue is a regression of: Bug #9676.

- Using `DROP TABLE` with concurrent queries causes `mysqld` to crash. (Bug #21784)
- `INSERT ... SELECT` sometimes generated a spurious `Column count doesn't match value count` error. (Bug #21774)
- `UPGRADE` was treated as a reserved word, although it is not. (Bug #21772)

- A function result in a comparison was replaced with a constant by the optimizer under some circumstances when this optimization was invalid. (Bug #21698)
- Selecting from `INFORMATION_SCHEMA.FILES` could crash the server. (Bug #21676)
- Errors could be generated during the execution of certain prepared statements that ran queries on partitioned tables. (Bug #21658)
- The presence of a subquery in the `ON` clause of a join in a view definition prevented the `MERGE` algorithm from being used for the view in cases where it should be permitted. (Bug #21646)
- When records are merged from the insert buffer and the page needs to be reorganized, `InnoDB` used incorrect column length information when interpreting the records of the page. This caused a server crash due to apparent corruption of secondary indexes in `ROW_FORMAT=COMPACT` that contain prefix indexes of fixed-length columns. Data files should not be corrupted, but the crash was likely to repeat every time the server was restarted. (Bug #21638)
- For character sets having a `mbmaxlen` value of 2, any `ALTER TABLE` statement changed `TEXT` columns to `MEDIUMTEXT`. (Bug #21620)
- `mysql` displayed an empty string for `NULL` values. (Bug #21618)
- Selecting from a `MERGE` table could result in a server crash if the underlying tables had fewer indexes than the `MERGE` table itself. (Bug #21617, Bug #22937)
- A loaded storage engine plugin did not load after a server restart. (Bug #21610)
- For `INSERT ... ON DUPLICATE KEY UPDATE`, use of `VALUES(col_name)` within the `UPDATE` clause sometimes was handled incorrectly. (Bug #21555)
- Subqueries with aggregate functions but no `FROM` clause could return incorrect results. (Bug #21540)
- `mysqldump` incorrectly tried to use `LOCK TABLES` for tables in the `INFORMATION_SCHEMA` database. (Bug #21527)
- The server could crash for the second execution of a function containing a `SELECT` statement that uses an aggregating `IN` subquery. (Bug #21493)
- Memory overruns could occur for certain kinds of subqueries. (Bug #21477)
- A `DATE` can be represented as an integer (such as `20060101`) or as a string (such as `'2006.01.01'`). When a `DATE` (or `TIME`) column is compared in one `SELECT` against both representations, constant propagation by the optimizer led to comparison of `DATE` as a string against `DATE` as an integer. This could result in integer comparisons such as `2006` against `20060101`, erroneously producing a false result. (Bug #21475)
- `myisam_ftdump` produced bad counts for common words. (Bug #21459)
- Adding `ORDER BY` to a `SELECT DISTINCT(expr)` query could produce incorrect results. (Bug #21456)
- The URL into the online manual that is printed in the stack trace message by the server was out of date. (Bug #21449)
- Database and table names have a maximum length of 64 characters (even if they contain multibyte characters), but were truncated to 64 bytes.

**Note**

An additional fix was made in MySQL 5.1.18.

(Bug #21432)

- With `max_sp_recursion` set to 0, a stored procedure that executed a `SHOW CREATE PROCEDURE` statement for itself triggered a recursion limit exceeded error, though the statement involves no recursion. (Bug #21416)
- After `FLUSH TABLES WITH READ LOCK` followed by `UNLOCK TABLES`, attempts to drop or alter a stored routine failed with an error that the routine did not exist, and attempts to execute the routine failed with a lock conflict error. (Bug #21414)
- On 64-bit Windows, a missing table generated error 1017, not the correct value of 1146. (Bug #21396)
- Table aliases in multiple-table `DELETE` statements sometimes were not resolved. (Bug #21392)
- The optimizer sometimes produced an incorrect row-count estimate after elimination of `const` tables. This resulted in choosing extremely inefficient execution plans in some cases when distribution of data in joins were skewed. (Bug #21390)
- For multiple-table `UPDATE` statements, storage engines were not notified of duplicate-key errors. (Bug #21381)
- Using relative paths for `DATA DIRECTORY` or `INDEX DIRECTORY` with a partitioned table generated a warning rather than an error, and caused “junk” files to be created in the server's data directory. (Bug #21350)
- Using `EXPLAIN PARTITIONS` with a query on a table whose partitioning expression was based on the value of a `DATE` column could sometimes cause the server to crash. (Bug #21339)
- The feature of being able to recover a temporary table named `#sql_id` in InnoDB by creating a table named `rsql_id_recover_innodb_tmp_table` was broken by the introduction of the new identifier encoding in MySQL 5.1.6 (Bug #21313)
- It was possible for a stored routine with a non-`latin1` name to cause a stack overrun. (Bug #21311)
- A query result could be sorted improperly when using `ORDER BY` for the second table in a join. (Bug #21302)
- Query results could be incorrect if the `WHERE` clause contained `t.key_part NOT IN (val_list)`, where `val_list` is a list of more than 1000 constants. (Bug #21282)
- Queries that used the `index_merge` and `sort_union` methods to access an InnoDB table could produce inaccurate results. This issue was introduced in MySQL 5.1.10 when a new handler and bitmap interface was implemented. (Bug #21277)
- For user-defined functions created with `CREATE FUNCTION`, the `DEFINER` clause is not legal, but no error was generated. (Bug #21269)
- The `SELECT` privilege was required for an insert on a view, instead of the `INSERT` privilege. (Bug #21261)

References: This issue is a regression of: Bug #20989.

- `mysql_config --libmysqld-libs` did not produce any SSL options necessary for linking `libmysqld` with SSL support enabled. (Bug #21239)
- Subqueries on `INFORMATION_SCHEMA` tables could erroneously return an empty result. (Bug #21231)
- `mysql_upgrade` created temporary files in a possibly insecure way. (Bug #21224)

- When `DROP DATABASE` or `SHOW OPEN TABLES` was issued while concurrently in another connection issuing `DROP TABLE`, `RENAME TABLE`, `CREATE TABLE LIKE` or any other statement that required a name lock, the server crashed. (Bug #21216, Bug #19403)
- The `--master-data` option for `mysqldump` requires certain privileges, but `mysqldump` generated a truncated dump file without producing an appropriate error message or exit status if the invoking user did not have those privileges. (Bug #21215)
- Using `ALTER TABLE ... REORGANIZE PARTITIONS` to reduce the number of subpartitions to 1 caused the server to crash. (Bug #21210)
- In the package of pre-built time zone tables that is available for download at <http://dev.mysql.com/downloads/timezones.html>, the tables now explicitly use the `utf8` character set so that they work the same way regardless of the system character set value. (Bug #21208)
- 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.

Use of more than 1024 simultaneous cursors server wide also could lead to memory corruption. This applies to both stored procedure cursors and C API cursors. (Bug #21206)
- When run with the `--use-threads` option, `mysqlimport` returned a random exit code. (Bug #21188)
- A subquery that uses an index for both the `WHERE` and `ORDER BY` clauses produced an empty result. (Bug #21180)
- Running `SHOW TABLE STATUS` on any `InnoDB` table having at least one record could crash the server. Note that this was not due to any issue in the `InnoDB` storage engine, but rather with `AUTO_INCREMENT` handling in the partitioning code—however, the table did not have to have an `AUTO_INCREMENT` column for the bug to manifest. (Bug #21173)
- Some prepared statements caused a server crash when executed a second time. (Bug #21166)
- The optimizer assumed that if `(a=x AND b=x)` is true, `(a=x AND b=x) AND a=b` is also true. But that is not always so if `a` and `b` have different data types. (Bug #21159)
- Some `ALTER TABLE` statements affecting a table's subpartitioning could hang. (Bug #21143)
- Certain malformed `INSERT` statements could crash the `mysql` client. (Bug #21142)
- `SHOW INNODB STATUS` contained some duplicate output. (Bug #21113)
- `InnoDB` was slow with more than 100,000 `.idb` files. (Bug #21112)
- Creating a `TEMPORARY` table with the same name as an existing table that was locked by another client could result in a lock conflict for `DROP TEMPORARY TABLE` because the server unnecessarily tried to acquire a name lock. (Bug #21096)
- Performing an `INSERT` on a view that was defined using a `SELECT` that specified a collation and a column alias caused the server to crash. (Bug #21086)
- Incorrect results could be obtained from re-execution of a parametrized prepared statement or a stored routine with a `SELECT` that uses `LEFT JOIN` with a second table having only one row. (Bug #21081)
- `ALTER VIEW` did not retain existing values of attributes that had been originally specified but were not changed in the `ALTER VIEW` statement. (Bug #21080)
- The `myisam_stats_method` variable was mishandled when set from an option file or on the command line. (Bug #21054)

- With `query_cache_type` set to 0, `RESET QUERY CACHE` was very slow and other threads were blocked during the operation. Now a cache reset is faster and nonblocking. (Bug #21051)
- `mysql` crashed for very long arguments to the `connect` command. (Bug #21042)
- When creating a table using `CREATE...SELECT` and a stored procedure, there would be a mismatch between the binary log and transaction cache which would cause a server crash. (Bug #21039)
- 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 #21019)
- `mysqldump` sometimes did not select the correct database before trying to dump views from it, resulting in an empty result set that caused `mysqldump` to die with a segmentation fault. (Bug #21014)
- Performance during an import on a table with a trigger that called a stored procedure was severely degraded. (Bug #21013)
- `mysql_upgrade` produced a malformed `upgrade_defaults` file by overwriting the `[client]` group header with a `password` option. This prevented `mysqlcheck` from running successfully when invoked by `mysql_upgrade`. (Bug #21011)
- A query of the form shown here caused the server to crash:

```
SELECT * FROM t1 NATURAL JOIN (  
    t2 JOIN (  
        t3 NATURAL JOIN t4,  
        t5 NATURAL JOIN t6  
    )  
    ON (t3.id3 = t2.id3 AND t5.id5 = t2.id5)  
);
```

(Bug #21007)

- A `SELECT` that used a subquery in the `FROM` clause that did not select from a table failed when the subquery was used in a join. (Bug #21002)
- `REPLACE ... SELECT` for a view required the `INSERT` privilege for tables other than the table being modified. (Bug #20989)
- `STR_TO_DATE()` sometimes would return `NULL` if the `%D` format specifier was not the last specifier in the format string. (Bug #20987)
- 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)
- 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)
- Closing of temporary tables failed if binary logging was not enabled. (Bug #20919)
- Use of zero-length variable names caused a server crash. (Bug #20908)
- Building `mysql` on Windows with `CMake` 2.4 failed to create `libmysqld` correctly. (Bug #20907)
- Creating a partitioned table that used the `InnoDB` storage engine and then restarting `mysqld` with `--skip-innodb` caused MySQL to crash. (Bug #20871)

- For certain queries, the server incorrectly resolved a reference to an aggregate function and crashed. (Bug #20868)
- If the binary logging format was changed between the times when a locked table was modified and when it was unlocked, the binary log contents were incorrect. (Bug #20863)
- It was possible to provide the `ExtractValue()` function with input containing “tags” that were not valid XML; for example, it was possible to use tag names beginning with a digit, which are not permitted by the W3C's XML 1.0 specification. Such cases caused the function to return “junk” output rather than an error message signalling the user as to the true nature of the problem. (Bug #20854)
- **InnoDB** (Partitioning): Updating an **InnoDB** table using **HASH** partitioning with a composite primary key would cause the server to hang. (Bug #20852)
- `mysqldump` did not add version-specific comments around **WITH PARSE** and **TABLESPACE ... STORAGE DISK** clauses for **CREATE TABLE** statements, causing dump files from servers where these features were in use to fail when loaded into older servers. (Bug #20841)
- For multiple **INSERT DELAYED** statements executed in a batch by the delayed-insert handler thread, not all rows were written to the binary log. (Bug #20821)
- The `ExtractValue()` function did not accept XML tag names containing a period (.) character. (Bug #20795)
- Using aggregate functions in subqueries yielded incorrect results under certain circumstances due to incorrect application of `MIN()/MAX()` optimization. (Bug #20792)
- On Windows, inserting into a **MERGE** table after renaming an underlying **MyISAM** table caused a server crash. (Bug #20789)
- Within stored routines, some error messages were printed incorrectly. A nonnull-terminated string was passed to a message-printing routine that expected a null-terminated string. (Bug #20778)
- Merging multiple partitions having subpartitions into a single partition with subpartitions, or splitting a single partition having subpartitions into multiple partitions with subpartitions, could sometimes crash the server. These issues were associated with a failure reported in the `partition_range` test. (Bug #20767, Bug #20893, Bug #20766, Bug #21357)
- Searches against a **ZEROFILL** column of a partitioned table could fail when the **ZEROFILL** column was part of the table's partitioning key. (Bug #20733)
- 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)
- An **UPDATE** that referred to a key column in the **WHERE** clause and activated a trigger that modified the column resulted in a loop. (Bug #20670)
- Issuing a **SHOW CREATE FUNCTION** or **SHOW CREATE PROCEDURE** statement without sufficient privileges could crash the `mysql` client. (Bug #20664)
- **INSERT DELAYED** did not honor `SET INSERT_ID` or the `auto_increment_*` system variables. (Bug #20627, Bug #20830)
- A buffer overwrite error in Instance Manager caused a crash. (Bug #20622)
- Loading a plugin caused any an existing plugin with the same name to be lost. (Bug #20615)
- A query selecting records from a single partition of a partitioned table and using **ORDER BY ic DESC** (where *ic* represents an indexed column) could cause errors or crash the server. (Bug #20583)

- Valgrind revealed several issues with `mysqld` that were corrected: A dangling stack pointer being overwritten; possible uninitialized data in a string comparison; `syscall()` write parameter pointing to an uninitialized byte. (Bug #20579, Bug #20769, Bug #20783, Bug #20791)
- If the `auto_increment_offset` setting causes MySQL to generate a value larger than the column's maximum possible value, the `INSERT` statement is accepted in strict SQL mode, whereas but should fail with an error. (Bug #20573)
- In a view defined with `SQL SECURITY DEFINER`, the `CURRENT_USER()` function returned the invoker, not the definer. (Bug #20570)
- 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)
- Scheduled events that invoked stored procedures executing DDL operations on partitioned tables could crash the server. (Bug #20548)
- Users who had the `SHOW VIEW` privilege for a view and privileges on one of the view's base tables could not see records in `INFORMATION_SCHEMA` tables relating to the base table. (Bug #20543)
- The `fill_help_tables.sql` file did not load properly if the `ANSI_QUOTES` SQL mode was enabled. (Bug #20542)
- 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)
- If a partitioned `InnoDB` table contained an `AUTO_INCREMENT` column, a `SHOW` statement could cause an assertion failure with more than one connection. (Bug #20493)
- Using `EXPLAIN PARTITIONS` with a `UNION` query could crash the server. This could occur whether or not the query actually used any partitioned tables. (Bug #20484)
- Creation of a view as a join of views or tables could fail if the views or tables are in different databases. (Bug #20482)
- `SELECT` statements using `GROUP BY` against a view could have missing columns in the output when there was a trigger defined on one of the base tables for the view. (Bug #20466)
- For connections that required a `SUBJECT` value, a check was performed to verify that the value was correct, but the connection was not refused if not. (Bug #20411)
- `mysql_upgrade` was missing from binary MySQL distributions. (Bug #20403, Bug #18516, Bug #20556)
- Some user-level errors were being written to the server's error log, which is for server errors. (Bug #20402)
- Using `ALTER TABLE ... ENGINE = x`, where `x` was not a storage engine supported by the server, would cause `mysqld` to crash. (Bug #20397)
- User names have a maximum length of 16 characters (even if they contain multibyte characters), but were being truncated to 16 bytes. (Bug #20393)
- Some queries using `ORDER BY ... DESC` on subpartitioned tables could crash the server. (Bug #20389)

- `mysqslap` did not enable the `CLIENT_MULTI_RESULTS` flag when connecting, which is necessary for executing stored procedures. (Bug #20365)
- 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)
- When a statement used a stored function that inserted into an `AUTO_INCREMENT` column, the generated `AUTO_INCREMENT` value was not written into the binary log, so a different value could in some cases be inserted on the slave. (Bug #20341)
- Partitions were represented internally as the wrong data type, which led in some cases to failures of queries such as `SELECT COUNT(*) FROM INFORMATION_SCHEMA.PARTITIONS WHERE PARTITION_NAME = 'partition_name'`. (Bug #20340)
- `PROCEDURE ANALYSE()` returned incorrect values of `M FLOAT(M, D)` and `DOUBLE(M, D)`. (Bug #20305)
- Defining a table partitioned by `LIST` with a single `PARTITION ... VALUES IN (NULL)` clause could lead to server crashes, particularly with queries having `WHERE` conditions comparing the partitioning key with a constant. (Bug #20268, Bug #19801)
- Partition pruning could cause incorrect results from queries, such missing rows, when the partitioning expression relied on a `BIGINT UNSIGNED` column. (Bug #20257)
- For a `MyISAM` table locked with `LOCK TABLES ...WRITE`, queries optimized using the `index_merge` method did not show rows inserted with the lock in place. (Bug #20256)
- `mysqldump` produced a malformed dump file when dumping multiple databases that contained views. (Bug #20221)
- Running `InnoDB` with many concurrent threads could cause memory corruption and a seg fault due to a bug introduced in MySQL 5.1.11. (Bug #20213)
- `SUBSTRING()` results sometimes were stored improperly into a temporary table when multibyte character sets were used. (Bug #20204)
- The thread for `INSERT DELAYED` rows was maintaining a separate `AUTO_INCREMENT` counter, resulting in incorrect values being assigned if `DELAYED` and non-`DELAYED` inserts were mixed. (Bug #20195)
- The `--default-storage-engine` server option did not work. (Bug #20168)
- For a table having `LINEAR HASH` subpartitions, the `LINEAR` keyword did not appear in the `SUBPARTITION_METHOD` column of the `INFORMATION_SCHEMA.PARTITIONS` table. (Bug #20161)
- For a `DATE` parameter sent using 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)
- `pererror` crashed on Solaris due to `NULL` return value of `strerror()` system call. (Bug #20145)
- `FLUSH TABLES` followed by a `LOCK TABLES` statement to lock a log table and a nonlog table caused an infinite loop and high CPU use. Now `FLUSH TABLES` ignores log tables. To flush the log tables, use `FLUSH LOGS` instead. (Bug #20139)
- On Linux, `libmysqlclient` when compiled with yaSSL using the `icc` compiler had a spurious dependency on C++ libraries. (Bug #20119)

- For an `ENUM` column that used the `ucs2` character set, using `ALTER TABLE` to modify the column definition caused the default value to be lost. (Bug #20108)
- For `mysql`, escaping with backslash sometimes did not work. (Bug #20103)
- Queries on tables that were partitioned by `KEY` and had a `VARCHAR` column as the partitioning key produced an empty result set. (Bug #20086)
- A number of dependency issues in the RPM `bench` and `test` packages caused installation of these packages to fail. (Bug #20078)
- Use of `MIN()` or `MAX()` with `GROUP BY` on a `ucs2` column could cause a server crash. (Bug #20076)
- `mysqld --flush` failed to flush `MyISAM` table changes to disk following an `UPDATE` statement for which no updated column had an index. (Bug #20060)
- In MySQL 5.1.11, the `--with-openssl` and `--with-yassl` options were replaced by `--with-ssl`. But no message was issued if the old options were given. Now `configure` produces a message indicating that the new option should be used and exits. (Bug #20002)
- When a statement is executed that does not generate any rows, an extra table map event and associated binrows event would be generated and written to the binary log. (Bug #19995)
- Join conditions using index prefixes on `utf8` columns of `InnoDB` tables incorrectly ignored rows where the length of the actual value was greater than the length of the index prefix. (Bug #19960)
- `AUTHORS` and `CONTRIBUTORS` were not treated as reserved words. (Bug #19939)
- The `query` command for `mysqltest` did not work. (Bug #19890)
- Identifiers with embedded escape characters were not handled correctly by some `SHOW` statements due to some old code that was doing some extra unescaping. (Bug #19874)
- When executing a `SELECT` with `ORDER BY` on a view that is constructed from a `SELECT` statement containing a stored function, the stored function was evaluated too many times. (Bug #19862)
- Using `SELECT` on a corrupt `MyISAM` table using the dynamic record format could cause a server crash. (Bug #19835)
- Using cursors with `READ COMMITTED` isolation level could cause `InnoDB` to crash. (Bug #19834)
- `CREATE DATABASE`, `RENAME DATABASE`, and `DROP DATABASE` could deadlock in cases where there was a global read lock. (Bug #19815)
- The yaSSL library bundled with `libmysqlclient` had some conflicts with OpenSSL. Now macros are used to rename the conflicting symbols to have a prefix of `ya`. (Bug #19810)
- The `WITH CHECK OPTION` was not enforced when a `REPLACE` statement was executed against a view. (Bug #19789)
- Multiple-table updates with `FEDERATED` tables could cause a server crash. (Bug #19773)
- 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)
- `make install` tried to build files that should already have been built by `make all`, causing a failure if installation was performed using a different account than the one used for the initial build. (Bug #19738)
- `InnoDB` unlocked its data directory before committing a transaction, potentially resulting in nonrecoverable tables if a server crash occurred before the commit. (Bug #19727)

- An issue with yaSSL prevented Connector/J clients from connecting to the server using a certificate. (Bug #19705)
- For a `MyISAM` table with a `FULLTEXT` index, compression with `myisampack` or a check with `myisamchk` after compression resulted in table corruption. (Bug #19702)
- The `ENGINE` clause was displayed in the output of `SHOW CREATE TABLE` for partitioned tables when the SQL mode included `no_table_options`. (Bug #19695)
- A cast problem caused incorrect results for prepared statements that returned float values when MySQL was compiled with `gcc` 4.0. (Bug #19694)
- `EXPLAIN PARTITIONS` would produce illegible output in the `partitions` column if the length of text to be displayed in that column was too long. This could occur when very many partitions were defined for the table, partitions were given very long names, or due to a combination of the two. (Bug #19684)
- The `mysql_list_fields()` C API function returned the incorrect table name for views. (Bug #19671)
- If a query had a condition of the form `tableX.key = tableY.key`, which participated in equality propagation and also was used for `ref` access, then early `ref`-access `NULL` filtering was not performed for the condition. This could make query execution slower. (Bug #19649)
- Re-execution of a prepared multiple-table `DELETE` statement that involves a trigger or stored function can result in a server crash. (Bug #19634)
- File size specifications for `InnoDB` data files were case sensitive. (Bug #19609)
- `CHECK TABLE` on a `MyISAM` table briefly cleared its `AUTO_INCREMENT` value, while holding only a read lock. Concurrent inserts to that table could use the wrong `AUTO_INCREMENT` value. `CHECK TABLE` no longer modifies the `AUTO_INCREMENT` value. (Bug #19604)
- Some yaSSL public function names conflicted with those from OpenSSL, causing conflicts for applications that linked against both OpenSSL and a version of `libmysqlclient` that was built with yaSSL support. The yaSSL public functions now are renamed to avoid this conflict. (Bug #19575)
- In the `INFORMATION_SCHEMA.FILES` table, the `INITIAL_SIZE`, `MAXIMUM_SIZE`, and `AUTOEXTEND_SIZE` columns incorrectly were being stored as `VARCHAR` rather than `BIGINT`. (Bug #19544)
- `InnoDB` failed to increment the `handler_read_prev` counter. (Bug #19542)
- Portions of statements related to partitioning were not surrounded by version-specific comments by `mysqldump`, breaking backward compatibility for dump files. (Bug #19488)
- Repeated `DROP TABLE` statements in a stored procedure could sometimes cause the server to crash. (Bug #19399)
- Renaming a database to itself caused a server crash. (Bug #19392)
- Race conditions on certain platforms could cause the Instance Manager to fail to initialize. (Bug #19391)
- When not running in strict mode, the server failed to convert the invalid years portion of a `DATE` or `DATETIME` value to `'0000'` when inserting it into a table.

**Note**

This fix was reverted in MySQL 5.1.18.

(Bug #19370)

References: See also: Bug #25301.

- Use of the `--no-pager` option caused `mysql` to crash. (Bug #19363)
- Multiple calls to a stored procedure that altered a partitioned `MyISAM` table would cause the server to crash. (Bug #19309)
- `ALTER TABLE ... COALESCE PARTITION` did not delete the files associated with the partitions that were removed. (Bug #19305)
- Adding an index to a partitioned table that had been created using `AUTO_INCREMENT = value` caused the `AUTO_INCREMENT` value to be reset. (Bug #19281)
- Multiple-table `DELETE` statements containing a subquery that selected from one of the tables being modified caused a server crash. (Bug #19225)
- The final parenthesis of a `CREATE INDEX` statement occurring in a stored procedure was omitted from the binary log when the stored procedure was called. (Bug #19207)
- An `ALTER TABLE` operation that does not need to copy data, when executed on a table created prior to MySQL 4.0.25, could result in a server crash for subsequent accesses to the table. (Bug #19192)
- SSL connections using `yaSSL` on `OpenBSD` could fail. (Bug #19191)
- `ALTER TABLE ... REBUILD PARTITION` could cause the server to hang or crash. (Bug #19122)
- Using `ALTER TABLE` on a subpartitioned table caused the server to crash. (Bug #19067)
- Trying to execute a query having a `WHERE` clause using `int_col = "string_value" OR int_col IS NULL` on a partitioned table whose partitioning or subpartitioning function used the integer column `int_col` would crash the server. (Bug #19055)
- A `SELECT` with a subquery that was bound to the outer query over multiple columns returned different results when a constant was used instead of one of the dependant columns. (Bug #18925)
- It was possible using `ALTER EVENT ... RENAME ...` to move an event to a database on which the user did not have the `EVENT` privilege. (Bug #18897)
- When used in the `DO` clause of a `CREATE EVENT` statement, the statements `CREATE EVENT`, `CREATE FUNCTION`, and `CREATE PROCEDURE` caused the server to crash. (These statements are not permitted inside `CREATE EVENT`.) (Bug #18896, Bug #16409)
- `BIT` columns in a table could cause joins that use the table to fail. (Bug #18895)
- The build process incorrectly tried to overwrite `sql/lex_hash.h`. This caused the build to fail when using a shadow link tree pointing to original sources that were owned by another account. (Bug #18888)
- Setting `myisam_repair_threads` caused any repair operation on a `MyISAM` table to fail to update the cardinality of indexes, instead making them always equal to 1. (Bug #18874)
- 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)
- `InnoDB`: Quoted Unicode identifiers were not handled correctly. This included names of tables, columns, and foreign keys. (Bug #18800)

- Intermediate tables created during the execution of an `ALTER TABLE` statement were visible in the output of `SHOW TABLES`. (Bug #18775)
- `FEDERATED` tables raised invalid duplicate key errors when attempting on one server to insert rows having the same primary key values as rows that had been deleted from the linked table on the other server. (Bug #18764)
- Memory used by scheduled events was not freed when the events were dropped. (Bug #18683)
- The implementation for `UNCOMPRESS()` did not indicate that it could return `NULL`, causing the optimizer to do the wrong thing. (Bug #18539)
- Referring to a stored function qualified with the name of one database and tables in another database caused a “table doesn't exist” error. (Bug #18444)
- Identifiers could not contain bytes with a value of 255, though that should be permitted as of the identifier-encoding changes made in MySQL 5.1.6. (Bug #18396)
- Triggers on tables in the `mysql` database caused a server crash. Triggers for tables in this database are no longer permitted. (Bug #18361, Bug #18005)
- Incorrect type aggregation for `IN()` and `CASE` expressions could lead to an incorrect result. (Bug #18360)
- The length of the pattern string prefix for `LIKE` operations was calculated incorrectly for multibyte character sets. As a result, the scanned range was wider than necessary if the prefix contained any multibyte characters, and rows could be missing from the result set. (Bug #18359, Bug #16674)
- On Windows, corrected a crash stemming from differences in Visual C runtime library routines from POSIX behavior regarding invalid file descriptors. (Bug #18275)
- Linking the `pthread`s library to single-threaded MySQL libraries caused `dlopen()` to fail at runtime on HP-UX. (Bug #18267)
- The source distribution failed to compile when configured with the `--with-libwrap` option. (Bug #18246)
- On Windows, terminating `mysqld` with **Control+C** could result in a crash during shutdown. (Bug #18235)
- Selecting data from a `MEMORY` table with a `VARCHAR` column and a `HASH` index over it returned only the first row matched. (Bug #18233)
- The use of `MIN()` and `MAX()` on columns with an index prefix produced incorrect results in some queries. (Bug #18206)
- A `UNION` over more than 128 `SELECT` statements that use an aggregate function failed. (Bug #18175)
- The optimizer did not take advantage of indexes on columns used for the second or third arguments of `BETWEEN`. (Bug #18165)
- Performing `INSERT ... SELECT ... JOIN ... USING` without qualifying the column names caused `ERROR 1052 "column 'x' in field list is ambiguous"` even in cases where the column references were unambiguous. (Bug #18080)
- 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)

- Race conditions on certain platforms could cause the Instance Manager to try to restart the same instance multiple times. (Bug #18023)
- Changing the definition of a `DECIMAL` column with `ALTER TABLE` caused loss of column values. (Bug #18014)
- For table-format output, `mysql` did not always calculate columns widths correctly for columns containing multibyte characters in the column name or contents. (Bug #17939)
- The character set was not being properly initialized for `CAST()` with a type such as `CHAR(2) BINARY`, which resulted in incorrect results or a server crash. (Bug #17903)
- Checking a `MyISAM` table (using `CHECK TABLE`) having a spatial index and only one row would wrongly indicate that the table was corrupted. (Bug #17877)
- For a reference to a nonexistent index in `FORCE INDEX`, the error message referred to a column, not an index. (Bug #17873)
- A stored procedure that created and invoked a prepared statement was not executed when called in a `mysqld` init-file. (Bug #17843)
- It is possible to create `MERGE` tables into which data cannot be inserted (by not specifying a `UNION` clause. However, when an insert was attempted, the error message was confusing. Now an error occurs indicating that the table is read only. (Bug #17766)
- Attempting to insert a string of greater than 4096 bytes into a `FEDERATED` table resulted in the error `ERROR 1296 (HY000) at line 2: Got error 10000 'Error on remote system: 1054: Unknown column 'string-value' from FEDERATED`. This error was raised regardless of the type of column involved (`VARCHAR`, `TEXT`, and so on.) (Bug #17608)
- If a file name was specified for the `--log` or `--log-slow-queries` options but the server was logging to tables and not files, the server produced no error message. (Bug #17599)
- If the general log table reached a large enough file size (27GB), `SELECT COUNT(*)` on the table caused a server crash. (Bug #17589)
- Using the extended syntax for `TRIM()`—that is, `TRIM(... FROM ...)`—in a `SELECT` statement defining a view caused an invalid syntax error when selecting from the view. (Bug #17526)
- Use of the `--prompt` option or `prompt` command caused `mysql` to be unable to connect to the Instance Manager. (Bug #17485)
- `OPTIMIZE TABLE` and `REPAIR TABLE` yielded incorrect messages or warnings when used on partitioned tables. (Bug #17455)
- `mysqldump` would not dump views that had become invalid because a table named in the view definition had been dropped. Instead, it quit with an error message. Now you can specify the `--force` option to cause `mysqldump` to keep going and write an SQL comment containing the view definition to the dump output. (Bug #17371)
- `N'xxx'` and `_utf8'xxx'` were not treated as equivalent because `N'xxx'` failed to unescape backslashes (`\`) and doubled apostrophe/single quote characters (`'`). (Bug #17313)
- Following a failed attempt to add an index to an `ARCHIVE` table, it was no longer possible to drop the database in which the table had been created. (Bug #17310)
- Assignments of values to variables of type `TEXT` were handled incorrectly in stored routines. (Bug #17225)

- Views created from prepared statements inside of stored procedures were created with a definition that included both `SQL_CACHE` and `SQL_NO_CACHE`. (Bug #17203)
- `mysqldump` wrote an extra pair of `DROP DATABASE` and `CREATE DATABASE` statements if run with the `--add-drop-database` option and the database contained views. (Bug #17201)
- A `Table ... doesn't exist` error could occur for statements that called a function defined in another database. (Bug #17199)
- A prepared statement that altered partitioned table within a stored procedure failed with the error `Unknown prepared statement handler`. (Bug #17138)
- `myisam_ftdump` failed when trying to open a MyISAM index file that you did not have write permissions to access, even though the command would only be reading from the file. (Bug #17122)
- `ALTER TABLE` on a table created prior to 5.0.3 would cause table corruption if the `ALTER TABLE` did one of the following:
 - Change the default value of a column.
 - Change the table comment.
 - Change the table password.(Bug #17001)
- For statements that have a `DEFINER` clause such as `CREATE TRIGGER` or `CREATE VIEW`, long user names or host names could cause a buffer overflow. (Bug #16899)
- The `PASSWORD()` function returned invalid results when used in some `UNION` queries. (Bug #16881)
- `ORDER BY RAND() LIMIT 1` always set a user variable to the last possible value from the table. (Bug #16861)
- Queries containing a subquery that used aggregate functions could return incorrect results. (Bug #16792)
- Concatenating the results of multiple constant subselects produced incorrect results. (Bug #16716)
- When performing a `GROUP_CONCAT()`, the server transformed `BLOB` columns `VARCHAR` columns, which could cause erroneous results when using Connector/J and possibly other MySQL APIs. (Bug #16712)
- Stored procedures did not use the character set defined for the database in which they were created. (Bug #16676)
- Some server errors were not reported to the client, causing both to try to read from the connection until a hang or crash resulted. (Bug #16581)
- If the files for an open table were removed at the OS level (external to the server), the server exited with an assertion failure. (Bug #16532)
- On Windows, a definition for `mysql_set_server_option()` was missing from the C client library. (Bug #16513)
- `mysqlcheck` tried to check views instead of ignoring them. (Bug #16502)
- Updating a column of a `FEDERATED` table to `NULL` sometimes failed. (Bug #16494)
- 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)

- A scheduled event that took longer to execute than the length of time scheduled between successive executions could “skip” executions. For example, an event defined with `EVERY 1 SECOND`—but which required longer than 1 second to complete—might be executed only once every 2 seconds. (Bug #16417)
- A subselect used in the `ON SCHEDULE` clause of a `CREATE EVENT` or `ALTER EVENT` statement caused the server to crash, rather than producing an error as expected. (Bug #16394)
- Grant table modifications sometimes did not refresh the in-memory tables if the host name was `' '` or not specified. (Bug #16297)
- A subquery in the `WHERE` clause of the outer query and using `IN` and `GROUP BY` returned an incorrect result. (Bug #16255)
- A query could produce different results with and without an index, if the `WHERE` clause contained a range condition that used an invalid `DATETIME` constant. (Bug #16249)
- `TIMESTAMPDIFF()` examined only the date and ignored the time when the requested difference unit was months or quarters. (Bug #16226)
- Using tables from MySQL 4.x in MySQL 5.x, in particular those with `VARCHAR` fields and using `INSERT DELAYED` to update data in the table would result in either data corruption or a server crash. (Bug #16218, Bug #17294, Bug #16611)
- The value returned by a stored function returning a string value was not of the declared character set. (Bug #16211)
- The `index_merge/Intersection` optimizer could experience a memory overrun when the number of table columns covered by an index was sufficiently large, possibly resulting in a server crash. (Bug #16201)
- Row equalities (such as `WHERE (a,b) = (c,d)`) were not taken into account by the optimizer, resulting in slow query execution. Now they are treated as conjunctions of equalities between row elements. (Bug #16081)
- Some memory leaks in the `libmysqld` embedded server were corrected. (Bug #16017)
- Values greater than 2 gigabytes used in the `VALUES LESS THAN` clause of a table partitioned by `RANGE` were treated as negative numbers. (Bug #16002)
- A `CREATE TABLE` that produced a `The PARTITION function returns the wrong type` error also caused an `Incorrect information in file` to be printed to `STDERR`, and a junk file to be left in the database directory. (Bug #16000)
- The `max_length` metadata value for columns created from `CONCAT()` could be incorrect when the collation of an argument differed from the collation of the `CONCAT()` itself. In some contexts such as `UNION`, this could lead to truncation of the column contents. (Bug #15962)
- When `NOW()` was used in a `BETWEEN` clause of the definition for a view, it was replaced with a constant in the view. (Bug #15950)
- The server's handling of the number of partitions or subpartitions specified in a `PARTITIONS` or `SUBPARTITIONS` clause was changed. Beginning with this release, the number of partitions must:

- be a positive, nonzero integer
- not have any leading zeros
- not be an expression

Also beginning with this version, no attempt is made to convert, truncate, or evaluate a `PARTITIONS` or `SUBPARTITIONS` value; instead, the `CREATE TABLE` or `ALTER TABLE` statement containing the `PARTITIONS` or `SUBPARTITIONS` clause now fails with an appropriate error message. (Bug #15890)

- Long multiple-row `INSERT` statements could take a very long time for some multibyte character sets. (Bug #15811)
- The C API failed to return a status message when invoking a stored procedure. (Bug #15752)
- `mysqlimport` sends a `set @@character_set_database=binary` statement to the server, but this is not understood by pre-4.1 servers. Now `mysqlimport` encloses the statement within a `/*!40101 ... */` comment so that old servers will ignore it. (Bug #15690)
- `DELETE` with `LEFT JOIN` for InnoDB tables could crash the server if `innodb_locks_unsafe_for_binlog` was enabled. (Bug #15650)
- `BIN()`, `OCT()`, and `CONV()` did not work with BIT values. (Bug #15583)
- Nested natural joins worked executed correctly when executed as a nonprepared statement could fail with an `Unknown column 'col_name' in 'field list'` error when executed as a prepared statement, due to a name resolution problem. (Bug #15355)
- 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)
- Invalid escape sequences in option files caused MySQL programs that read them to abort. (Bug #15328)
- `SHOW GRANTS FOR CURRENT_USER` did not return definer grants when executed in `DEFINER` context (such as within a stored procedure defined with `SQL SECURITY DEFINER`), it returned the invoker grants. (Bug #15298)
- The `--collation-server` server option was being ignored. With the fix, if you choose a nondefault character set with `--character-set-server`, you should also use `--collation-server` to specify the collation. (Bug #15276)
- Re-executing a stored procedure with a complex stored procedure cursor query could lead to a server crash. (Bug #15217)
- The server crashed if it tried to access a `CSV` table for which the data file had been removed. (Bug #15205)
- When using tables containing `VARCHAR` columns created under MySQL 4.1 with a 5.0 or later server, for some queries the metadata sent to the client could have an empty column name. (Bug #14897)
- An invalid comparison between keys with index prefixes over multibyte character fields could lead to incorrect result sets if the selected query execution plan used a range scan by an index prefix over a `UTF8` character field. This also caused incorrect results under similar circumstances with many other character sets. (Bug #14896)

- When setting a column to its implicit default value as the result of inserting a `NULL` into a `NOT NULL` column as part of a multi-row insert or `LOAD DATA` operation, the server returned a misleading warning message. (Bug #14770)
- For `BOOLEAN` mode full-text searches on nonindexed columns, `NULL` rows generated by a `LEFT JOIN` caused incorrect query results. (Bug #14708, Bug #25637)
- The parser rejected queries that selected from a table twice using a `UNION` within a subquery. The parser now supports arbitrary subquery, join, and parenthesis operations within `EXISTS` subqueries. A limitation still exists for scalar subqueries: If the subquery contains `UNION`, the first `SELECT` of the `UNION` cannot be within parentheses. For example, `SELECT (SELECT a FROM t1 UNION SELECT b FROM t2)` will work, but `SELECT ((SELECT a FROM t1) UNION (SELECT b FROM t2))` will not. (Bug #14654)
- Using `SELECT` and a table join while running a concurrent `INSERT` operation would join incorrect rows. (Bug #14400)
- Prepared statements caused general log and server memory corruption. (Bug #14346)
- The binary log lacked character set information for table names when dropping temporary tables. (Bug #14157)
- `libmysqld` produced some warnings to `stderr` which could not be silenced. These warnings now are suppressed. (Bug #13717)
- RPM packages had spurious dependencies on Perl modules and other programs. (Bug #13634)
- `InnoDB` locking was improved by removing a gap lock for the case that you try to delete the same row twice within a transaction. (Bug #13544)
- `REPLACE` statements caused activation of `UPDATE` triggers, not `DELETE` and `INSERT` triggers. (Bug #13479)
- The source distribution failed to compile when configured with the `--without-geometry` option. (Bug #12991)
- With settings of `read_buffer_size` \geq 2G and `read_rnd_buffer_size` \geq 2G, `LOAD DATA INFILE` failed with no error message or caused a server crash for files larger than 2GB. (Bug #12982)
- A `B-TREE` index on a `MEMORY` table erroneously reported duplicate entry error for multiple `NULL` values. (Bug #12873)
- Instance Manager didn't close the client socket file when starting a new `mysqld` instance. `mysqld` inherited the socket, causing clients connected to Instance Manager to hang. (Bug #12751)
- On Mac OS X, zero-byte `read()` or `write()` calls to an SMB-mounted file system could return a nonstandard return value, leading to data corruption. Now such calls are avoided. (Bug #12620)
- `DATE_ADD()` and `DATE_SUB()` returned `NULL` when the result date was on the day '9999-12-31'. (Bug #12356)
- For very complex `SELECT` statements could create temporary tables that were too large, and for which the temporary files were not removed, causing subsequent queries to fail. (Bug #11824)
- After an `INSERT ... ON DUPLICATE KEY UPDATE` statement that updated an existing row, `LAST_INSERT_ID()` could return a value not in the table. (Bug #11460)
- `USE` did not refresh database privileges when employed to re-select the current database. (Bug #10979)

- The server returns a more informative error message when it attempts to open a `MERGE` table that has been defined to use non-`MyISAM` tables. (Bug #10974)
- The type of the value returned by the `VARIANCE()` function varied according to the type of the input value. The function should always return a `DOUBLE` value. (Bug #10966)
- The same trigger error message was produced under two conditions: The trigger duplicated an existing trigger name, or the trigger duplicated an existing combination of action and event. Now different messages are produced for the two conditions so as to be more informative. (Bug #10946)
- A locking safety check in `InnoDB` reported a spurious error `stored_select_lock_type is 0 inside ::start_stmt()` for `INSERT ... SELECT` statements in `innodb_locks_unsafe_for_binlog` mode. The safety check was removed. (Bug #10746)
- `CREATE USER` did not respect the 16-character user name limit. (Bug #10668)
- A server or network failure with an open client connection would cause the client to hang even though the server was no longer available.

As a result of this change, the `MYSQL_OPT_READ_TIMEOUT` and `MYSQL_OPT_WRITE_TIMEOUT` options for `mysql_options()` now apply to TCP/IP connections on all platforms. Previously, they applied only to Windows. (Bug #9678)

- `INSERT INTO ... SELECT ... LIMIT 1` could be slow because the `LIMIT` was ignored when selecting candidate rows. (Bug #9676)
- The optimizer could produce an incorrect result after `AND` with collations such as `latin1_german2_ci`, `utf8_czech_ci`, and `utf8_lithuanian_ci`. (Bug #9509)
- The `DATA DIRECTORY` table option did not work for `TEMPORARY` tables. (Bug #8706)
- A stored procedure with a `CONTINUE` handler that encountered an error continued to execute a statement that caused an error, rather with the next statement following the one that caused the error. (Bug #8153)
- For ODBC compatibility, MySQL supports use of `WHERE col_name IS NULL` for `DATE` or `DATETIME` columns that are `NOT NULL`, to permit column values of `'0000-00-00'` or `'0000-00-00 00:00:00'` to be selected. However, this was not working for `WHERE` clauses in `DELETE` statements. (Bug #8143)
- A user variable set to a value selected from an unsigned column was stored as a signed value. (Bug #7498)
- The `--with-collation` option was not honored for client connections. (Bug #7192)
- With `TRADITIONAL` SQL mode, assignment of out-of-bound values and rounding of assigned values was done correctly, but assignment of the same numbers represented as strings sometimes was handled differently. (Bug #6147)
- On an `INSERT` into an updatable but noninsertable view, an error message was issued stating that the view was not updatable. Now the message says the view is not insertable-into. (Bug #5505)
- `EXPLAIN` sometimes returned an incorrect `select_type` for a `SELECT` from a view, compared to the `select_type` for the equivalent `SELECT` from the base table. (Bug #5500)
- 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)
- Incorporated portability fixes into the definition of `__attribute__` in `my_global.h`. (Bug #2717)

- User-created tables having a name beginning with `#sql` were not visible to `SHOW TABLES` and could collide with internal temporary table names. Now they are not hidden and do not collide. (Bug #1405)

Changes in MySQL 5.1.11 (2006-05-26)

This is a new Beta development release, fixing recently discovered bugs.

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 Enterprise* (a commercial MySQL offering). For more details, please see (<http://www.mysql.com/products/enterprise>).

- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

Functionality Added or Changed

- **Incompatible Change:** The Event Scheduler can now be in one of three states (on, off, or the new suspended state). In addition, due to the fact that `SET GLOBAL event_scheduler;` now acts in a synchronous rather than asynchronous manner, the Event Scheduler thread can be no longer be activated or deactivated at run time.

For more information regarding these changes, see [Event Scheduler Overview](#). (Bug #17619)

- **MySQL Cluster:** The limit of 2048 ordered indexes per cluster has been lifted. There is now no upper limit on the number of ordered indexes (including `AUTO_INCREMENT` columns) that may be used. (Bug #14509)
- Added the `log_queries_not_using_indexes` system variable. (Bug #19616)
- Added the `ssl_ca`, `ssl_capath`, `ssl_cert`, `ssl_cipher`, and `ssl_key` system variables, which display the values given using the corresponding command options. See [Command Options for Secure Connections](#). (Bug #19606)
- The `ENABLE KEYS` and `DISABLE KEYS` clauses for the `ALTER TABLE` statement are now supported for partitioned tables. (Bug #19502)
- Added the `--ssl-verify-server-cert` option to MySQL client programs. This option causes the server's Common Name value in its certificate to be verified against the host name used when connecting to the server, and the connection is rejected if there is a mismatch. Added `MYSQL_OPT_SSL_VERIFY_SERVER_CERT` option for the `mysql_options()` C API function to enable this verification. This feature can be used to prevent man-in-the-middle attacks. Verification is disabled by default. (Bug #17208)
- The default for the `innodb_thread_concurrency` system variable was changed to 8. (Bug #15868)
- It is now possible to use `NEW.var_name` values within triggers as `INOUT` parameters to stored procedures. (Bug #14635)
- Added the `--angel-pid-file` option to `mysqlmanager` for specifying the file in which the angel process records its process ID when `mysqlmanager` runs in daemon mode. (Bug #14106)
- Previously, to build MySQL from source with SSL support enabled, you would invoke `configure` with either the `--with-openssl` or `--with-yassl` option. Those options both have been replaced by the `--with-ssl` option. By default, `--with-ssl` causes the bundled yaSSL library to be used. To select OpenSSL instead, give the option as `--with-ssl=path`, where `path` is the directory where the OpenSSL header files and libraries are located.

- The `mysql_get_ssl_cipher()` C API function was added.
- `mysql_explain_log` (a third-party program) is no longer included in MySQL distributions.

Bugs Fixed

- **Security Fix:** An SQL-injection security hole has been found in multibyte encoding processing. The bug was in the server, incorrectly parsing the string escaped with the `mysql_real_escape_string()` C API function.

This vulnerability was discovered and reported by Josh Berkus <josh@postgresql.org> and Tom Lane <tgl@sss.pgh.pa.us> as part of the inter-project security collaboration of the OSDB consortium. For more information about SQL injection, please see the following text.

Discussion. An SQL injection security hole has been found in multibyte 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 multibyte 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). (Bug #8378, CVE-2006-2753)

References: See also: Bug #8303.

- **MySQL Cluster; Partitioning:** `SELECT MIN(unique_column)` from a Cluster table with user-defined partitioning crashed the server. (Bug #18730)
- **MySQL Cluster; Replication:** (Replication): Memory was not freed after some `ALTER TABLE` operations, which could cause `mysqld` processes to crash. (Bug #19885)
- **MySQL Cluster:** Running `ALL START` in the `NDB` management client or restarting multiple nodes simultaneously could under some circumstances cause the cluster to crash. (Bug #19930)

- **MySQL Cluster:** `TRUNCATE TABLE` failed on tables having `BLOB` or `TEXT` columns with the error `Lock wait timeout exceeded`.

**Note**

This issue affected both in-memory and Disk Data tables.

(Bug #19201)

- **MySQL Cluster:** `ALTER TABLE ENGINE=...` failed when used to change a MySQL Cluster table having no explicit primary key to use a different storage engine.

**Note**

As a consequence of this fix, `SHOW CREATE TABLE` no longer displays auto-partitioning information for `NDBCLUSTER` tables.

(Bug #19010)

- **MySQL 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)
- **MySQL 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 nondefault character sets or collations. (Bug #14918)
- **MySQL Cluster:** Data node failures could cause excessive CPU usage by `ndb_mgmd`. (Bug #13987)
- **Replication:** The embedded server crashed with row-based replication enabled. (Bug #18518)
- **Cluster Replication:** `mysqld` processes did not always detect cluster shutdown, leading to issues with Cluster replication and schema distribution. (Bug #19395)
- **Cluster API:** On big-endian platforms, `NdbOperation::write_attr()` did not update 32-bit fields correctly. (Bug #19537)
- **Cluster API:** The `Ndb::dropEventOperation()` method failed to clean up all objects used, which could cause memory leaks to occur. (Bug #17610)
- The `Data_free` column in the output of `SHOW TABLE STATUS` always displayed 0 for partitioned tables. (Bug #19501)
- Altering a `VARCHAR` column in a `MyISAM` table to make it longer could cause corruption of the following column. (Bug #19386)
- It was not possible to invoke a stored routine containing dynamic SQL from a scheduled event. (Bug #19264)
- Adding an index to a table created using partitioning by `KEY` and the `MEMORY` storage engine caused the server to crash. (Bug #19140)
- Use of uninitialized user variables in a subquery in the `FROM` clause resulted in invalid entries in the binary log. (Bug #19136)
- A `CREATE TABLE` statement that created a table from a materialized view did not inherit default values from the underlying table. (Bug #19089)

- Premature optimization of nested subqueries in the `FROM` clause that refer to aggregate functions could lead to incorrect results. (Bug #19077)
- When creating a table using `CREATE TABLE ... PARTITION BY ... SELECT ...`, the partitioning clause was ignored. (Bug #19062)
- For dates with 4-digit year parts less than 200, an implicit conversion to add a century was applied for date arithmetic performed with `DATE_ADD()`, `DATE_SUB()`, `+ INTERVAL`, and `- INTERVAL`. (For example, `DATE_ADD('0050-01-01 00:00:00', INTERVAL 0 SECOND)` became `'2050-01-01 00:00:00'`.) Now these operations return `NULL` rather than an incorrect non-`NULL` value. (Bug #18997)
- `BLOB` or `TEXT` arguments to or values returned from stored functions were not copied properly if too long and could become garbled. (Bug #18587)
- The client libraries were not compiled for position-independent code on Solaris-SPARC and AMD x86_64 platforms. (Bug #18091, Bug #13159, Bug #14202)
- Returning the value of a system variable from a stored function caused a server crash. (Bug #18037)
- Revised memory allocation for local objects within stored functions and triggers to avoid memory leak for repeated function or trigger invocation. (Bug #17260)
- Symlinking `.mysql_history` to `/dev/null` to suppress statement history saving by `mysql` did not work. (`mysql` deleted the symlink and recreated `.mysql_history` as a regular file, and then wrote history to it.) (Bug #16803)
- `IS_USED_LOCK()` could return an incorrect connection identifier. (Bug #16501)
- Simultaneous scheduled events whose actions conflicted with one another could crash the server. (Bug #16428)
- Concurrent reading and writing of privilege structures could crash the server. (Bug #16372)
- The server no longer uses a signal handler for signal 0 because it could cause a crash on some platforms. (Bug #15869)
- `EXPLAIN ... SELECT INTO` caused the client to hang. (Bug #15463)
- `CREATE TABLE ... SELECT ...` statements that used a stored function explicitly or implicitly (through a view) resulted in a `Table not locked` error. (Bug #15137, Bug #12472)
- Display better error message for `ALTER TABLE` operations that will result in duplicate keys due to `AUTO_INCREMENT` resequencing. (Bug #14573)
- The result from `CONV()` is a string, but was not always treated the same way as a string when converted to a real value for an arithmetic operation. (Bug #13975)
- Within a trigger, `SET` statements used the SQL mode of the invoking statement, not the mode in effect at trigger creation time. (Bug #6951)
- Corrected several problems with the treatment of the `--log-error` option by `mysqld_safe`. These problems were manifest as differences from `mysqld` in error log handling.
 - If a file name was given for `--log-error`, `mysqld_safe` ignored it and did not pass it to `mysqld`, which then wrote error information to `stderr` and resulted in incorrect log rotation when `FLUSH LOGS` was used.
 - `mysqld_safe` now adds `.err` to the end of the file name if no extension is present (the same as `mysqld`).

- `mysqld_safe` treated a relative path name as relative to its own current working directory. Now it treats a relative path name as relative to the data directory (the same as `mysqld`).

In addition, some argument quoting problems were corrected. (Bug #6061)

- `mysqld_safe` treated a relative path name as relative to its own current working directory. Now it treats a relative path name as relative to the data directory (the same as `mysqld`).
- `mysql_safe` now adds `.err` to the end of the file name if no extension is present (the same as `mysqld`).
- If a file name was given for `--log-error`, `mysqld_safe` ignored it and did not pass it to `mysqld`, which then wrote error information to `stderr` and resulted in incorrect log rotation when `FLUSH LOGS` was used.
- The `basedir` and `tmpdir` system variables could not be accessed using `@@var_name` syntax. (Bug #1039)
- The patch for Bug #8303 broke the fix for Bug #8378 and was reverted.

In string literals with an escape character (`\`) followed by a multibyte character that had (`\`) as its second byte, the literal was not interpreted correctly. Now only next byte now is escaped, and not the entire multibyte character. This means it is a strict reverse of the `mysql_real_escape_string()` function.

Changes in MySQL 5.1.10 (Not released)



Note

This was an internal release only, and no binaries were published.

MySQL 5.1.10 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> for finding and reporting these to us.

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 Enterprise* (a commercial MySQL offering). For more details, please see (<http://www.mysql.com/products/enterprise>).

- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

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)
- **MySQL Cluster:** It is now possible to restore a MySQL Cluster backup between big-endian and little-endian machines. (Bug #19255)
- **MySQL Cluster:** It is now possible to perform a partial start of a cluster. That is, it is now possible to bring up the cluster without first running `ndbd --initial` on all configured data nodes. (Bug #18606)

- **MySQL Cluster:** It is now possible to install MySQL with Cluster support to a nondefault location and change the search path for font description files using either the `--basedir` or `--character-sets-dir` options. (Previously in MySQL 5.1, `ndbd` searched only the default path for character sets.)
- **Packaging:** The `MySQL-shared-compat-5.1.X-.i386.rpm` shared compatibility RPMs no longer contain libraries for MySQL 5.0. This avoids a conflict because the 5.0 and 5.1 libraries share the same `soname` number. They now contain libraries for MySQL 3.23, 4.0, 4.1, and 5.1. (Bug #19288)
- SQL syntax for prepared statements now supports `ANALYZE TABLE`, `OPTIMIZE TABLE`, and `REPAIR TABLE`. (Bug #19308)
- The `ONLY_FULL_GROUP_BY` SQL mode now also applies to the `HAVING` clause. That is, columns not named in the `GROUP BY` clause cannot be used in the `HAVING` clause if not used in an aggregate function. (Bug #18739)
- XPath expressions passed to the `ExtractValue()` and `UpdateXML()` functions can now include the colon character (":"). This enables use of these functions with XML which employs namespaces. (Bug #18170)
- On Windows, some names such as `nul`, `prn`, and `aux` could not be used as file names because they are reserved as device names. These are now permissible names in MySQL. They are encoded by appending `@@@` to the name when the server creates the corresponding file or directory. This occurs on all platforms for portability of the corresponding database object between platforms. (Bug #17870)
- The bundled yaSSL library was upgraded to version 1.3.5. This improves handling of certain problems with SSL-related command options. (Bug #17737)
- You must now have the `DROP` privilege to drop table partitions. (Bug #17139)
- Server and clients ignored the `--sysconfdir` option that was passed to `configure`. The directory specified by this option, if set, now is used as one of the standard locations in which to look for option files. (Bug #15069)
- In result set metadata, the `MYSQL_FIELD.length` value for `BIT` columns now is reported in number of bits. For example, the value for a `BIT(9)` column is 9. (Formerly, the value was related to number of bytes.) (Bug #13601)
- The following statements now cause an implicit commit: `ANALYZE TABLE`, `CHECK TABLE`, `OPTIMIZE TABLE`, and `REPAIR TABLE`.
- Added the `KEY_BLOCK_SIZE` table option and index option. This can be used in `CREATE TABLE`, `ALTER TABLE`, and `CREATE INDEX` statements to provide a hint to the storage engine about the size to use for index key blocks. The engine is permitted to change the value if necessary.
- Added the `sql_big_selects` system variable to the output of `SHOW VARIABLES`.
- The `mysql_upgrade` command has been converted from a shell script to a C program, so it is available on non-Unix systems such as Windows. This program should be run for each MySQL upgrade. See [mysql_upgrade — Check and Upgrade MySQL Tables](#).
- Added the `REFERENTIAL_CONSTRAINTS` table to `INFORMATION_SCHEMA`. It provides information about foreign keys.
- Added the `have_dynamic_loading` system variable that indicates whether the server supports dynamic loading of plugins.
- Added `--debug` option to Instance Manager.
- Binary distributions that include SSL support now are built using yaSSL when possible.

Bugs Fixed

- **Security Fix:** A `NUL` byte within a comment in a statement string caused the rest of the string not to be written to the query log, permitting logging to be bypassed. (Bug #17667, CVE-2006-0903)
- **Security Fix:** A malicious client, using specially crafted invalid `COM_TABLE_DUMP` packets was able to trigger an exploitable buffer overflow on the server. Thanks to Stefano Di Paola <stefano.dipaola@wisec.it> for finding and reporting this bug. (CVE-2006-1518)
- **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. (,) Thanks to Stefano Di Paola <stefano.dipaola@wisec.it> for finding and reporting this bug. (CVE-2006-1516, CVE-2006-1517)
- **MySQL Cluster; Replication:** (Replication): Delete and update of rows in a table without a primary key failed on the slave. (Bug #17400)
- **MySQL Cluster:** A 5.1.6 or newer server did not read local checkpoints recorded by any other 5.1 version, thus preventing a system restart following an upgrade. (Bug #19333)
- **MySQL Cluster:** Concurrent `INSERT` and `ROLLBACK` statements from different connections could cause node failures. (Bug #19245)
- **MySQL Cluster:** (Disk Data): Running an `INSERT` and a `DELETE` on a Disk Data table in the same transaction could cause a deadlock. (Bug #19244)
- **MySQL Cluster:** Starting `mysqld` without `--log-bin` caused DDL statements on `NDB` tables to time out. (Bug #19214)
- **MySQL Cluster:** `mysql-test-run.pl` started `NDB` even for test cases that did not need it. (Bug #19083)
- **MySQL Cluster:** Stopping multiple nodes could cause node failure handling not to be completed. (Bug #19039)
- **MySQL Cluster:** The Cluster binlog `mysqld` accepted updates even though the binary log was not set up, which could lead to updates missing from the binary log. (Bug #18932)
- **MySQL Cluster:** `mysqld` could crash when attempting an update if the cluster had failed previously. (Bug #18798)
- **MySQL Cluster:** An `INSERT` or `UPDATE` of more than 128 bytes of data in a 4-replica cluster could cause data nodes to crash. (Bug #18622)
- **MySQL Cluster:** (Disk Data): `CREATE LOGFILE GROUP` accepted values other than `NDB` or `NDBCLUSTER` in the `ENGINE` clause. (Bug #18604)
- **MySQL Cluster:** (Disk Data): Omitting the required `ENGINE` clause from a `CREATE LOGFILE GROUP` or `CREATE TABLESPACE` statement caused the server to crash. An appropriate error message is now returned instead. (Bug #18603)
- **MySQL Cluster:** Queries using `ORDER BY pkN` failed against a `LIST`-partitioned Cluster table having a multi-column primary key, where `pkN` represents one of the columns making up the primary key. (Bug #18598)
- **MySQL Cluster:** A simultaneous `DROP TABLE` and table update operation utilising a table scan could trigger a node failure. (Bug #18597)
- **MySQL Cluster:** Fragment IDs were not logged correctly, causing `ndb_restore_log` to fail. (Bug #18594)

- **MySQL Cluster:** Repeated use of the `SHOW` and `ALL STATUS` commands in the `ndb_mgm` client could cause the `mgmd` process to crash. (Bug #18591)
- **MySQL Cluster:** `ndbd` sometimes failed to start with the error `Node failure handling not completed` following a graceful restart. (Bug #18550)
- **MySQL Cluster:** `ndb_restore` failed to restore a backup made from a 5.0 cluster to a 5.1 cluster. (Bug #18210)
- **MySQL Cluster:** Adding an index to an unsigned integer column did not work correctly. (Bug #18133)
- **MySQL Cluster:** A `SELECT` from an `NDB` table with `ORDER BY indexed_column` and a `LIMIT` clause failed following `ALTER TABLE`. (Bug #18094)
- **MySQL Cluster:** `mysqldump` included in its output data from the internal `cluster` database. (Bug #17840)
- **MySQL Cluster:** Backups could fail for large clusters with many tables, where the number of tables approached `MaxNoOfTables`. (Bug #17607)
- **MySQL Cluster:** Some queries having a `WHERE` clause of the form `c1=val1 OR c2 LIKE 'val2'` were not evaluated correctly. (Bug #17421)
- **MySQL Cluster:** An issue with `ndb_mgmd` prevented more than 27 `mysqld` processes from connecting to a single cluster at one time. (Bug #17150)
- **MySQL 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).

A similar issue could occur when the cluster was first started with a sufficiently low value for this parameter. (Bug #16447, Bug #18612)
- **MySQL Cluster:** Performing multiple `ALTER TABLE` operations on the same `NDB` table from different `mysqld` processes in the same cluster led to schema versioning errors when trying to access the table again following the restart of one of the `mysqld` processes. (Bug #16445)
- **MySQL 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)
- **MySQL Cluster:** On slow networks or CPUs, the management client `SHOW` command could sometimes erroneously show all data nodes as being master nodes belonging to nodegroup 0. (Bug #15530)
- **MySQL Cluster:** Unused open handlers for tables in which the metadata had changed were not properly closed. This could result in stale results from `NDB` tables following an `ALTER TABLE` statement. (Bug #13228)
- **MySQL Cluster:** Uninitialized internal variables could lead to unexpected results. (Bug #11033, Bug #11034)
- **MySQL Cluster:** When attempting to create an index on a `BIT` or `BLOB` column, `Error 743: Unsupported character set in table or index` was returned instead of `Error 906: Unsupported attribute type in index`.
- **Partitioning; Cluster Replication:** Attempting to create an index using multiple columns on an explicitly partitioned table in a replicated Cluster database could cause the master `mysqld` process to crash. (Bug #18284)

- **Replication; Cluster Replication:** An issue with replication caused a `mysqld` connected to a replicated cluster to crash when entering single user mode. (Bug #18535)
- **Replication:** `CREATE VIEW` statements would not be replicated to the slave if the `--replicate-wild-ignore-table` rule was enabled. (Bug #18715)
- **Replication:** Updating a field value when also requesting a lock with `GET_LOCK()` would cause slave servers in a replication environment to terminate. (Bug #17284)
- **Replication:** The binary log would create an incorrect `DROP` query when creating temporary tables during replication. (Bug #17263)
- **Disk Data:** Issuing a `CREATE LOGFILE GROUP` statement during the drop of an `NDB` table would cause database corruption. (Bug #19141)
- **Disk Data:** Concurrent table schema operations and operations on log file groups, tablespaces, data files, or undo files could lead to data node failures. (Bug #18575)
- **Cluster Replication:** Using the `--binlog-do-db` option caused problems with `CREATE TABLE` on the cluster acting as the replication master. (Bug #19492)
- **Cluster Replication:** When taking part in Cluster replication of tables containing `BLOB` columns, `mysqld` falsely reported a large memory leak in the replication buffers when there was none. (Bug #19247)
- **Cluster Replication:** Trying to restore the `apply_status` table from a 5.0 cluster backup failed on a 5.1 server. (Bug #18935)
- **Cluster API:** Passing a nonexistent index name to `NdbIndexScanOperation::setBound()` caused a segmentation fault. (Bug #19088)
- A compatibility issue with NPTL (Native POSIX Thread Library) on Linux could result in a deadlock with `FLUSH TABLES WITH READ LOCK` under some conditions. (Bug #20048)
- Some outer joins were incorrectly converted to inner joins. (Bug #19816)

References: This issue is a regression of: Bug #17146.

- A view definition that referred to an alias in the `HAVING` clause could be saved in the `.frm` file with the alias replaced by the expression that it referred to, causing failure of subsequent `SELECT * FROM view_name` statements. (Bug #19573)
- `mysql` displayed `NULL` for strings that are empty or contain only spaces. (Bug #19564)
- Selecting from a view that used `GROUP BY` on a nonconstant temporal interval (such as `DATE(col) + INTERVAL TIME_TO_SEC(col) SECOND`) could cause a server crash. (Bug #19490)
- An outer join of two views that was written using `{ OJ ... }` syntax could cause a server crash. (Bug #19396)
- An issue with file handling in the partitioning code could cause `mysqld` to crash when started and then stopped within a very short period of time. (Bug #19313)
- `myisamchk` and `myisam_ftdump` should permit either table names or `.MYI` file names as arguments, but permitted only table names. (Bug #19220)
- `InnoDB` could read a delete mark from its system tables incorrectly. (Bug #19217)
- Executing a `CREATE EVENT` statement could cause 100% CPU usage. (Bug #19170)

- Eliminated some memory corruption problems that resulted in `double free or corruption` errors and a server crash. (Bug #19154)
- Attempting to set the default value of an `ENUM` or `SET` column to `NULL` caused a server crash. (Bug #19145)
- Index corruption could occur in cases when `key_cache_block_size` was not a multiple of the `myisam-block-size` value (for example, with `--key_cache_block_size=1536` and `--myisam-block-size=1024`). (Bug #19079)
- Instance Manager now finds the version numbers, so that it works properly when the executable name isn't the same as what the Instance Manager launched (such as when wrapping a `libtool`-wrapped executable from the source tree). (Bug #19059)
- Some fast `ALTER TABLE` operations (requiring no temporary table) did not work for all tables. (Bug #19011)
- Successive `ALTER TABLE ... DROP PARTITION` statements on the same subpartitioned table could eventually cause the server to crash. (Bug #18962)
- 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)
- The parser leaked memory when its stack needed to be extended. (Bug #18930)
- MySQL would not compile on Linux distributions that use the `tinfo` library. (Bug #18912)
- The server attempted to flush uninitialized log tables during `SIGHUP` processing, causing a crash. (Bug #18848)
- For a reference to a nonexistent stored function in a stored routine that had a `CONTINUE` handler, the server continued as though a useful result had been returned, possibly resulting in a server crash. (Bug #18787)
- 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)
- Inserts failed with duplicate key errors on a table partitioned using an `AUTO_INCREMENT` column for the partitioning key. (Bug #18753, Bug #18552)
- It was possible to create a `RANGE`-partitioned table with a partition defined using the clause `VALUES LESS THAN (NULL)`, even though such a partition could never contain any values whatsoever. (Bug #18752)
- Delimited identifiers for partitions were not being treated the same as delimited identifiers for other database objects (such as tables and columns) with regard to permitted characters. (Bug #18750)
- Conversion of a number to a `CHAR UNICODE` string returned an invalid result. (Bug #18691)
- If the second or third argument to `BETWEEN` was a constant expression such as `'2005-09-01 - INTERVAL 6 MONTH` and the other two arguments were columns, `BETWEEN` was evaluated incorrectly. (Bug #18618)
- `LOAD DATA FROM MASTER` failed when trying to load the `INFORMATION_SCHEMA` database from the master, because the `INFORMATION_SCHEMA` system database would already exist on the slave. (Bug #18607)

- Running an `ALTER TABLE` on a partitioned table simultaneously experiencing a high number of concurrent DML statements could crash the server. (Bug #18572)
- 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)
- `mysqltest` incorrectly interpreted some `ER_xxx` error names given in the `error` command. (Bug #18495)
- `InnoDB`: `ALTER TABLE` to add or drop a foreign key for an `InnoDB` table had no effect. (Bug #18477)
- `InnoDB` did not use a consistent read for `CREATE ... SELECT` when `innodb_locks_unsafe_for_binlog` was set. (Bug #18350)
- `DROP DATABASE` did not drop stored routines associated with the database if the database name was longer than 21 characters. (Bug #18344)
- A query on a table partitioned or subpartitioned by `HASH` did not display all results when using a `WHERE` condition involving a column used in the hashing expression. (Bug #18329, Bug #18423)
- In `mysqltest`, `--sleep=0` had no effect. Now it correctly causes `sleep` commands in test case files to sleep for 0 seconds. (Bug #18312)
- The `ExtractValue()` function did not return character data within `<![CDATA[]>` as expected. (Bug #18285)
- A recent change caused the `mysql` client not to display `NULL` values correctly and to display numeric columns left-justified rather than right-justified. The problems have been corrected. (Bug #18265)
- Updates to a `MEMORY` table caused the size of `BTREE` indexes for the table to increase. (Bug #18160)
- A failed `ALTER TABLE` operation could fail to clean up a temporary `.frm` file. (Bug #18129)
- Event-creation statements enclosed in multi-line comments using `/*!version_number ... */` syntax were not parsed correctly. (Bug #18078)
- `SELECT DISTINCT` queries sometimes returned only the last row. (Bug #18068)
- `InnoDB`: A `DELETE` followed by an `INSERT` and then by an `UPDATE` on a partitioned `InnoDB` table caused subsequent queries to return incorrect results. (Bug #17992)
- It was possible to use trailing spaces in the names of partitions and subpartitions. Attempting to do so now raises the error `Incorrect partition name`. (Bug #17973)
- `LIKE` searches failed on a `CHAR` column used as the partitioning column of a table partitioned by `KEY`. (Bug #17946)
- Executing `SELECT` on a large table that had been compressed within `myisampack` could cause a crash. (Bug #17917)
- The `sql_big_selects` system variable was not displayed by `SHOW VARIABLES`. (Bug #17849)
- `REPAIR TABLE` did not restore the length for packed keys in tables created under MySQL 4.x, which caused them to appear corrupt to `CHECK TABLE` but not to `REPAIR TABLE`. (Bug #17810)
- A range access optimizer heuristic was invalid, causing some queries to be much slower in MySQL 5.0 than in 4.0. (Bug #17379, Bug #18940)
- Logging to the `mysql.general_log` and `mysql.slow_log` tables did not work for Windows builds because the `CSV` storage engine was unavailable. The `CSV` engine now is enabled in Windows builds. (Bug #17368)

- If the `WHERE` condition of a query contained an `OR`-ed `FALSE` term, the set of tables whose rows cannot serve for null-complements in outer joins was determined incorrectly. This resulted in blocking possible conversions of outer joins into joins by the optimizer for such queries. (Bug #17164)
- Casting a string to `DECIMAL` worked, but casting a trimmed string (using `LTRIM()` or `RTRIM()`) resulted in loss of decimal digits. (Bug #17043)
- `MyISAM` table deadlock was possible if one thread issued a `LOCK TABLES` request for write locks and then an administrative statement such as `OPTIMIZE TABLE`, if between the two statements another client meanwhile issued a multiple-table `SELECT` for some of the locked tables. (Bug #16986)
- `ALTER TABLE ... REBUILD PARTITION` returned an inaccurate error message. (Bug #16819)
- Use of `--default-storage-engine=innodb` resulted in an error with the server reporting that `InnoDB` was an unknown table type. (Bug #16691)
- `MySQL-shared-compat-5.1.9-0.i386.rpm` incorrectly depended on `glibc` 2.3 and could not be installed on a `glibc` 2.2 system. (Bug #16539)
- The presence of multiple equalities in a condition after reading a constant table could cause the optimizer not to use an index. This resulted in certain queries being much slower than in MySQL 4.1. (Bug #16504)
- Within a trigger, `CONNECTION_ID()` did not return the connection ID of the thread that caused the trigger to be activated. (Bug #16461)
- The XPath `string-length()` function was not implemented for use with `ExtractValue()`. (Bug #16319)
- The `ExtractValue()` function failed with a syntax error when the XPath expression used special characters such as Ñ ("N-tilde"). (Bug #16233)
- The `sql_notes` and `sql_warnings` system variables were not always displayed correctly by `SHOW VARIABLES` (for example, they were displayed as `ON` after being set to `OFF`). (Bug #16195)
- If the first argument to `BETWEEN` was a `DATE` or `TIME` column of a view and the other arguments were constants, `BETWEEN` did not perform conversion of the constants to the appropriate temporary type, resulting in incorrect evaluation. (Bug #16069)
- After calling `FLUSH STATUS`, the `max_used_connections` variable did not increment for existing connections and connections which use the thread cache. (Bug #15933)
- `DELETE` and `UPDATE` statements that used large `NOT IN (value_list)` clauses could use large amounts of memory. (Bug #15872)
- `InnoDB` failure to release an adaptive hash index latch could cause a server crash if the query cache was enabled. (Bug #15758)
- `LAST_INSERT_ID()` in a stored function or trigger returned zero. . (Bug #15728)
- The `system_time_zone` and `version_*` system variables could not be accessed using `SELECT @@var_name` syntax. (Bug #15684, Bug #12792)
- If the server were built without partition support, it was possible to run partitioning-related statements with no errors or warnings, even though these statements would have no effect. Now such statements are not permitted unless the server has been compiled using the `--with-partition` option. (Bug #15561)

- Use of `CONVERT_TZ()` in a view definition could result in spurious syntax or access errors. (Bug #15153)
- Prevent recursive views caused by using `RENAME TABLE` on a view after creating it. (Bug #14308)
- Some queries were slower in 5.0 than in 4.1 because some 4.1 cost-evaluation code had not been merged into 5.0. (Bug #14292)
- Avoid trying to include `<asm/atomic.h>` when it doesn't work in C++ code. (Bug #13621)
- Running `myisampack` followed by `myisamchk` with the `--unpack` option would corrupt the `AUTO_INCREMENT` key. (Bug #12633)
- Use of `CONVERT_TZ()` in a stored function or trigger (or in a stored procedure called from a stored function or trigger) caused an error. (Bug #11081)
- When `myisamchk` needed to rebuild a table, `AUTO_INCREMENT` information was lost. (Bug #10405)

Changes in MySQL 5.1.9 (2006-04-12)

This is a new Beta development release, fixing recently discovered bugs.



Note

This Beta 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.

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 Enterprise* (a commercial MySQL offering). For more details, please see (<http://www.mysql.com/products/enterprise>).

- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

Functionality Added or Changed

- **MySQL Cluster:** The `NDB` storage engine now supports `CREATE TABLE` statements of arbitrary length. (Previously, `CREATE TABLE` statements for MySQL Cluster tables could contain a maximum of 4096 characters only.) (Bug #17813)
- **MySQL Cluster:** Added the `--nowait-nodes` startup option for `ndbd`, making it possible to skip specified nodes without waiting for them to start when starting the cluster. See [ndbd — The MySQL Cluster Data Node Daemon](#).
- `mysqld_safe` no longer checks for a `mysqld-max` binary. Instead, `mysqld_safe` nows checks only for the standard `mysqld` server unless another server binary is specified explicitly using `--mysqld` or `--mysqld-version`. If you previously relied on the implicit invocation of `mysqld-max`, you should use an appropriate option now. (Bug #17861)
- For partitioned tables, the output of `SHOW TABLE STATUS` now shows in the `Engine` column the name of the storage engine used by all partitions for the table; in the `Create_options` column, the output

now shows `partitioned` for a partitioned table. This change also affects the values shown in the corresponding columns of the `INFORMATION_SCHEMA.TABLES` table. (Bug #17631)

- `SHOW PLUGIN` was renamed to `SHOW PLUGINS`. `SHOW PLUGIN` now is deprecated and generates a warning. (Bug #17112)
- Large file support was re-enabled for the MySQL server binary for the AIX 5.2 platform. (Bug #13571)
- Binary MySQL distributions now include a `mysqld-max` server, in addition to the usual `mysqld` optimized server and the `mysqld-debug` debugging server.

Bugs Fixed

- **Security Fix:** Invalid arguments to `DATE_FORMAT ()` caused a server crash. 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. (Bug #20729, CVE-2006-3469)
- **MySQL Cluster; Partitioning:** `BLOB` columns did not work correctly with user-partitioned `NDB` tables. (Bug #16796)
- **MySQL Cluster:** An uninitialized internal variable could lead to unexpected results. (Bug #18831)
- **MySQL Cluster:** `TRUNCATE TABLE` did not reset the `AUTO_INCREMENT` counter for `MyISAM` tables when issued inside a stored procedure.



Note

This bug did not affect `InnoDB` tables.

In addition, `TRUNCATE TABLE` does not reset the `AUTO_INCREMENT` counter for `NDB` tables regardless of when it is called.

(Bug #14945)

References: See also: Bug #18864.

- For full-text searches in boolean mode, and when a full-text parser plugin was used, a `MYSQL_FTPARSER_PARAM::ftparser_state` could have been corrupted by recursive calls to the plugin. (Bug #18836)
- `mysql_reconnect ()` sent a `SET NAMES` statement to the server, even for pre-4.1 servers that do not understand the statement. (Bug #18830)
- A query against a partitioned table using `WHERE col IS NULL` could produce incorrect results given the following conditions:
 - The table had partitions and subpartitions
 - The partitioning function depended on a single column `col` of one of the MySQL integer types
 - The partitioning function was not monotonically increasing

The same issue could cause the server to crash when run in debug mode. (Bug #18659)

- Partition pruning did not work properly for some kinds of partitioning and subpartitioning, with certain `WHERE` clauses. (Partitions and subpartitions that should have been marked as used were not so marked.) The error could manifest as incorrect content in `EXPLAIN PARTITIONS` output as well as missing rows in the results of affected queries. (Bug #18558)

- Building the server using `--with-example-storage-engine` failed to enable the `EXAMPLE` storage engine in the server. (Bug #18464)
- If `InnoDB` encountered a `ER_LOCK_TABLE_FULL` error and rolled back a transaction, the transaction was still written to the binary log. (Bug #18283)
- Complex queries with nested joins could cause a server crash. (Bug #18279)
- `COUNT(*)` on a `MyISAM` table could return different results for the base table and a view on the base table. (Bug #18237)
- `EXTRACT(QUARTER FROM date)` returned unexpected results. (Bug #18100)
- Queries using `WHERE ... IS NULL` returned incorrect results from partitioned tables. (Bug #18070)
- Partition pruning did not perform correctly with partitions on `NULL`, and could potentially crash the server. (Bug #18053)
- `MEDIUMINT` columns were not handled in the same way as other column types by partition pruning. Partition pruning would sometimes use inappropriate columns in performing queries. Both of these issues were rectified as part of the same bug fix. (Bug #18025)
- For tables created in a MySQL 4.1 installation upgraded to MySQL 5.0 and up, multiple-table updates could update only the first matching row. (Bug #16281)
- For `mysql.server`, if the `basedir` option was specified after `datadir` in an option file, the setting for `datadir` was ignored and assumed to be located under `basedir`. (Bug #16240)
- Triggers created in one version of the server could not be dropped after upgrading to a newer version. (Bug #15921)
- `CAST(double AS SIGNED INT)` for large `double` values outside the signed integer range truncated the result to be within range, but the result sometimes had the wrong sign, and no warning was generated. (Bug #15098)
- Quoted values could not be used for partition option values. (Bug #13520)
- Delimited identifiers could not be used in defining partitions. (Bug #13433)
- `mysql_config` returned incorrect libraries on `x86_64` systems. (Bug #13158)
- The server was always built as though `--with-extra-charsets=complex` had been specified. (Bug #12076)

Changes in MySQL 5.1.8 (Not released)



Note

This was an internal release only, and no binaries were published.

- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

Functionality Added or Changed

- **Incompatible Change; Cluster Replication:** The `cluster_replication` database has been renamed to `cluster`. This will effect replication between MySQL Clusters where one cluster is running

MySQL 5.1.8 or later, and the other is running MySQL 5.1.7 or earlier. See [MySQL Cluster Replication](#), and especially [MySQL Cluster Replication Schema and Tables](#).

- **Incompatible Change:** The semantics of `ALTER TABLE t ENGINE=X` for partitioned tables is changed, and now means that the storage engine used for table *t* is changed to *X*.

The previous statement formerly (prior to MySQL 5.1.8) meant that all partitioning was removed from the table. To remove the partitioning of a table, the syntax `ALTER TABLE t REMOVE PARTITIONING` is introduced. The `REMOVE PARTITIONING` option can be used in combination with existing `ALTER TABLE` options such as those employed for adding or dropping columns or indexes. (Bug #17754)
- **Incompatible Change:** For purposes of determining placement, `RANGE` partitioning now treats `NULL` as less than any other value. (Formerly, `NULL` was treated as equal to zero.) See [How MySQL Partitioning Handles NULL](#). (Bug #15447)
- **MySQL Cluster:** The stability of `CREATE` and `DROP` operations on `NDB` tables containing `BLOB` columns has been improved. (Bug #17761)
- **MySQL Cluster:** The `NDBCLUSTER` storage engine now supports `INSERT IGNORE` and `REPLACE` statements. Previously, these statements failed with an error. (Bug #17431)
- **Replication:** Triggers from older servers that included no `DEFINER` clause in the trigger definition now execute with the privileges of the invoker (which on the slave is the slave SQL thread). Previously, replication slaves could not replicate such triggers. (Bug #16266)
- **Replication:** The `binlog_format` system variable now can be set to a third format, `MIXED`, as described in [Replication Formats](#).
- **Replication:** The `binlog_format` system variable now is dynamic and can be changed at runtime, as described in [Replication Formats](#).
- **Replication:** A slave server may now switch the replication format automatically. This happens when the server is running in either `STATEMENT` or `MIXED` format and encounters a row in the binary log that is written in `ROW` logging format. In that case, the slave switches to row-based replication temporarily for that event, and switches back to the previous format afterward.
- **Disk Data:** You can now have only one log file group at any one time. See [CREATE LOGFILE GROUP Syntax](#). (Bug #16386)
- Builds for Windows, Linux, and Unix (except AIX) platforms now have SSL support enabled, in the server as well as in the client libraries. Because part of the SSL code is written in C++, this does introduce dependencies on the system's C++ runtime libraries in several cases, depending on compiler specifics. (Bug #18195)
- Partition pruning was made more stable, particularly in cases involving queries using tests for `NULL` values in the `WHERE` clause against subpartitioned tables which were partitioned by `LIST(some_function(col1, ... ,colN))`. (Bug #17891)
- The output of `SHOW CREATE EVENT` no longer qualifies the event name with the name of the schem to which the event belongs. (Bug #17714)
- The following deprecated constructs now generate warnings, and they are removed as of MySQL 5.5. Where alternatives are shown, applications should be updated to use them. Existing applications that depend on the deprecated constructs should be converted to make use of the current equivalents as soon as possible. You should not employ them in new applications.
 - The `log_bin_trust_routine_creators` system variable (use `log_bin_trust_function_creators`).

- The `table_type` system variable (use `storage_engine`).
- The `TYPE` table option to specify the storage engine for `CREATE TABLE` or `ALTER TABLE` (use `ENGINE`).
- The `SHOW TABLE TYPES` SQL statement (use `SHOW ENGINES`).
- The `SHOW INNODB STATUS` and `SHOW MUTEX STATUS` SQL statements (use `SHOW ENGINE INNODB STATUS` and `SHOW ENGINE INNODB MUTEX`).
- The `SHOW PLUGIN` SQL statement (use `SHOW PLUGINS`).
- The `LOAD TABLE ... FROM MASTER` and `LOAD DATA FROM MASTER` SQL statements (use `mysqldump` or `mysqlhotcopy` to dump tables and `mysql` to reload dump files).
- The `BACKUP TABLE` and `RESTORE TABLE` SQL statements (use `mysqldump` or `mysqlhotcopy` to dump tables and `mysql` to reload dump files).
- `TIMESTAMP(N)` data type: The ability to specify a display width of `N` (use without `N`).
- The `--master-xxx` server options to set replication parameters (use the `CHANGE MASTER TO` statement instead): `--master-host`, `--master-user`, `--master-password`, `--master-port`, `--master-connect-retry`, `--master-ssl`, `--master-ssl-ca`, `--master-ssl-capath`, `--master-ssl-cert`, `--master-ssl-cipher`, `--master-ssl-key`.

In addition, `SHOW BDB LOGS` and `SHOW LOGS` are removed as of MySQL 5.1.12.



Important

TYPE vs ENGINE . In order not to break legacy applications, support for `TYPE = engine_name` —deprecated since MySQL 4.0—has been restored, but now generates a warning.

Beginning with MySQL 5.5, `TYPE = engine_name` will no longer be available and will produce a syntax error.

You should not use `TYPE` in any new applications, and you should immediately begin conversion of existing applications to use the `ENGINE = engine_name` syntax instead.

(Bug #17501)

- Temporary tables may no longer be partitioned. (Bug #17497)
- More specific error messages are now given when attempting to create an excessive number of partitions or subpartitions. (Previously, no distinction was made between an excessive number of partitions and an excessive number of subpartitions.) (Bug #17393)
- Added the `--events` option to `mysqldump` to enable events to be included in the dump output. (Bug #16853)
- For an event having no `STARTS` time specified when it was created, the `mysql.event` table's `start` column now displays the creation time rather than `NULL`.

In addition, both the `SHOW EVENTS` statement's `Starts` column and the `STARTS` column of the `INFORMATION_SCHEMA.EVENTS` table are now empty rather than `NULL` when `STARTS` was not used in the `CREATE EVENT` statement. (Bug #16537)

- Event names are now case-insensitive. That is (for example), you cannot have events with the names `Myevent` and `MyEvent` belonging to the same database and definer. (Bug #16415)
- Description of the `EVENT` privilege has been changed to `To create, alter, drop, and execute events`. (Bug #16412)
- `MICROSECOND` intervals are no longer permitted for events. (Bug #16411)
- Events no longer support times past the end of the Unix epoch. (Formerly, such dates were interpreted as being at the beginning of the Unix epoch.) (Bug #16396)
- The XPath `last()` function is now implemented for use with `ExtractValue()`. (Bug #16318)
- The `ExtractValue()` function with `contains()` now uses the SQL collation in making comparisons. Previously, comparisons were always binary (that is, case-sensitive). (Bug #16316)
- Names of subpartitions must now be unique for an entire table, and not merely within the same partition. (Bug #15408)
- Added the `--sysdate-is-now` option to `mysqld` to enable `SYSDATE()` to be treated as an alias for `NOW()`. See [Date and Time Functions](#). (Bug #15101)
- `mysqldump` now surrounds the `DEFINER`, `SQL SECURITY DEFINER` and `WITH CHECK OPTION` clauses of a `CREATE VIEW` statement with "not in version" comments to prevent errors in earlier versions of MySQL. (Bug #14871)
- The `mysql_ping()` function will now retry if the `reconnect` flag is set and error `CR_SERVER_LOST` is encountered during the first attempt to ping the server. (Bug #14057)
- The `mysqltest` utility now converts all `CR/LF` combinations to `LF` to enable test cases intended for Windows to work properly on UNIX-like systems. (Bug #13809)
- The output from `SHOW CREATE TABLE` is more consistent about using uppercase for keywords. Data types still are in lowercase. (Bug #10460)
- The client API now attempts to reconnect using TCP/IP if the `reconnect` flag is set, as is the case with sockets. (Bug #2845)
- The syntax for `CREATE PROCEDURE` and `CREATE FUNCTION` statements now includes a `DEFINER` clause. The `DEFINER` value specifies the security context to be used when checking access privileges at routine invocation time if the routine has the `SQL SECURITY DEFINER` characteristic. See [CREATE PROCEDURE and CREATE FUNCTION Syntax](#), for more information.

When `mysqldump` is invoked with the `--routines` option, it now dumps the `DEFINER` value for stored routines.

Bugs Fixed

- **MySQL Cluster; Partitioning:** Trying to insert a value into a nonexistent `LIST` partition of an `NDB` table would cause the server to crash.



Note

Beginning with MySQL 5.1.12, user-defined partitioning types other than `KEY` or `LINEAR KEY` were disabled for `NDB` tables.

(Bug #17763)

- **MySQL Cluster; Partitioning:** A repeated `SELECT` on a partitioned table that used the `NDB` storage engine could cause the server to crash. (Bug #17390)

- **MySQL Cluster; Replication:** `AUTO_INCREMENT` values were not propagated correctly in statement-based replication. (Bug #18208)
- **MySQL Cluster; Replication:** Memory was mistakenly freed for `NdbRecAttr` objects during addition of an index while replicating the cluster, which could cause `mysqld` to crash. (Bug #18106)
- **MySQL Cluster; Replication:** Row-based replication could fail with tables using `VARCHAR` columns for primary keys and having `BLOB` columns. (Bug #18067)
- **MySQL Cluster; Replication:** (Replication): The binary log on the secondary master was not being set up correctly following a table rename. (Bug #17838)
- **MySQL Cluster:** Attempting to restart a node with dropped events still pending failed. (Bug #18491)
- **MySQL Cluster:** Two `mysqld` processes starting at the same time could cause a race condition. (Bug #18472)
- **MySQL Cluster:** A timeout in the handling of an `ABORT` condition with more than 32 operations could yield a node failure. (Bug #18414)
- **MySQL Cluster:** Two `mysqld` processes did not synchronize `DROP TABLE` binary log events correctly. (Bug #18395)
- **MySQL Cluster:** A node restart immediately following a `CREATE TABLE` failed.

**Important**

This fix supports 2-node Clusters only.

(Bug #18385)

- **MySQL 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)
- **MySQL Cluster:** When multiple node restarts were attempted without permitting each restart to complete, the error message returned was `Array index out of bounds` rather than `Too many crashed replicas`. (Bug #18349)
- **MySQL 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)
- **MySQL Cluster:** Issuing a `DROP LOGFILE GROUP` statement would cause `ndbd` processes to crash if MySQL had been compiled with `gcc4`. (Bug #18295)
- **MySQL 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)
- **MySQL Cluster:** Insufficient `StringBuffer` memory when attempting to create a trigger caused the server to crash. (Bug #18101)
- **MySQL Cluster:** Variable-length columns used as primary keys were not handled correctly. (Bug #18075)
- **MySQL Cluster:** `CREATE UNIQUE INDEX` on a column containing nonunique data could cause one or more `ndbd` nodes to hang or crash. (Bug #18040)

- **MySQL Cluster:** Node recovery of tables with `VARCHAR` columns using character sets was inconsistent, which could cause a number of issues, including the data nodes failing to restart and `ALTER TABLE` statements to hang. (Bug #18026)
- **MySQL Cluster:** A `SELECT ... ORDER BY` query on an explicitly partitioned Cluster table with no explicit indexes would crash the server. (Bug #17899)
- **MySQL Cluster:** `ALTER TABLE ... ADD INDEX` failed with `ERROR 756: Index on disk column is not supported` when run against a Disk Data table having a primary key. (Bug #17888)
- **MySQL Cluster:** In some cases, a single `ndbd` node failed following a system restart. (Bug #17854)
- **MySQL Cluster:** A simultaneous `RENAME` of several tables was logged multiple times. (Bug #17827)
- **MySQL Cluster:** Trying to perform a `DELETE` from an `NDB` table following a `LOCK TABLES` caused the `ndbd` processes to hang. (Bug #17812)
- **MySQL Cluster:** Trying to update very large partitioned tables using the `NDB` storage engine sometimes caused the server to crash. (Bug #17806, Bug #16385)
- **MySQL Cluster:** Using `ALTER TABLE ... ADD PARTITION` on a table partitioned by `LIST` would cause the client to hang. (Bug #17701)
- **MySQL Cluster:** With a single replica, transactions waiting in the log synchronization queue were not being restarted, causing them to be aborted. (Bug #17536)
- **MySQL Cluster:** `ALTER TABLE` on a partitioned `NDB` table could cause the server to crash. (Bug #17499)
- **MySQL Cluster:** `DELETE` operations on `NDB` tables could cause memory leaks. (Bug #16874)
- **MySQL Cluster:** Some query cache statistics were not always correctly reported for Cluster tables. (Bug #16795)
- **MySQL Cluster:** Restarting nodes were permitted to start and join the cluster too early. (Bug #16772)
- **MySQL Cluster:** `UNDO_BUFFER_SIZE` was limited to 17 MB. (Bug #16657, Bug #17890)
- **MySQL Cluster:** Inserting and deleting `BLOB` column values while a backup was in process could cause data nodes to shut down. (Bug #14028)
- **Replication:** Replication of data stored in a partitioned table would cause slave servers to issue an assertion and terminate. (Bug #18436)
- **Replication:** 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)
- **Replication:** Slave servers would retry the execution of an SQL statement an infinite number of times, ignoring the value `SLAVE_TRANSACTION_RETRIES` when using the `NDB` engine. (Bug #16228)
- **Replication:** The `DEFINER` value for stored routines was not replicated. (Bug #15963)
- **Disk Data:** `CREATE UNIQUE INDEX` failed with `Error 4243: Index not found`. (Bug #18039)
- **Disk Data:** It was not possible to create more than 9 tablespaces. (Bug #16913)
- A `SELECT ... ORDER BY ...` from a view defined using a function could crash the server. An example of such a view is `CREATE VIEW v1 AS SELECT SQRT(c1) FROM t1`. (Bug #18386)
- The server would crash when `SHOW STATUS` was called on a server linked with `yaSSL`. (Bug #18310)

- The `ExtractValue()` function did not return an error when passed an invalid XPath string. (Bug #18172)
- Using the `position()` function in the XPath argument to `ExtractValue()` crashed the server. (Bug #18171)
- `REPAIR TABLE`, `OPTIMIZE TABLE`, and `ALTER TABLE` operations on transactional tables (or on tables of any type on Windows) could corrupt triggers associated with those tables. (Bug #18153)
- Connecting to a server with a UCS2 default character set with a client using a non-UCS2 character set crashed the server. (Bug #18004)
- Using `ALTER TABLE ... REBUILD PARTITION` without specifying the name of the partition caused the server to crash, rather than reporting a syntax error. (Bug #17947)
- `ALTER TABLE ... REBUILD PARTITION` with no partition name specified would crash the server. (Bug #17940)
- A query with a `WHERE date_column > date_value` condition failed on a table partitioned by `RANGE`. (Bug #17894)
- Renaming and adding a new column to a partitioned table in the same `ALTER TABLE` statement caused the server to crash. (Bug #17772)
- **MyISAM**: Performing a bulk insert on a table referenced by a trigger would crash the table. (Bug #17764)
- Using triggers with partitioned **InnoDB** tables led to incorrect results. (Bug #17744)
- Updating a view that filters certain rows to set a filtered out row to be included in the table caused infinite loop. For example, if the view has a `WHERE` clause of `salary > 100` then issuing an `UPDATE` statement of `SET salary = 200 WHERE id = 10`, caused an infinite loop. (Bug #17726)
- A security enhancement in Visual Studio 8 could cause a MySQL debug server compiled with it to hang when running `SELECT` queries against partitioned tables. (Bug #17722)
- The **EXAMPLE** storage engine did not work on Windows. (Bug #17721)
- `ALTER TABLE ... REORGANIZE PARTITION` failed with `Error on rename of filename ...` on Windows. (Bug #17720)
- The MySQL server could crash with out of memory errors when performing aggregate functions on a **DECIMAL** column. (Bug #17602)
- `NULL` values were written to the `mysql.slow_log` table incorrectly. (Bug #17600)
- `mysql_fix_privilege_tables` did not create the `mysql.plugin` table. (Bug #17568)
- Improper checking of binary log statements could result in a server crash. (Bug #17457)
- Repeated invocations of a stored procedure containing a `SHOW CREATE EVENT` statement would result in the error `Packets out of order`. (Bug #17403)
- For **FEDERATED** tables, a `SELECT` statement with an `ORDER BY` clause did not return rows in the proper order. (Bug #17377)
- `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)
- Calling `CREATE TABLE` or `ALTER TABLE` twice on a partitioned table in a stored procedure or a prepared statement resulted in errors and sometimes server crashes. (Bug #17290)

- Checks for permissions on database operations could be performed in a case-insensitive manner (a user with permissions on database `MYDATABASE` could by accident get permissions on database `myDataBase`), if the privilege data were still cached from a previous check. (Bug #17279)
- Stored procedures that call UDFs and pass local string variables caused server crashes. (Bug #17261)
- A problem with `NULLs` and interval mapping sometimes caused incorrect results or crashes when trying to use less-than searches on partitioned tables. (Bug #17173)
- Attempting to add a new partition to a table partitioned by a unique key would cause an `Out of memory` error. (Bug #17169)
- Creating a table with the same name as the mapped name of another table caused a server crash. For example, if MySQL maps the table name `txu#P#p1` to `txu@0023P@0023p1` on disk, creating another table named `txu@0023P@0023p1` crashed the server. (Bug #17142)
- Trying to add a partition to a table having subpartitions could crash the server. (Bug #17140)
- Attempting to use a conflicting `VALUES` clause in `ALTER TABLE ... ADD PARTITION` caused the server to crash. An example of such a conflicting clause would be that uses `VALUES LESS THAN (constant)` (which indicates a range) with a table that is partitioned by `LIST`. (Bug #17127)
- A failed `ALTER TABLE ... ADD PRIMARY KEY` on a partitioned table would result in bad table metadata and could possibly crash the server. (Bug #17097)
- Stored routine names longer than 64 characters were silently truncated. Now the limit is properly enforced and an error occurs. (Bug #17015)
- Cursors in stored routines could cause a server crash. (Bug #16887)
- Triggers created without `BEGIN` and `END` clauses resulted in “You have an error in your SQL syntax” errors when dumping and replaying a binary log. (Bug #16878)
- Using `ALTER TABLE` to increase the length of a `BINARY(M)` column caused column values to be padded with spaces rather than `0x00` bytes. (Bug #16857)
- `ALTER TABLE ... COALESCE PARTITION` failed with an `Out of Memory` error. (Bug #16810)
- `ALTER TABLE ... ADD COLUMN ... AFTER ...` failed when used on partitioned tables. (Bug #16806)
- If the server was started with the `--skip-grant-tables` option, it was impossible to create a trigger or a view without explicitly specifying a `DEFINER` clause. (Bug #16777)
- In a highly concurrent environment, a server crash or deadlock could result from execution of a statement that used stored functions or activated triggers coincident with alteration of the tables used by these functions or triggers. (Bug #16593)
- Clients compiled from source with the `--without-readline` did not save command history from session to session. (Bug #16557)
- Using `ORDER BY intvar` within a stored procedure (where `intvar` is an integer variable or expression) would crash the server.

**Note**

The use of an integer `i` in an `ORDER BY i` clause for sorting the result by the `ith` column is deprecated (and nonstandard). It should *not* be used in new applications. See [SELECT Syntax](#).

(Bug #16474)

- Slow queries executed by scheduled events were not being written to the slow query log. (Bug #16426)
- `INSERT` statements executed by scheduled events were not written to the general log. (Bug #16413)
- Repeated invocations of a stored procedure containing a `CREATE EVENT` or `ALTER EVENT` statement would crash the server. (Bug #16408)
- Names of subpartitions were not displayed in the output of `SHOW CREATE TABLE`. (Bug #16370)
- The `ExtractValue()` function would not accept expressions which matched element names containing an underscore character. (Bug #16320)
- The `self()` XPath function was not handled correctly by `ExtractValue()`. (Bug #16315)
- The `ExtractValue()` function permitted the use of the `!` character in identifiers by ignoring the illegal character. This is now correctly reported as a syntax error. (Bug #16313)
- A memory leak caused warnings on slaves for certain statements that executed without warning on the master. (Bug #16175)
- No error was reported when subpartitions were defined for a nonsubpartitioned table. (Bug #15961)
- 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)
- The `mysql_close()` C API function leaked handles for shared-memory connections on Windows. (Bug #15846)
- A `SELECT` using a function against a nested view would crash the server. (Bug #15683)
- Setting up subpartitions on at least one but not all the partitions of a partitioned table caused the server to crash. (Bug #15407)
- During conversion from one character set to `ucs2`, multibyte characters with no `ucs2` equivalent were converted to multiple characters, rather than to `0x003F QUESTION MARK`. (Bug #15375)
- `CREATE TABLE ... PARTITION ... AS SELECT ...` would cause the server to crash. (Bug #15336)
- When attempting to insert a `0` into a `LIST`-partitioned table that had no value-list containing `0`, no error was reported. (Bug #15253)
- `SELECT COUNT(*)` for a `MyISAM` table could return different results depending on whether an index was used. (Bug #14980)
- Stored routines that contained only a single statement were not written properly to the dumpfile when using `mysqldump`. (Bug #14857)
- Execution of a stored function or trigger which inserted data into a table while running concurrent selects on the same table could result in storing incorrect data in the query cache. (Bug #14767)
- Naming a partition using the characters `Ç` or `ç` ("c-cedilla"; Unicode `00C7` or `00E7`) made unreadable the table containing the partition. (Bug #14527)
- Searches on indexed columns of partitioned tables failed to find all matching rows following updates of the indexed columns. (Bug #14526)

- Creating a partition which depends on an expression containing a column using the UTF8 character set would cause the server to crash. (Bug #14367)
- On Linux, creation of table partitions failed within a stored procedure. (Bug #14363)
- Invoking more than once a prepared statement that creates a partitioned table would crash the server. (Bug #14350)
- The `RENAME TABLE` statement did not move triggers to the new table. (Bug #13525)
- The server would execute stored routines that had a nonexistent definer. (Bug #13198)
- 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)
- Loading of UDFs in a statically linked MySQL caused a server crash. UDF loading is now blocked if the MySQL server is statically linked. (Bug #11835)
- Setting the `myisam_repair_threads` system variable to a value larger than 1 could cause corruption of large `MyISAM` tables. (Bug #11527)
- Issuing `GRANT EXECUTE` on a procedure would display any warnings related to the creation of the procedure. (Bug #7787)

Changes in MySQL 5.1.7 (2006-02-27)

- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

Functionality Added or Changed

- **Incompatible Change:** The `mysql_stmt_attr_get()` C API function now returns a boolean rather than an unsigned int for `STMT_ATTR_UPDATE_MAX_LENGTH`. (Bug #16144)
- **Incompatible Change:** Due to a change in the naming scheme for partitioning and subpartitioning files, it is not possible for the server to read partitioned tables created in previous MySQL versions. Attempting to read pre-5.1.6 partitioned tables with a MySQL 5.1.7 or later server now generates a suitable warning message.

Two possible workarounds are:

- 1. Create a nonpartitioned table with the same table schema using a standard `CREATE TABLE` statement (that is, with no partitioning clauses)

2. Issue a `SELECT INTO` to copy the data into the nonpartitioned table before the upgrade
Following the upgrade, you can partition the new table using `ALTER TABLE ... PARTITION BY ...`
- Alternatively, you can dump the table using `mysqldump` prior to upgrading and reload it afterward with `LOAD DATA`.

In either case, you should drop the pre-5.1.6 partitioned tables before upgrading to 5.1.6 or later.



Important

If any partitioned tables that were created prior to MySQL 5.1.6 are present following an upgrade to MySQL 5.1.6 or later, it is also not possible to read from

the `INFORMATION_SCHEMA.PARTITIONS` table, nor will you be able to drop those tables or the database or databases in which they are located. In this event, you must:

1. Shut down `mysqld`
2. Manually delete the table, partition, and (if any) subpartition files
3. Restart the MySQL Server

(Bug #13437, Bug #16695)

- **Incompatible Change:** `TYPE = engine_name` is no longer accepted as a synonym for the `ENGINE = engine_name` table option. (`TYPE` has been deprecated since MySQL 4.0.)
- **MySQL Cluster:** Attempting to `SELECT ... FROM INFORMATION_SCHEMA.FILES` now raises a warning in the event that the cluster has crashed. (Bug #17087)
- **Replication:** In row-based replication, when executing a `Rows_log_event`, the associated table was locked, the rows applied and the lock released. This did not work since there are storage engines that count locks and perform an autocommit when the number of locks reach zero. Now we ensure that all table maps come before all `ROWS` events in a statement.
- **Disk Data:** Status messages have been added to `ndb_restore` to enable users to know that data files for Disk Data are being created. (Bug #16873)
- **Cluster Replication:** It is now possible to replicate `NDB` tables having no explicit primary key. See [MySQL Cluster Replication](#).
- Creator privileges are now checked for all events before execution. (Bug #17289)
- `CREATE EVENT`, `DROP EVENT`, and `ALTER EVENT` statements are not permitted in triggers. (Bug #16410)
- The SQL mode in effect at the time an event is created or altered is recorded and used during event execution. (Bug #16407)
- 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)
- Added the `--wait-timeout` option to `mysqlmanager` to enable configuration of the timeout for dropping an inactive connection, and increased the default timeout from 30 seconds to 28,800 seconds (8 hours). (Bug #15980, Bug #12674)
- All subpartitions within a given partitioned table are now guaranteed to have unique names. (Bug #15408)
- `mysqlimport` now has a `--use-threads=N` option for loading data files in parallel using `N` threads.
- Added the `RENAME DATABASE` statement.
- Added the `PROCESSLIST` table to `INFORMATION_SCHEMA`.
- Several changes were made to make upgrades easier:
 - Added the `mysql_upgrade` program that checks all tables for incompatibilities with the current version of MySQL Server and repairs them if necessary. This program should be run for each MySQL upgrade (rather than `mysql_fix_privilege_tables`). See [mysql_upgrade — Check and Upgrade MySQL Tables](#).

- Added the `FOR UPGRADE` option for the `CHECK TABLE` statement. This option checks whether tables are incompatible with the current version of MySQL Server.
- Added the `--check-upgrade` to `mysqlcheck` that invokes `CHECK TABLE` with the `FOR UPGRADE` option. Added the `--fix-db-names` and `--fix-table-names` options to `mysqlcheck`.
- Removed the `have_isam` and `have_raid` system variables.
- Added the `IN NATURAL LANGUAGE MODE` and `IN NATURAL LANGUAGE MODE WITH QUERY EXPANSION` modifiers for full-text searches. See [Full-Text Search Functions](#).

Bugs Fixed

- **MySQL Cluster:** Creating `NDB` tables containing `BLOB` columns but no primary key caused unpredictable behavior. (Bug #17559)
- **MySQL Cluster:** Inserting the output of `REPEAT('some_string', some_int)` into a `BLOB` column resulted in the error `Invalid blob attributes or invalid blob parts table`. (Bug #17505)
- **MySQL Cluster:** `ndbd` restarts could sometimes fail due to incorrect memory access. (Bug #17417)
- **MySQL Cluster:** Sharing of table names containing special characters between multiple SQL nodes was not handled correctly when binary logging was enabled (a timeout error resulted). (Bug #17415)
- **MySQL Cluster:** Table definitions were not shared between multiple SQL nodes in a cluster without binary logging being enabled. (Bug #17414)
- **MySQL Cluster:** Cluster log file paths were truncated to 128 characters. They may now be as long as `MAX_PATH` (the maximum path length permitted by the operating system). (Bug #17411)
- **MySQL Cluster:** `SHOW CREATE TABLE` failed when run against a table created in a different session. (Bug #17340)
- **MySQL Cluster:** Following multiple forced shutdowns and restarts of data nodes, `DROP DATABASE` could fail. (Bug #17325)
- **MySQL Cluster:** The `REDO` log would become corrupted (and thus unreadable) in some circumstances, due to a failure in the query handler. (Bug #17295)
- **MySQL 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)
- **MySQL 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)
- **MySQL Cluster:** `CREATE TEMPORARY TABLE` of a Cluster table failed with an `Unsupported` error or crash the server. (Bug #17210, Bug #16552)
- **MySQL Cluster:** The storage engine did not permit views to be updated. (Bug #17206)
- **MySQL Cluster:** When attempting to import data into an `NDB` table using `LOAD DATA INFILE`, the server would hang in the event of a duplicate key error. (Bug #17154)
- **MySQL Cluster:** In some cases, `LOAD DATA INFILE` did not load all data into `NDB` tables. (Bug #17081)
- **MySQL Cluster:** `CREATE TABLE new_tbl LIKE old_tbl;` failed when `old_tbl` used the `NDB` storage engine. (Bug #17005)

- **MySQL Cluster:** An unhandled resources issue could cause node failure with a `DELETE FROM TABLE` affecting thousands of rows. (Bug #16492)
- **MySQL Cluster:** `UNIQUE` keys in Cluster tables were limited to 225 bytes in length. (Bug #15918)
- **MySQL Cluster:** `REPLACE` failed when attempting to update a primary key value in a Cluster table. (Bug #14007)
- **MySQL Cluster:** No error message was generated for setting `NoOfFragmentLogFiles` too low. (Bug #13966)
- **MySQL Cluster:** No error message was generated for setting `MaxNoOfAttributes` too low. (Bug #13965)
- **MySQL Cluster:** Performing large numbers of data manipulation statements on cluster tables using Disk Data could lead to a server crash.
- **Replication; Cluster Replication:** Row-based replication of a cluster failed to take `--binlog-ignore-db` settings into account. (Bug #17188)
- **Replication:** An `ALTER DATABASE` statement on a replication master crashed the slaves. (Bug #17521)
- **Replication:** For a transaction that used `MyISAM` and `InnoDB` tables, interruption of the transaction due to a dropped connection on a master server caused slaves to lose synchrony. (Bug #16559)
- **Replication:** Previously, a stored function invocation was written to the binary log as `DO func_name()` if the invocation changes data and occurs within a nonlogged statement, or if the function invokes a stored procedure that produces an error. These invocations now are logged as `SELECT func_name()` instead for better control over error code checking (slave servers could stop due to detecting a different error than occurred on the master). (Bug #14769)
- **Replication:** `BIT` fields were not properly handled when using row-based replication. (Bug #13418)
- **Disk Data:** In some cases, a cluster using Disk Data tables could not be restarted following a normal shutdown. (Bug #16872)
- **Cluster Replication:** Row-based replication was not set up correctly if a backup was already in progress. For example, connecting a `mysqld` instance to a cluster which was being backed up would result in the message `NDB: skipping setup table tbl_name` being written to the error log. (Bug #17459)
- **Cluster Replication:** Cluster tables not having an explicit primary key could not be replicated. (Bug #14541)
- Column counts were encoded incorrectly in the binary log for row-based logging format. (Bug #17678)
- Data truncations on non-`UNIQUE` indexes could crash `InnoDB` when using multibyte character sets. (Bug #17530)
- Execution times for scheduled events were not calculated correctly: the last execution time was used as a base rather than the actual start time. (Bug #17494)
- Creating an event and using a whitespace character other than space following the `DO` keyword caused a server crash. (Bug #17453)
- Partitioning with certain `SUBPARTITION BY HASH` clauses caused an error when querying for a partitioned column using an `IS NULL` comparison. (Bug #17430, Bug #17432)

- Race conditions between event creation, dropping, and execution could result in a server crash or hang. (Bug #17373)
- Trying to create a partitioned table with more than 32 attributes failed. (Bug #17179)
- Attempting to add a new partition to a table partitioned by a unique key would cause an `Out of memory` error. (Bug #17169)
- `myisam_ftdump` did not work for `FULLTEXT` indexes associated with a parser plugin. (Bug #17116)
- On Windows platforms, some attempts to create partitioned tables from the command line would cause the `mysql` client to hang. (Bug #17082)
- A `SELECT` from the last partition of a subpartitioned table having a `UNIQUE KEY` could crash the MySQL Server. (Bug #16907)
- Statements that contained Unicode characters were not logged to the log tables correctly. (Bug #16905)
- A `SELECT` on a subpartitioned table having a multiple-column `PRIMARY` or `UNIQUE KEY`, and whose partitioning function used only the first column of the key, could cause `mysqld` to crash. (Bug #16901)
- A `RETURN` statement within a trigger caused a server crash. `RETURN` is no longer permitted within triggers. To exit immediately, use `LEAVE`. (Bug #16829)
- Using `REPLACE INTO` on a partitioned table having a primary key would crash the server in the event of a duplicate key error. (Bug #16782)
- `DROP TABLE` would sometimes fail on a table having subpartitions that used the default storage engine. (Bug #16775)
- If the query optimizer transformed a `GROUP BY` clause in a subquery, it did not also transform the `HAVING` clause if there was one, producing incorrect results. (Bug #16603)
- Querying the `INFORMATION_SCHEMA.PARTITIONS` table on a nonmax server caused a server crash. This also happened following the creation of a table with a very large number (hundreds) of partitions. (Bug #16591, Bug #17141)
- `SHOW CREATE EVENT` displayed no output. (Bug #16423)
- `DROP DATABASE` did not drop events for the database. (Bug #16406)
- The `mysql_fix_privilege_tables.sql` script did not properly initialize the `Event_priv` column to 'Y' for those accounts that should have the `EVENT` privilege. (Bug #16400)
- `SELECT` with `GROUP BY` on a view could cause a server crash. (Bug #16382)
- MySQL server dropped client connection for certain `SELECT` statements against views defined that used `MERGE` algorithm. (Bug #16260)
- Using an XPath expression containing `=` with `ExtractValue()` caused the server to crash. (Bug #16242)
- When used with the `ExtractValue()` function, an XPath expression having no leading `/` character would crash the server. (Bug #16234)
- Using `GROUP BY` on column used in `WHERE` clause could cause empty set to be returned. (Bug #16203)
- `CAST(... AS TIME)` operations returned different results when using versus not using prepared-statement protocol. (Bug #15805)

- The `SELECT` privilege was required for triggers that performed no selects. (Bug #15196)
- The `UPDATE` privilege was required for triggers that performed no updates. (Bug #15166)
- 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)
- Killing a long-running query containing a subquery could cause a server crash. (Bug #14851)
- `SUBSTRING_INDEX()` could yield inconsistent results when applied with the same arguments to consecutive rows in a query. (Bug #14676)
- `SET sql_mode = N`, where $N > 31$, did not work properly. (Bug #13897)
- `SHOW CREATE TABLE` produced extraneous spaces following the keywords `PRIMARY KEY`. (Bug #13883)
- `InnoDB` could display an incorrect error message for a cascading update. (Bug #9680)
- `CHECKSUM TABLE` returned different values for `MyISAM` tables depending on whether the `QUICK` or `EXTENDED` option was used. (Bug #8841)
- `SET TRANSACTION ISOLATION LEVEL` acted like `SET SESSION TRANSACTION ISOLATION LEVEL`. That is, it set the isolation level for longer than the next transaction. (Bug #7955)
- Repeated invocation of `my_init()` and `my_end()` caused corruption of character set data and connection failure. (Bug #6536)

Changes in MySQL 5.1.6 (2006-02-01)

- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

Functionality Added or Changed

- **Incompatible Change:** Words with apostrophes are now matched in a `FULLTEXT` search against nonapostrophe words (for example, a search for `Jerry` will match against the term `Jerry's`). Users upgrading to this version must issue `REPAIR TABLE ... QUICK` statements for tables containing `FULLTEXT` indexes. (Bug #14194)
- **Incompatible Change:** This release introduces the `TRIGGER` privilege. Previously, the `SUPER` privilege was needed to create or drop triggers. Now those operations require the `TRIGGER` privilege. This is a security improvement because you no longer need to grant users the `SUPER` privilege to enable them to create triggers. However, the requirement that the account named in a trigger's `DEFINER` clause must have the `SUPER` privilege has changed to a requirement for the `TRIGGER` privilege. After upgrading, be sure to update your grant tables by running `mysql_upgrade`. This will assign the `TRIGGER` privilege to all accounts that had the `SUPER` privilege. (After updating, you might also consider whether any of those accounts no longer need the `SUPER` privilege.) If you fail to update the grant tables, triggers may fail when activated. (Bug #9412)
- **Incompatible Change:** Before MySQL 5.1.6, the server writes general query log and slow query log entries to log files. As of MySQL 5.1.6, the server's logging capabilities for these logs are more flexible. Log entries can be written to log files (as before) or to the `general_log` and `slow_log` tables in the `mysql` database. If logging is enabled, either or both destinations can be selected. The `--log-output` option controls the destination or destinations of log output. See [Selecting General Query and Slow Query Log Output Destinations](#).

If logging is enabled, the default destination now is to log to tables, which differs from earlier versions. If you had the server configured for logging to log files formerly, use `--log-output=FILE` to preserve this behavior after an upgrade to MySQL 5.1.6 or higher.

- **Important Change; MySQL Cluster; Replication:** Replication between MySQL Clusters is now supported. It is now also possible to replicate between a MySQL Cluster and a noncluster database. See [MySQL Cluster Replication](#), for more information.
- **MySQL Cluster:** Added the `ndb_extra_logging` system variable.
- **MySQL Cluster:** The `NDB` storage engine now supports the `CREATE INDEX` and `DROP INDEX` statements.
- **Packaging:** MySQL 5.1.6 introduces some changes to distribution packaging:
 - Distributions include both a `mysqld` optimized server and `mysqld-debug` debugging server. There is no separate debug distribution.
 - There is no longer a `mysqld-max` server. (Note: This changed in MySQL 5.1.9: The `mysqld-max` server also is included in binary distributions.)
 - Server binaries no longer are stripped, except for RPM distributions.
 - Binary distributions for Unix and Unix-like systems no longer include `safe_mysqld` as a link to `mysqld_safe`. `safe_mysqld` has been deprecated since MySQL 4.0 and now is removed.
- The `mysqldump` utility now supports an option for dumping tablespaces. Use `-Y` or `--all-tablespaces` to enable this functionality. (Bug #16753)
- Partition support is not an “engine”, but it was included in the output of `SHOW ENGINES`. Now it is not. The `have_partition_engine` variable was renamed to `have_partitioning`. (Bug #14355, Bug #16718)
- `ANALYZE TABLE` is now supported for partitioned tables. (Bug #13441)
- Added the `--use-threads` option for `mysqlslap`.
- Queries against partitioned tables can now take advantage of partition pruning. In some cases, this can result in query execution that is an order of magnitude faster than the same query against a nonpartitioned version of the same table.
- There is no longer a `mysqld-max` server. (Note: This changed in MySQL 5.1.9: The `mysqld-max` server also is included in binary distributions.)
- Added the `FILES` table to `INFORMATION_SCHEMA`.
- Binary distributions for Unix and Unix-like systems no longer include `safe_mysqld` as a link to `mysqld_safe`. `safe_mysqld` has been deprecated since MySQL 4.0 and now is removed.
- Special characters in database and table identifiers now are encoded when creating the corresponding directory names and file names. This relaxes the restrictions on the characters that can appear in identifiers. See [Mapping of Identifiers to File Names](#).
- Added the `event_scheduler` system variable.
- MySQL 5.1.6 introduces the *Event Scheduler* which enables statements to be scheduled for execution at predetermined times. Events can be transient (one-time-only) or recurrent at regular intervals, and may execute queries and statements permitted in stored routines, including compound statements.

Events can be altered after creation, and dropped when no longer needed.

Information about scheduled events can be obtained using the statements `SHOW EVENTS` and `SHOW CREATE EVENT`, or by querying the `INFORMATION_SCHEMA.EVENTS` table. All of these are available beginning in MySQL 5.1.6.

Users must have the `EVENT` privilege (also added in 5.1.6) to create events.

For more information, see [Using the Event Scheduler](#).

- Distributions include both a `mysqld` optimized server and `mysqld-debug` debugging server. There is no separate debug distribution.
- Server binaries no longer are stripped, except for RPM distributions.
- The `ARCHIVE` storage engine now supports the `AUTO_INCREMENT` column attribute and the `AUTO_INCREMENT` table option. [The ARCHIVE Storage Engine](#).
- Server plugins can register their own status variables to be displayed by the `SHOW STATUS` statement.
- Added the `PARTITIONS` table to `INFORMATION_SCHEMA`.
- Added the `EVENTS` table to `INFORMATION_SCHEMA`.

Bugs Fixed

- **MySQL Cluster:** `NDB` leaked disk space when performing repeated `INSERT` or `DELETE` statements. (Bug #16771)
- **MySQL Cluster:** `ndb_delete_all` ran out of memory when processing tables containing `BLOB` columns. (Bug #16693)
- **MySQL Cluster:** Trying to import too many dumped tables requiring resources beyond those allocated in the cluster configuration file caused the server to crash instead of reporting an insufficient resources error. (Bug #16455)
- **MySQL Cluster:** A `BIT` column whose offset and length totaled 32 caused the cluster to crash. (Bug #16125)
- **MySQL Cluster:** The `ndb_autodiscover` test failed sporadically due to a node not being permitted to connect to the cluster. (Bug #15619)
- **MySQL Cluster:** `NDB` returned an incorrect `Can't find file error` for OS error 24; this has been changed to `Too many open files`. (Bug #15020)
- **MySQL Cluster:** `CREATE TABLESPACE` statements were incorrectly parsed on 64-bit platforms. (`INITIAL SIZE size` worked, but `INITIAL SIZE = size` failed.) (Bug #13556)
- **MySQL Cluster:** Using `mysqldump` to obtain a dump of a partitioned table employing the `NDB` storage engine produced a nonfunctional table creation statement. (Bug #13155)
- **Disk Data:** Tablespaces created using parameters with relatively low values (10 MB or less) produced filesizes much smaller than expected. (Bug #16742)
- **Disk Data:** `NDB` returned the wrong error when the tablespace on disk was full. (Bug #16738)
- **Disk Data:** The error message generated by a failed `ADD UNDOFILE` did not provide any reasons for the failure. (Bug #16267)

- **Disk Data:** `DROP LOGFILE GROUP` corrupted the cluster file system and caused `ndbd` to fail when running more than one node on the same system. (Bug #16193)
- **Cluster API:** 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 incomplete error handling for the failed node. (Bug #15908)
- When the full-text search parser plugin returned more words than half of the length (in bytes) of the query string, the server would crash. (Bug #16722)
- An indexing error sometimes caused values to be assigned to the wrong `RANGE` partition. (Bug #16684)
- An `INSERT` statement in a stored procedure corrupted the binary log. (Bug #16621)
- Trying to add more than one partition in a single `ALTER TABLE ... ADD PARTITION` statement caused the server to crash. (Bug #16534)
- Parallel builds occasionally failed on Solaris. (Bug #16282)
- Inserting a negative value into an integer column used as the partitioning key for a table partitioned by `HASH` could cause the server to crash. (Bug #15968)
- Creating a partitioned table using a storage engine other than the session default storage engine caused the server to crash. (Bug #15966)
- The error message for specifying values for which no partition exists returned wrong values on certain platforms. (Bug #15910)
- Specifying a value for `--tmpdir` without a trailing slash had unpredictable results. (Bug #15904)
- `STR_TO_DATE(1,NULL)` caused a server crash. (Bug #15828, CVE-2006-3081)
- `ALTER TABLE ... ADD PARTITIONS` on a table with one partition crashed the server. (Bug #15820)
- The `mysql_real_connect()` C API function incorrectly reset the `MYSQL_OPT_RECONNECT` option to its default value. (Bug #15719)
- In some cases the query optimizer did not properly perform multiple joins where inner joins followed left joins, resulting in corrupted result sets. (Bug #15633)
- Certain permission management statements could create a `NULL` host name for a user, resulting in a server crash. (Bug #15598)
- Improper memory handling for stored routine variables could cause memory overruns and binary log corruption. (Bug #15588)
- The absence of a table in the left part of a left or right join was not checked prior to name resolution, which resulted in a server crash. (Bug #15538)
- An `ALTER TABLE ... PARTITION BY ...` statement did not have any effect. (Bug #15523)
- Using `RANGE` partitioning with a `CASE` expression as the partitioning function would cause records to be placed in the wrong partition. (Bug #15393)
- Certain subqueries where the inner query was the result of an aggregate function would return different results with MySQL 5.0 or 5.1 than with MySQL 4.1.

Subselects could also return wrong results when the query cache and grouping were involved. (Bug #15347)

- Attempting to insert data into a partitioned table that used the `BLACKHOLE` storage engine caused `mysqld` to crash. (Bug #14524)
- A `FULLTEXT` query in a prepared statement could result in unexpected behavior. (Bug #14496)
- With a table partitioned by `LIST`, inserting a value which was smaller than any value shown in the partitioning value-lists could cause the server to crash. (Bug #14365)
- The `DATA DIRECTORY` and `INDEX DIRECTORY` clauses of a `CREATE TABLE` statement involving partitions did not work. (Bug #14354)
- `SHOW CREATE TABLE` did not display the `PARTITIONS` clause for tables partitioned by `HASH` or `KEY`. (Bug #14327)
- `ALTER TABLE ... DROP PARTITION` would truncate all `DATE` column values in the table's remaining partitions to `NULL`. (Bug #13644)
- `ALTER TABLE ... ADD PARTITION` could crash the server or cause an `Out of memory` error in some circumstances. (Bug #13447)
- The server would permit foreign keys to be declared in the definition of a partitioned table despite the fact that partitioned tables do not support foreign keys (see [Restrictions and Limitations on Partitioning](#)). (Bug #13446)
- A `SELECT` from a key-partitioned table with a multi-column key could cause the server to crash. (Bug #13445)
- Issuing a `TRUNCATE TABLE` statement twice in succession on the same partitioned table would cause the server to crash. (Bug #13442)
- Using a `REPLACE` statement on a partitioned table caused the server to crash. (Bug #13440)
- Using an identifier rather than a literal integer value in the `LESS THAN` clause of a range-partitioned table could cause the server to crash and corruption of tables. (Bug #13439)
- Using `ENGINE=...` within a `PARTITION` clause could cause the server to crash. (Bug #13438)
- `CREATE TABLE ... LIKE` did not work if the table whose schema was to be copied was a partitioned table. (Bug #13435)
- Multibyte path names for `LOAD DATA` and `SELECT ... INTO OUTFILE` caused errors. Added the `character_set_filesystem` system variable, which controls the interpretation of string literals that refer to file names. (Bug #12448)
- Temporary table aliasing did not work inside stored functions. (Bug #12198)
- Using the `TRUNCATE()` function with a negative number for the second argument on a `BIGINT` column returned incorrect results. (Bug #8461)
- Certain Japanese table names were not properly saved during a `CREATE TABLE` statement. (Bug #3906)

Changes in MySQL 5.1.5 (2006-01-10)

- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

Functionality Added or Changed

- **Replication:** Added the `binlog_format` system variable that controls whether to use row-based or statement-based binary logging. Added the `--binlog-format` and `--binlog-row-event-max-size` server options for binary logging control. See [Replication Formats](#).
- Added the `--port-open-timeout` option to `mysqld` to control how many seconds the server should wait for the TCP/IP port to become free if it cannot be opened. (Bug #15591)
- A new statement, `BINLOG`, is generated by `mysqlbinlog` to represent row-based events in binary log files. The statement argument, a base 64-encoded string, is decoded by the server to determine the data change indicated by the corresponding event.
- Added the `--create-schema`, `--lock-directory`, `--number-of-queries`, `--only-print`, `--preserve-schema`, and `--slave` options for `mysqlslap`.
- If `innodb_locks_unsafe_for_binlog` is enabled or if the transaction isolation mode is `READ COMMITTED`, InnoDB can use “semi-consistent” reads. This affects treatment by `UPDATE` statements for rows that are already locked by another transaction. If a row is locked, InnoDB returns the latest committed version to MySQL so that MySQL can determine whether the row matches the `WHERE` condition of the `UPDATE`. If the row matches (must be updated), MySQL reads the row again and this time InnoDB either locks it or waits for a lock on it.

References: See also: Bug #3300.

- Added the `--base64-output` option to `mysqlbinlog` to print all binary log entries using base64 encoding. This is for debugging only. Logs produced using this option should not be applied on production systems.
- Added the `INFORMATION_SCHEMA PLUGINS` table and the `SHOW PLUGIN` statement.
- Two new Hungarian collations are included: `utf8_hungarian_ci` and `ucs2_hungarian_ci`. These support the correct sort order for Hungarian vowels. However, they do not support the correct order for sorting Hungarian consonant contractions; we expect to fix this issue in a future release.
- Plugins now can have status variables that are displayed in the output from `SHOW STATUS`. See [Writing Plugins](#).
- Added the `INFORMATION_SCHEMA ENGINES` table.
- Added the XML functions `ExtractValue()` and `UpdateXML()`. `ExtractValue()` returns the content of a fragment of XML matching a given XPath expression. `UpdateXML()` replaces the element selected from a fragment of XML by an XPath expression supplied by the user with a second XML fragment (also user-supplied), and returns the modified XML. See [XML Functions](#).

Bugs Fixed

- `INSERT DELAYED` caused `mysqld` to crash. (Bug #16095)
- The `--plugin_dir` option was not working. Specifying the parser name for full-text also did not work correctly. (Bug #16068)
- Attempting to insert into a table partitioned by `LIST` a value less than any specified in one of the table's partition definitions resulted in a server crash. In such cases, `mysqld` now returns `ERROR 1500 (HY000): Table has no partition for value v`, where `v` is the out-of-range value. (Bug #15819)
- Issuing a `DROP USER` statement could cause some users to encounter a `hostname is not permitted to connect to this MySQL server` error. (Bug #15775)

- The output of `mysqldump --triggers` did not contain the `DEFINER` clause in dumped trigger definitions. (Bug #15110)
- The output of `SHOW TRIGGERS` contained extraneous whitespace. (Bug #15103)
- Creating a trigger caused a server crash if the table or trigger database was not known because no default database had been selected. (Bug #14863)
- **InnoDB:** Comparison of indexed `VARCHAR CHARACTER SET ucs2 COLLATE ucs2_bin` columns using `LIKE` could fail. (Bug #14583)
- A `COMMIT` statement followed by a `ALTER TABLE` statement on a BDB table caused server crash. (Bug #14212)
- 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)
- **InnoDB:** A semi-consistent read for an `UPDATE` statement with no index column in the `WHERE` condition locked all the rows in the table. (Bug #3300)

Changes in MySQL 5.1.4 (2005-12-21)

- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

Functionality Added or Changed

- Added the `--server-id` option to `mysqlbinlog` to enable only those events created by the server having the given server ID to be extracted. (Bug #15485)
- It is now possible to build the server such that `MyISAM` tables can support up to 128 keys rather than the standard 64. This can be done by configuring the build using the option `--with-max-indexes=N`, where $N \leq 128$ is the maximum number of indexes to permit per table. (Bug #10932)
- Added the `myisam_use_mmap` system variable.
- Added the `--bdb-data-direct` and `--bdb-log-direct` server options.
- Added the `mysqlslap` program, which is designed to emulate client load for a MySQL server and report the timing of each stage. It works as if multiple clients are accessing the server.
- The bundled `BDB` library was upgraded to version 4.4.16.
- Added the `cp1250_polish_ci` collation for the `cp1250` character set.

Bugs Fixed

- **MySQL Cluster:** The `--ndb` option for `perror` did not function. (Bug #15486)
- **MySQL Cluster:** Using `ORDER BY primary_key_column` when selecting from a table having the primary key on a `VARCHAR` column caused a forced shutdown of the cluster. (Bug #15240, Bug #15682, Bug #14828, Bug #15517)
- Server could not be built on default Debian systems with BDB enabled. (Bug #15734)
- `SHOW ENGINES` output showed the `FEDERATED` engine as `DISABLED` even for builds with `FEDERATED` support. (Bug #15559)
- **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)

- It was not possible to reorganize a partition reusing a discarded partition name.

Now, for example, you can create a table such as this one:

```
CREATE TABLE t1 (a INT)
  PARTITION BY RANGE (a) (
    PARTITION p0 VALUES LESS THAN (10),
    PARTITION p1 VALUES LESS THAN (20),
    PARTITION p2 VALUES LESS THAN MAXVALUE
  );
```

and then repartition it as shown here:

```
ALTER TABLE t1 REORGANIZE PARTITION p2 INTO (
  PARTITION p2 VALUES LESS THAN (30)
);
```

Previously, attempting to do so would produce the error `All partitions must have unique names in the table`. (Bug #15521)

- The `BLACKHOLE` storage engine did not handle transactions properly: Rolled-back transactions were written to the binary log. Now they are not. (Bug #15406)
- A left join on a column that having a `NULL` value could cause the server to crash. (Bug #15268)
- Selecting from a view processed with the temptable algorithm caused a server crash if the query cache was enabled. (Bug #15119)
- Creating a view that referenced a stored function that selected from a view caused a crash upon selection from the view. (Bug #15096)
- 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)
- `ROW_COUNT()` returned an incorrect result after `EXECUTE` of a prepared statement. (Bug #14956)
- `ANALYZE TABLE` did not properly update table statistics for a `MyISAM` table with a `FULLTEXT` index containing stopwords, so a subsequent `ANALYZE TABLE` would not recognize the table as having already been analyzed. (Bug #14902)
- Creating a view within a stored procedure could result in an out of memory error or a server crash. (Bug #14885)
- `SELECT` queries that began with an opening parenthesis were not being placed in the query cache. (Bug #14652)
- Space truncation was being ignored when inserting into `BINARY` or `VARBINARY` columns. Now space truncation results in a warning, or an error in strict mode. (Bug #14299)
- The maximum value of `MAX_ROWS` was handled incorrectly on 64-bit systems. (Bug #14155)
- For binary string data types, `mysqldump --hex-blob` produced an illegal output value of `0x` rather than `'\'`. (Bug #13318)
- Some comparisons for the `IN()` operator were inconsistent with equivalent comparisons for the `=` operator. (Bug #12612)

- Attempts to assign `NULL` to a `NOT NULL` column in strict mode now result in a message of `Column 'col_name' cannot be null`, rather than `Column set to default value; NULL supplied to NOT NULL column 'col_name' at row n`. (Bug #11491)
- `SHOW CREATE DATABASE` was sometimes refused when the client had privileges for the database. (Bug #9785)
- Invalid casts to `DATE` values now result in a message of `Incorrect datetime value`, rather than `Truncated incorrect datetime value`. (Bug #8294)
- `mysql` ignored the `MYSQL_TCP_PORT` environment variable. (Bug #5792)

Changes in MySQL 5.1.3 (2005-11-29)

- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

Functionality Added or Changed

- **Incompatible Change; Plugin API:** MySQL 5.1 adds support for a very flexible plugin API that enables loading and unloading of various components at runtime, without restarting the server. Although the work on this is not finished yet, *plugin full-text parsers* are a first step in this direction. This enables users to implement their own input filter on the indexed text, enabling full-text search capability on arbitrary data such as PDF files or other document formats. A pre-parser full-text plugin performs the actual parsing and extraction of the text and hands it over to the built-in MySQL full-text search. (Author: Sergey Vojtovich)

The plugin API requires the `mysql.plugin` table. When upgrading from an older version of MySQL, you should run the `mysql_fix_privilege_tables` command to create this table. See [mysql_fix_privilege_tables — Upgrade MySQL System Tables](#).

Plugins are installed in the directory named by the `plugin_dir` system variable. This variable also controls the location from which the server loads user-defined functions (UDFs), which is a change from earlier versions of MySQL. That is, all UDF library files now must be installed in the plugin directory. When upgrading from an older version of MySQL, you must migrate your UDF files to the plugin directory.

- **Incompatible Change:** Renamed the `table_cache` system variable to `table_open_cache`. Any scripts that refer to `table_cache` should be updated to use the new name.
- **MySQL Cluster:** `VARCHAR` columns used in MySQL Cluster tables are now variable-sized; that is, they now only allocate as much space as required to store the data. Previously, a `VARCHAR(n)` column allocated `n+2` bytes (aligned to 4 bytes), regardless of whether the actual inserted value required that much space. (In other words, a `VARCHAR` column always required the same, fixed, amount of storage as a `CHAR` column of the same size.)
- **Partitioning:** MySQL Server now supports user-defined table partitioning, which enables distributing portions of individual tables across a file system, according to rules which can be set when the table is created. In effect, different portions of a table are stored as separate tables in different locations, but from the user point of view, the partitioned table is still a single table. See [Partitioning](#), for further information on this functionality. (Author: Mikael Ronström)
- `RAND()` no longer permits nonconstant initializers. (Previously, the effect of nonconstant initializers is undefined.) (Bug #6172)

- Added the `table_definition_cache` system variable. If you use a large number of tables, you can create a large table definition cache to speed up opening of tables. The table definition cache takes less space and does not use file descriptors, unlike the normal table cache.
- `SET instance_name.option_name=option_value` sets an option to the specified value and writes it to the config file. See [mysqlmanager — The MySQL Instance Manager](#), for more details on these new commands. (Author: Petr Chardin)
- `SHOW instance_name LOG FILES` provides a listing of all log files used by the instance. (Author: Petr Chardin)
- Added the `SHOW AUTHORS` statement.
- Fast `ALTER TABLE`: Operations that change only table metadata and not table data do not require a temporary table to be used, which improves performance. For example, renaming a column changes only the `.frm` file and no longer uses a temporary table.
- **The Instance Manager (IM)** now has some additional functionality:
 - `SHOW instance_name LOG FILES` provides a listing of all log files used by the instance. (Author: Petr Chardin)
 - `SHOW instance_name LOG {ERROR | SLOW | GENERAL} size` retrieves a part of the specified log file. (Author: Petr Chardin)
 - `SET instance_name.option_name=option_value` sets an option to the specified value and writes it to the config file. See [mysqlmanager — The MySQL Instance Manager](#), for more details on these new commands. (Author: Petr Chardin)
- `SHOW instance_name LOG {ERROR | SLOW | GENERAL} size` retrieves a part of the specified log file. (Author: Petr Chardin)
- Added the `SHOW FUNCTION CODE` and `SHOW PROCEDURE CODE` statements (available only for servers that have been built with debugging support). See [SHOW PROCEDURE CODE Syntax](#).
- The performance of boolean full-text searches (using the “+” Operator) has been improved. See [Full-Text Search Functions](#), for more details about full-text searching. (Author: Sergey Vojtovich)

Bugs Fixed

- `RESET MASTER` failed to delete log files on Windows. One consequence of this change is that server opens the general query and slow log files in shared mode, so now they can be renamed while the server has them open (something not true in previous versions). (Bug #13377)
- Set functions could not be aggregated in outer subqueries. (Bug #12762)

Changes in MySQL 5.1.2 (Not released)

- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

Functionality Added or Changed

- Added the `bdb_cache_parts` and `bdb_region_size` system variables, and permitted `bdb_cache_size` to be larger than 4GB on systems that support it. (Bug #14895)
- Added `MAXLOCKS`, `MINLOCKS`, `MAXWRITE`, and `MINWRITE` as permissible values of the `--bdb-lock-detect` option. (Bug #14876)

- Added `--replace` to `mysqldump`. This option uses `REPLACE INTO`, rather than `INSERT INTO`, when writing the dumpfile.
- Added `Transactions`, `XA`, and `Savepoints` columns to `SHOW ENGINES` output.

Bugs Fixed

- Foreign keys were not properly enforced in `TEMPORARY` tables. Foreign keys are no longer permitted in `TEMPORARY` tables. (Bug #12084)

Changes in MySQL 5.1.1 (Not released)

- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

Functionality Added or Changed

- The `mysql_tableinfo` script has been removed. Information similar to that produced by this script is available from tables in the `INFORMATION_SCHEMA` database.

Bugs Fixed

- **MySQL Cluster; Partitioning:** Specifying the wrong nodegroup in a `CREATE TABLE` statement using partitioning would lead to the table name being locked after the statement failed (that is, the table name could not be re-used). (Bug #12114)
- Using `ORDER BY` in a query with a partitioned table on a 64-bit operating system could crash the server. (Bug #12116)
- Performing a `CREATE TABLE` statement with a `PARTITION BY` clause in a prepared statement could crash a server running in debug mode. (Bug #12097)
- When two threads competed for the same table, a deadlock could occur if one thread also had a lock on another table through `LOCK TABLES` and the thread was attempting to remove the table in some manner while the other thread tried to place locks on both tables. (Bug #10600)