MySQL 8.4 Release Notes

Abstract

This document contains release notes for the changes in MySQL 8.4. For information about changes in a different version of MySQL, see the release notes for that version.

For additional MySQL 8.4 documentation, see the MySQL 8.4 Reference Manual, which includes an overview of features added in MySQL 8.4 (What Is New in MySQL 8.4 since MySQL 8.0), and discussion of upgrade issues that you may encounter while upgrading.

MySQL platform support evolves over time; please refer to https://www.mysql.com/support/supportedplatforms/ database.html for the latest updates.

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 (https://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 the MySQL Forums, where you can discuss your issues with other MySQL users.

Document generated on: 2025-06-04 (revision: 30098)

Table of Contents

Preface and Legal Notices	1
Changes in MySQL 8.4.6 (Not yet released, LTS Release)	. 3
Changes in MySQL 8.4.5 (2025-04-15, LTS Release)	3
Changes in MySQL 8.4.4 (2025-01-21, LTS Release)	8
Changes in MySQL 8.4.3 (2024-10-15, LTS Release)	14
Changes in MySQL 8.4.2 (2024-07-23, LTS Release)	21
Changes in MySQL 8.4.1 (2024-07-01, LTS Release)	21
Changes in MySQL 8.4.0 (2024-04-30, LTS Release)	28

Preface and Legal Notices

This document contains release notes for the changes in MySQL 8.4.

Legal Notices

Copyright © 1997, 2025, Oracle and/or its affiliates.

License Restrictions

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.

Warranty Disclaimer

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

Restricted Rights Notice

If this is software, software documentation, data (as defined in the Federal Acquisition Regulation), 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 embedded, installed, or activated on delivered hardware, and modifications of such programs) and Oracle computer documentation or other Oracle data delivered to or accessed by U.S. Government end users are "commercial computer software," "commercial computer software documentation," or "limited rights data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, reproduction, duplication, release, display, disclosure, modification, preparation of derivative works, and/or adaptation of i) Oracle programs (including any operating system, integrated software, any programs embedded, installed, or activated on delivered hardware, and modifications of such programs), ii) Oracle computer documentation and/ or iii) other Oracle data, is subject to the rights and limitations specified in the license contained in the applicable contract. The terms governing the U.S. Government's use of Oracle cloud services are defined by the applicable contract for such services. No other rights are granted to the U.S. Government.

Hazardous Applications Notice

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.

Trademark Notice

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

Intel and Intel Inside 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, Epyc, and the AMD logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

Third-Party Content, Products, and Services Disclaimer

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.

Use of This Documentation

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.

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 for Accessibility

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.

Changes in MySQL 8.4.6 (Not yet released, LTS Release)

Version 8.4.6 has no release notes, or they have not been published because the product version has not been released.

Changes in MySQL 8.4.5 (2025-04-15, LTS Release)

- Account Management Notes
- Audit Log Notes
- Compilation Notes
- SQL Function and Operator Notes
- INFORMATION_SCHEMA Notes
- Functionality Added or Changed
- Bugs Fixed

Account Management Notes

 It was possible in some cases to grant a user the EXECUTE privilege, but subsequently to be unable to revoke it from the same user. (Bug #37570206)

Audit Log Notes

• <COMMAND_CLASS> was not populated for <NAME>Execute</NAME>.

For more information, see Logging Specific Event Classes. (Bug #36686351)

Compilation Notes

• **Group Replication:** The OpenSSL Engine interface is deprecated, and is no longer being included in OpenSSL v3 main packages by some Linux distributions, including Fedora.

To avoid build issues, the usage of the OpenSSL Engine interface by the Group Communication System (GCS) is now restricted to OpenSSL versions previous to 1.1. (Bug #37475769)

- Linux: Use /usr/bin/gcc (GCC 14.2.1) when building the server on Oracle Linux 10. (Bug #37616148)
- Upgraded the bundled Curl library to version 8.12.1. (Bug #37633587)
- Abseil could not be built on FreeBSD. (Bug #37611924)
- The bundled version of opentelemetry-cpp was upgraded to version 1.19.0. (Bug #37506554)
- In order to use xxhash functions independently from the lz4 library (bundled or source), we compiled xxhash.c into our own binaries, which required using a great many CMake directives. Instead, we now build an interface library for xxhash, and link with that wherever such functions are used. (Bug #37417386)
- Use xxHash-0.8.2 from GitHub rather than the version bundled with Iz4. (Bug #37387318)

SQL Function and Operator Notes

• **Important Change:** When an SQL function is improved from one release to the next, it may throw SQL errors in situations in which it previously did not. If this happens in a table's constraints, default expressions, partitioning expressions, or virtual columns, the table could not be opened. This prevented both analyzing the problem (using, for example, SHOW CREATE TABLE) and addressing it (such as with an ALTER TABLE ... DROP ... statement).

Now, on server upgrade, we scan the data dictionary for tables that use any of the features just mentioned. We then try to open such tables, and if we fail to do so, we alert the user. This patch addresses this. The <u>--check-table-functions</u> server option introduced in this release helps to address this problem by making it possible to specify the server's behavior when encountering an error with such a function. Set this option to WARN in order to log a warning for each table which the server could not open; setting it to ABORT also logs these warnings as WARN, but aborts the server upgrade if any issues were found.

ABORT is the default; this enables the user to fix the issue using the older version of the server before upgrading to the new one. WARN flags the issues, but allows the user to continue in interactive mode while addressing the problem. (Bug #36890891)

References: See also: Bug #37009318. This issue is a regression of: Bug #98950, Bug #98951, Bug #31031886, Bug #31031888.

INFORMATION_SCHEMA Notes

• Fixed a performance issue in the **PROCESSLIST** table. (Bug #36778475)

Functionality Added or Changed

- Important Change: For platforms on which OpenSSL libraries are bundled, the linked OpenSSL library for MySQL Server has been updated to version 3.0.16. For more information, see OpenSSL 3.0 Series Release Notes and OpenSSL Security Advisory (11th February 2025). (Bug #36033684)
- **InnoDB:** To improve debugging, the buf_page_t and buf_block_t structure's metadata is now printed to the error log. (Bug #35115629)

References: See also: Bug #35115601.

• Increased the historical 1024-byte limit when printing the current query during signal handling to 1073741824 (1024 * 1024 * 1024). (Bug #37603354)

Bugs Fixed

• InnoDB: Fixed a potential memory leak in several places in the innobase code. (Bug #37403052)

• **InnoDB:** Under certain circumstances, MySQL could crash during shutdown due to pages which were still fixed or dirty. Errors similar to the following were logged:

[ERROR] [MY-011908] [InnoDB] [FATAL] Page [page id: space=46, page number=75] still fixed or dirty [ERROR] [MY-013183] [InnoDB] Assertion failure: buf0buf.cc:5889:ib::fatal triggered thread 1399637056

(Bug #37391519)

References: See also: Bug #35115601.

• InnoDB: CHECK TABLE for spatial indexes did not verify the MBR against the geometry MBR stored in the clustered index record. This could result in incorrect behaviour of spatial indexes.

As of this release, CHECK TABLE EXTENDED verifies the MBR matches the MBR stored in the clustered index record. (Bug #37359538)

• InnoDB: Fixed an issue relating to pessimistic row update.

Our thanks to Mengchu Shi and the team at Alibaba for the contribution. (Bug #37292404)

- InnoDB: The CHECK TABLE operation could incorrectly report corruption in spatial indexes. (Bug #37286473)
- InnoDB: Fixed an issue relating to InnoDB redo log recovery. (Bug #37061960)
- InnoDB: Fixed an issue relating to reading index_id values. (Bug #36993445, Bug #37709706)
- InnoDB: Fixed an issue relating to lower_case_table_names. (Bug #32288105)
- **InnoDB:** Partition table indexes were not checked when retrieving a record count while that table's definition was being altered by another client session. The record count was executed without error.

As of this release, the index is checked to ensure it is usable when retrieving a record count. (Bug #117459, Bug #37617773)

 InnoDB: Refactored code related to BPR_PCUR_* positioning for restore operations. (Bug #117259, Bug #37505746)

References: This issue is a regression of: Bug #37318367.

- InnoDB: Changes made to innodb_spin_wait_delay in MySQL 8.0.30 negatively impacted performance. (Bug #116463, Bug #37212019)
- **InnoDB:** Under certain circumstances, using ALTER TABLE with INPLACE to modify the size of a column could result in an index which exceeds the valid size limit (767 bytes). This occurred for tables with a row format of Redundant or Compact and the row format was not explicitly defined in the table creation.

As of this release, a validation is performed and an error returned by any ALTER TABLE, INPLACE operation which will result in an invalid index size. (Bug #116353, Bug #37168132)

• InnoDB: Fixed a memory leak in the Clone_persist_gtid thread.

Our thanks to Baolin Huang and the team at Alibaba for the contribution. (Bug #107991, Bug #34454572)

- **Partitioning:** When inserting NOW() into a column not part of the partition key of a partitioned table, all partitions were retrived, and no pruning occurred. (Bug #37397306)
- **Replication:** When the log sanitizer analyzes relay log files, it first searches for a starting point (such as a rotation event or transaction end), but in some cases, it was possible for a binary log file containing a needed GTID to be deleted as having no relevant data; this resulted in the point-in-time recovery process hanging indefinitely while waiting for the missing GTID to be applied. Now in such

cases, the analysis skips parsing of transaction boundaries until the start point is established. (Bug #37635908)

• **Replication:** In a source-replica setup, the replica encountered irregular failures of UPDATE and DELETE statements with ER_KEY_NOT_FOUND errors on the same table. (The replica's binary log and GTID records showed that the row required was committed, and had not been deleted or updated.) This occurred on the replica when the row-matching algorithm used was HASH_SCAN and two rows in the same table had the same CRC32 value.

In the event of such a CRC32 collision, finding a matching CRC32 in the hash table does not guarantee that the correct row is being updated, so the algorithm iterates over any multiple entries having the same CRC32, and compares the full record for each of them in a loop; the problem occurred due to the fact that the logic to exit this loop was incorrect. This logic has now been corrected. (Bug #37462058)

- **Replication:** The asynchronous_connection_failover_delete_source() function did not always perform as expected in all cases. (Bug #36479088)
- **Replication:** In some cases, the asynchronous_connection_failover_add_source() function did not perform as expected. (Bug #36479083)
- **Replication:** In some cases, MASTER_POS_WAIT() did not perform as expected. (Bug #36421684, Bug #37709187)
- **Replication:** The asynchronous_connection_failover_add_managed() function in some cases did not produce the expected result. (Bug #34648589)
- **Replication:** When the server was under a heavy write load, the binary log position for gtid_executed as shown in the Performance Schema log_status table did not match that of the gtid shown in the binary log file.

We fix this by increasing the scope of the lock on the log_status table when querying it to ensure that transactions in the commit pipeline are completed. This ensures that a query against the log_status table waits until gtid_executed has been fully updated, thereby guaranteeing consistency with its position in the binary log. (Bug #102175, Bug #32442772)

• **Group Replication:** When a secondary joined the group, it might happen that all group members started to grow the value of the column COUNT_TRANSACTIONS_ROWS_VALIDATING column of the Performance Schema replication_group_member_stats table indefinitely. This impacted memory consumption in all group members, eventually leading to thrashing if not mitigated by restarting the secondary group member that triggered the behavior, or in some cases, by restarting the whole group.

Analysis pointed to issues with the Group Replication start operation, which checks whether there are partial transactions on the group_replication_applier channel from previous group participation; if any are found, this channel is stopped after applying all complete transactions and its relay logs purged, and then the channel is restarted. After this, distributed recovery is performed, applying any missing data from group members.

The issues arose when the Group Replication pipeline operation for stopping the group_replication_applier channel incorrectly stopped a periodic task from the certifier module, which caused some periodic internal operations not to take place. One of these tasks was the periodic sending of the committed transactions; this omission prevented garbage collection for certification, which in turn caused a continuous increase in COUNT_TRANSACTIONS_ROWS_VALIDATING in the Performance Schema replication_group_member_stats table.

To solve this problem, we have taken steps to ensure that the pipeline operation for stopping the group_replication_applier channel no longer interferes with the certifier module, which also stops spurious values from being added for COUNT_TRANSACTIONS_ROWS_VALIDATING. (Bug #37613510)

• **Group Replication:** When running Group Replication, some transactions may not have write sets, as with empty transactions with GTID_NEXT specified, or DDL statements. For such transactions, Group Replication cannot check conflicts; thus, it is not known whether they can be applied in parallel, and for this reason, Group Replication follows a pessimistic approach, and runs them sequentially, potentially leading to an impact on performance.

While DDL must be applied sequentially, there is no actual reason to force such behavior for empty transactions, so this fix makes it possible for empty transactions to be applied concurrently with other nondependent transactions. (Bug #37597512, Bug #37569333)

- The fprintf_string() function in mysqldump did not use the actual quote character for string escaping. (Bug #37607195)
- EXPLAIN did not always handle subqueries correctly. (Bug #37560280)
- If a demangled function name exceeded 512 bytes in a stack trace, the function name was truncated and a newline was not printed.

As of this release, long strings, such as filenames and demangled functions, are written directly to the output. (Bug #37543598)

- mysqldump did not escape certain special characters properly in its output. With this fix, mysqldump now follows the rules as described in String Literals. (Bug #37540722, Bug #37709163)
- Some operations on tables having functional indexes were not handled properly. (Bug #37523857)
- Attempting to install an unknown component using INSTALL COMPONENT was not always handled correctly. (Bug #37437317)
- The Audit Log plugin did not handle errors correctly when writing JSON output.

See MySQL Enterprise Audit, for more information. (Bug #37370439)

- An UPDATE subsequent to an INSERT affecting a table which had a BEFORE INSERT trigger was sometimes rejected with a null value error when the INSERT had set a NOT NULL column to NULL, even though this should have been allowed by the server sql_mode in effect. (Bug #37337527)
- In some cases, components could not reuse the same connection for running multiple queries. (Bug #37286895)
- Improved error handling for stored routines. (Bug #37193011)
- Stored routines were not always invoked correctly in prepared statements. (Bug #37077424, Bug #37292797)
- Increased the size of SEL_ROOT: :elements from uint16 to size_t. (Bug #36610878)
- Removed an issue with multibyte UTF8 handling. (Bug #36593253)
- An ORDER BY containing an aggregation was not always handled correctly. (Bug #36593244)
- An optimizer hint was ignored, unexpectedly requiring the use of FORCE INDEX, when querying a view that included a UNION. For more information, see Optimizer Hints. (Bug #36536936)
- Some subselects were not handled correctly. (Bug #36421690)
- An invalid DDL statement in certain cases was not always rejected as expected. (Bug #35721121)
- Improved the internal function append_identifier().(Bug #35633084)
- Normally, a view with an unused window definition should be updatable, but when it contained a subquery, it was marked as not updatable. At update time, the window was eliminated, but this was too late to allow an update to be performed.

We fix this by testing mergeability, by checking the presence of window functions, rather than that of window definitions; this allows the view to be updateable, and the problematic UPDATE to succeed. (Bug #35507777)

- In some cases, SET did not perform correctly in prepared statements. (Bug #35308309)
- This fix addresses the following issues:
 - Query_expression::is_set_operation() was not always executed properly.
 - Some sequences of DML statements could lead to an unplanned exit.
 - Some nested subselects were not always handled correctly.

(Bug #34361287, Bug #35889583, Bug #35996409, Bug #36404149, Bug #37611264)

- On Debian, dh_strip_nondeterminism is no longer executed on zip and gzip files within the packages. (Bug #33791880)
- Removed an issue relating to invalid UTF8 values. (Bug #27618273, Bug #37709687)
- Addressed an issue relating to an invalid identifier. (Bug #22958632, Bug #37709664)
- A negative impact in performance was observed when using a multivalued index with ORDER BY DESC and LIMIT in a query, where the value specified by LIMIT was greater than the number of rows actually in the result. (Bug #117085, Bug #37436310)

References: This issue is a regression of: Bug #104897, Bug #33334911.

• If one client session had a uncommitted transaction that caused a DROP TABLE statement in another client session to be blocked, a third client session hung when trying to issue a USE DATABASE statement. (Bug #115706, Bug #36892499)

Changes in MySQL 8.4.4 (2025-01-21, LTS Release)

- Account Management Notes
- C API Notes
- Character Set Support
- Compilation Notes
- Component Notes
- Firewall Notes
- Installation Notes
- Optimizer Notes
- Performance Schema Notes
- Functionality Added or Changed
- Bugs Fixed

Account Management Notes

- The database cache was not flushed properly following execution of DROP USER. (Bug #37132323)
- Failed password validation was not always handled correctly. (Bug #37041439)

C API Notes

• Process memory usage grew when the libmysqlclient API user tried to cache and reuse a prepared statement, preparing it once and then calling either of mysql_stmt_bind_param() or mysql_stmt_bind_named_param() followed by mysql_stmt_execute(), repeatedly without calling the matching mysql_stmt_close(), or calling it in the distant future (on application exit, for example).

We solve this by introducing a separate MEM_ROOT object for storing the bind parameters array, which object can be cleared (deallocating the memory) on each call to mysql_stmt_bind_param() or mysql_stmt_bind_named_param(). (Bug #37202066)

Character Set Support

• Zero rows resulted (where one row was expected) when selecting from a view created with its connection and client character set to latin1, where a query on the view used utf8 as its connection and client character set, the view contained literal values with non-ASCII characters, and the query performed a condition pushdown into a UNION of query blocks of the view.

This problem was related to a previous issue which fixed an error for a similar problem: The problem in that case was properly considering the character set of the view definition when pushing down conditions contained in the view to inner query blocks, but the fix implemented at that time did not take into account the possibility that the view might contain non-ASCII characters.

This meant that the condition to be pushed down was written to a text string with the wrong character set. We fix this oversight by ensuring that the string is created with the correct character set. (Bug #37111452)

References: See also: Bug #36246859.

Compilation Notes

- macOS: Removed obsolete CMake code from MacOS builds. (Bug #37258036)
- macOS: It is now possible to build MySQL using the Homebrew version of Clang. (Bug #37256912)
- macOS: Removed warnings of the form ld: warning: ignoring duplicate libraries and warnings specific to xcodebuild. (Bug #37065301)
- Microsoft Windows: Disabled windows compiler warnings C26445 and C26821 in cmake/ msvc_cppcheck.cmake. Both of these relate to MSVC substitution of gsl::span for std::span, which is not used for MySQL. (Bug #37158156)
- **Solaris:** The minimum required version of GCC to build MySQL on Solaris has been raised to 11.4. See Source Installation Prerequisites, for more information. (Bug #37256600)
- Starting with CMake 3.26, CMake writes the file CMakeFiles/CMakeConfigureLog.yaml, which supersedes CMakeError.log. References to CMakeError.log have therefore been removed. (Bug #37305289)
- Implemented the standards-compliant my_char_traits<unsigned char> for use as a drop-in replacement for std::char_traits<unsigned char>, which was deprecated in Clang 18 and removed in Clang 19. (Bug #37273525, Bug #37785339)
- Removed a **maybe-uninitialized** error found in sql/item.cc when building MySQL with GCC 14. (Bug #37157201)
- The version of libedit used to compile MySQL was upgraded to 20240808-3.1. (Bug #37101293)
- The server could not be built on Ubuntu 22.04 using Clang 13. (Bug #37075154)

- Removed an error in <code>mysql_prepare_create_table()</code> (in the file <code>sql/sql_table.cc</code>) found when compiling MySQL with XCode 16. (Bug #37068527)
- MySQL could not be compiled using Clang 19. (Bug #37014761)
- #include <chrono> was missing from plugin/group_replication/libmysqlgcs/ src/bindings/xcom/xcom/task.cc even though std::chrono::duration_cast() was referenced in this file. (Bug #116779, Bug #37329617)
- The server could not be built on Fedora 40 (and possibly other Linux platforms) using cmake 3.11, due to an issue with TIRPC. (Bug #116164, Bug #37080195)

Component Notes

• INSTALL COMPONENT issued concurrently with a SET PERSIST which used a subquery could sometimes lead to an unplanned exit of the server. (Bug #36559078)

References: See also: Bug #35647759.

Firewall Notes

• In some cases, after performing an upgrade, stored procedures relating to MySQL Enterprise Firewall were not processed correctly. (Bug #36084822)

Installation Notes

• When upgrading from MySQL 5.7 to a later MySQL release series, the system-created mysql.sys and mysql.session accounts are now modified to use the caching_sha2_password authentication plugin instead of the mysql_native_password plugin, which is deprecated in MySQL 8.0, and removed in MySQL 9.0. (Bug #36608636)

Optimizer Notes

• Pushing down a condition which had an aggregate function in a WHERE clause caused the aggregate function to be evaluated when it should not have been. (Bug #36421735)

Performance Schema Notes

• If a user other than root ran START REPLICA, PERFORMANCE_SCHEMA.PROCESSLIST assigned that user's name to the newly spawned foreground replication threads instead of system user.

As of this release, system user is assigned to all foreground system threads. (Bug #37357562)

• Under certain circumstances, a metadata lock can be upgraded or downgraded to a different LOCK_TYPE. This change was not reflected in the PERFORMANCE_SCHEMA.METADATA_LOCKS table.

Our thanks to George Ma and the Alibaba team for the contribution. (Bug #116625, Bug #37271768)

Functionality Added or Changed

• Binary packages that include curl rather than linking to the system curl library have been upgraded to use curl 8.11.1. (Bug #37389565)

Bugs Fixed

• **Incompatible Change:** Corruption occurred in a spatial index when an update of a geometry with a minimal change in the minimum bounding rectangle (MBR) was followed by a delete operation.

When upgrading to this release, it is recommended that you drop any spatial indexes beforehand, then re-create them after the upgrade is complete. Alternatively, you can drop and re-create such indexes immediately following the upgrade, but before making use of any of the tables in which they occur. You should also be aware that downgrading to any previous version reintroduces the original problem described previously.

For more information, see Creating Spatial Indexes. (Bug #36452528)

- InnoDB: Concurrently truncating a table while querying the Performance Schema sometimes cause MySQL to halt unexpectedly. (Bug #37271715)
- **InnoDB:** It was possible for an ALTER TABLE operation using the INPLACE algorithm on a table containing both a spatial index and an auto-increment column to cause corruption or, in debug builds, to trigger a debug assert. This was due to the auto-increment column value being overwritten in the old records of the spatial index while the new record was prepared. (Bug #37189985)
- InnoDB: Certain IO buffer serializations triggered an assertion in debug builds that caused the system to hang. (Bug #37139618)
- InnoDB: Improved InnoDB start up time. (Bug #36880863)

References: This issue is a regression of: Bug #36808732.

• **InnoDB:** An assertion failure was raised when creating a FULLTEXT index on a table with an FTS_DOC_ID value greater than 4294967295. (Bug #36879147)

References: See also: Bug #37387224.

• **INNODB:** Dropping a primary key, and then adding a new AUTO_INCREMENT column as a primary key in descending order using the INPLACE algorithm failed.

Our thanks to Shaohua Wang and the team at Alibaba for the contribution. (Bug #36658450)

- **InnoDB:** Extending a user tablespace produces file extension redo log records (MLOG_FILE_EXTEND), but they were not produced when extending the system tablespace. (Bug #36511673)
- InnoDB: A DELETE operation on a table with a self referential foreign key and full-text index could have triggered an assertion. (Bug #36234681)
- InnoDB: When clearing an AHI index from all buffer pool pages, the block state would potentially change to BUF_BLOCK_MEMORY before acquiring the block mutex thus causing an unexpected halt. (Bug #35037114)
- InnoDB: Common prefix compression for redo log inserts (MLOG_REC_INSERT) was disabled but is now enabled when the versions match. (Bug #34946626)

References: This issue is a regression of: Bug #13899.

- InnoDB: Virtual column information for a row containing an externally stored BLOB was not always logged during an UPDATE operation, which sometimes resulted in an Index PRIMARY is corrupted error. (Bug #34574604)
- InnoDB: ON DELETE CASCADE with generated columns containing secondary indexes sometimes failed, due to virtual column templates not being initialized before deletion.

Our thanks to Rahul Malik for the contribution. (Bug #33691659)

 InnoDB: An update operation attempted to update a virtual column while building an update node for a child table, but should not have done so since foreign key constraints cannot reference virtual columns. (Bug #33327093) • **InnoDB:** It was possible for ALTER TABLE, which rebuilds InnoDB tables using the INPLACE algorithm, to be rejected with a duplicate key error due to a non-duplicate record being inserted concurrently while the rebuild was paused to release a page latch.

Our thanks to Dmitry Lenev and the team at Percona for contributing to this fix. (Bug #115511, Bug #36808088)

- InnoDB: The check enforcing the rule that ALGORITHM=INSTANT cannot be used on a column referenced by a foreign key constraint from another table did not inspect the last field of said constraint. (Bug #115457, Bug #36796094)
- InnoDB: CPU usage statistics did not account for a processor count over 128, which could degrade performance on these larger systems. (Bug #115399, Bug #36765223)
- InnoDB: Executing ALTER TABLE with ADD COLUMN or DROP COLUMN against an empty table now uses the INPLACE algorithm by default instead of INSTANT. This change means the row version is no longer incremented for these simple operations. (Bug #113051, Bug #36004394)
- **INNODB:** AN ALTER TABLE operation that rebuilt an INNODB table using the INPLACE algorithm potentially led to losing a row of data if a purge occurred concurrently on the altered table that contained a clustered or spatial index.

Our thanks to Dmitry Lenev and the team at Percona for contributing to this fix. (Bug #110706, Bug #113812, Bug #115608, Bug #116764, Bug #35303494, Bug #36261480, Bug #36846567, Bug #37318367)

- InnoDB: Queries with a descending primary key and the index_merge optimization sometimes yielded incorrect results such as missing rows. (Bug #106207, Bug #33767814)
- **Replication:** In an InnoDB ClusterSet setup, when autocommit was set to OFF on all nodes in the cluster, a controlled switchover using MySQL Shell was rejected with Error 1105 (Unknown error).

To fix this, we now force a new transaction in

channel_change_source_connection_auto_failover() whenever autocommit=OFF to prevent table access deadlocks when an info repository transaction is executed after changing SOURCE_CONNECTION_AUTO_FAILOVER. (Bug #37173907)

- **Replication:** While large transactions were being received and applied, and a request to stop the replication channel was made using **STOP REPLICA**, MySQL did not do so properly, and subsequently did not process any channel commands. In addition, the server shutdown process did not complete gracefully, and required either the MySQL process to be killed or the host system to be restarted. (Bug #115966, Bug #37008345)
- **Replication:** The log message written when a replica reconnects to the source (when, for example, it is stopped and restarted by issuing STOP REPLICA followed by START REPLICA) While initializing dump thread for replica with UUID *uuid*, found a zombie dump thread with the same UUID. Source is killing the zombie dump thread(*thread_id*) has been improved to Upon reconnection with the replica, while initializing the dump thread for UUID *uuid*, an existing dump thread with the same UUID was detected. The source is terminating the previous dump thread (*thread_id*), which is normal and expected. (Bug #84358, Bug #25330090)
- **Group Replication:** Removed a potential race condition between the internal functions cs::apply::Commit_order_queue::front() and cs::apply::Commit_order_queue::remove(). (Bug #37223451)

References: See also: Bug #35206392.

• Group Replication: When the primary node unexpectedly left the group and quickly attempted to rejoin, the member which had been elected to remove other, faulty members tried to expel or

remove the faulty node but could not do so due to lack of a majority. When the expelled or removed node was the primary, this left the cluster without a primary, resulting in an unusable state. (Bug #36991859)

References: See also: Bug #37181867.

• Group Replication: In some cases, adding a new secondary caused existing secondaries to lag, leading to a deadlock which persisted with no more writes possible until the primary was restarted.

This deadlock occurred between the ticket manager, which ensures that transactions are committed on the correct side of a view change (before or after the view change), and the commit order manager on the inbound replication channel, which ensures that transactions are committed in the same order in which they are received, when these two managers required different orders. This meant that, in some cases, adding a new secondary caused the group primary to be unable to do writes.

We solve this issue by ignoring the commit order manager ordering and enforcing the ticket manager ordering for non-conflicting transactions when such a deadlock occurs. A consequence of this is that replica_preserve_commit_order may not be strictly honored near a View_change_log_event. In other words, replica_preserve_commit_order no longer provides a strict guarantee on an inbound channel on a Group Replication primary. replica_preserve_commit_order still ensures that transactions are ordered correctly, with the only exception being non-conflicting transactions around view change log events. (Bug #35206392)

References: See also: Bug #37223451.

- Group Replication: Improved garbage collection in Group Replication by eliminating non-essential calls to is_subset_not_equals(). (Bug #110673, Bug #35286974)
- **Group Replication:** Removing a group member from a group in which all members were running the same version of MySQL, upgrading it to a later version (from a later release series), and then requesting it to rejoin the group caused the upgraded group member to hang in the recovering state.
- The errors ER_DD_UPDATE_DATADIR_FLAG_FAIL, ER_IB_MSG_FIL_STATE_MOVED_PREV_OR_HAS_DATADIR, ER_RPL_KILL_OLD_DUMP_THREAD_ENCOUNTERED, and ER_RPL_MTA_ALLOW_COMMIT_OUT_OF_ORDER were originally defined in MySQL 8.0, but were subsequently assigned different error code numbers (but with the same names) in MySQL 8.4. The numbers assigned in MySQL 8.0 now apply to MySQL 8.0 only; in the MySQL 8.4 and later release series, only the numbers assigned in MySQL 8.4 are used. (Bug #37284176)
- Added missing checks for NULL input arguments in mle::validate() and validate_session_options(). (Bug #37267887)
- In sql/item_cmpfunc.cc, Item_bool_func2::resolve_type() made an unchecked call to Item_bool_func::resolve_type(); the call to Item_bool_func::resolve_type() ignored its return value, and execution continued even in case of an error. (Bug #37143289)
- AppArmor denied access to /proc/\$pid/task/\$thread_id/mem, a file required to generate a stack trace. (Bug #37063288)

References: See also: Bug #37387034.

- A filter condition in a subquery was sometimes ignored when a query used the index_subquery join type for subquery execution, and the subquery table used materialization in the execution plan. The derived table access path replaced the filter condition, resulting a final plan without the filter layer. To fix this, in such cases, we now add the derived table access path along with the filter access path instead of replacing the latter. (Bug #36918913)
- Some UNION operations similar to a UNION b UNION c ... consumed excessive memory. To help keep this from happening, we now flatten equal set operations at the parsing level, before

contextualization occurs, which should result in reduced resource usage by such operations. (Bug #36652610)

• Improved the internal function my_print_help(). (Bug #36615714)

References: See also: Bug #37387224.

- Removed incorrect code from Acl_cache. (Bug #36608160)
- A subquery containing an aggregate function WITH ROLLUP which was part of a row value comparator was not always processed correctly. (Bug #36593235)

References: See also: Bug #37387180. This issue is a regression of: Bug #30969045, Bug #30921780, Bug #26227613, Bug #29134467, Bug #30967158.

- It was possible for errors raised when persisting variables not to be reported correctly. (Bug #36574732)
- Some subqueries using WITH ROLLUP were not always processed correctly. (Bug #36421704)
- MyISAM block length calculations were not always performed correctly. (Bug #36347992)
- Fixed an issue relating to FTS and concurrent DDL or DML. (Bug #34633727)
- DROP VIEW name was rejected with ER_BAD_TABLE_ERROR if there existed a table with the same name. (Bug #33200087)
- Incorrect results were returned by some queries that used hash antijoins when the hash table did not fit in the join buffer and spilled to disk. (The query triggering the issue actually specified LEFT JOIN, but this was transformed internally from a left outer join to an antijoin.)

The problem was that some rows in the probe table were skipped when writing the probe rows to chunk files, the skipped rows being those that had NULL in part of the join key. Such rows can be skipped for inner joins and semijoins, as they are known to have no match in the build table, but for outer joins and antijoins, rows in the probe table which have no matching row in the build table should be part of the join result, so they must be included in the chunk files.

We already preserved these rows in the chunk files for outer joins. This fix extends the logic used for that purpose so that it also applies for antijoins. (Bug #116334, Bug #37161583)

• In MySQL 8.0 and later, queries of the form SELECT DISTINCT ... FROM t1 WHERE NOT IN(SELECT ...) were transformed into an antijoin if possible, causing the optimizer not to choose a group skip scan for table t1 whereas it would have been chosen in MySQL 5.7. This resulted in a performance degradation for such queries. Group skip scan is not chosen, since the query is now no longer a single-table query following the antijoin transformation, and this access method is enabled only for single table queries. The same behaviour can be seen for queries which are transformed into semijoins as well. In such cases, group skip scan access method can still be used if the access method is used only for duplicate removal (that is, with DISTINCT or GROUP BY, but without aggregate functions).

To fix this, we enable group skip scan when there is only one table in the original query, irrespective of the number of semijoin tables present after internal transformations as long as the query contains no aggregate functions. (Bug #112362, Bug #35842412)

• The mysql client did not allow using '#' or '--' inside an optimizer hint comment.

Our thanks to Kaiwang Chen for the contribution. (Bug #98521, Bug #30875669)

Changes in MySQL 8.4.3 (2024-10-15, LTS Release)

• C API Notes

- Compilation Notes
- Configuration Notes
- SQL Function and Operator Notes
- JavaScript Programs
- Performance Schema Notes
- sys Schema Notes
- Thread Pool Notes
- Functionality Added or Changed
- Bugs Fixed

C API Notes

• The asynchronous interface used unsafe static local variables to store state information. (Bug #115703, Bug #36891894)

Compilation Notes

- MySQL did not compile on Fedora 41. (Bug #37046924)
- MySQL did not compile on Ubuntu 24.10. (Bug #37042308)
- Added CONTRIBUTING.md and SECURITY.md files to the MySQL sources to conform to Oracle's Open Source guidelines. (Bug #36998165)
- Aligned CMAKE_MINIMUM_REQUIRED with the correct required CMake version (3.14.6) as well as CMake policies for third-party libraries used by MySQL. (Bug #36978193)
- Binary packages that include curl rather than linking to the system curl library have been upgraded to use curl 8.9.1. Important issues fixed in curl version 8.9.1 are described at https://curl.se/docs/security.html. (Bug #36967379, Bug #36955197)
- The included zlib library has been upgraded from version 1.2.13 to version 1.3.1. (Bug #36950863)
- CMake options BUILD_SHARED_LIBS and CMAKE_SKIP_INSTALL_ALL_DEPENDENCY are now set to OFF in the top level of the build. (Bug #36930664)
- Pulling in the most recent CMake packages for gRPC and Protobuf caused the builds for these libraries to fail. Fixed by rewriting the package lookup to work for both older and newer CMake packages. (Bug #36905657)
- Upgraded the bundled libcurl library to version 8.9.0. (Bug #36886877)
- The bundled 1z4 library was upgraded to version 1.10.0. (Bug #36886747)
- The -DWITHOUT_SERVER option for CMake did not work on Enterprise Linux 7. (Bug #36824515)
- Linking with mold did not work on Enterprise Linux 9. (Bug #36818902)

References: See also: Bug #34099162.

- The version of libfido used with MySQL was upgraded to 1.15.0. (Bug #36752604)
- The RE2 regular expression library is now bundled with the MySQL sources. (Bug #36729026)
- The version of clang-format used for the MySQL codebase was upgraded from 10 to 15. (Bug #36500268)

- The bundled version of TI-RPC was upgraded to 1.3.5. (Bug #115698, Bug #36886602)
- It was not possible to build MySQL using Protobuf 22 or later.

Our thanks to Gordon Wang for the contribution. (Bug #115163, Bug #36678092)

Configuration Notes

- **Microsoft Windows:** On Windows, *MySQL Configurator* revert button functionality did not always revert to the original value. (Bug #36702176)
- **Microsoft Windows:** On Windows, the *MySQL Configurator* data directory file browsing mechanism required the **ProgramData** folder but can now function outside of that folder. (Bug #36702072)
- Microsoft Windows: On Windows, clicking the [X] close button on a *MySQL Configurator* wizard's page now yields a confirmation popup if the wizard is busy executing an operation. (Bug #36671317)
- **Microsoft Windows:** On Windows, *MySQL Configurator* no longer defines a custom server_type
 variable in the generated MySQL Server configuration file. This information is now stored in the configurator_settings.xml file. (Bug #36670309)
- **Microsoft Windows:** On Windows, the Removing Windows Firewall step in *MySQL Configurator* would fail if the my.ini file was missing a mysqlx_port definition. (Bug #36666260)
- **Microsoft Windows:** On Windows, *MySQL Configurator* no longer performs unnecessary background operations after clicking **Next** if the selected options did not change. (Bug #36395605)
- **Microsoft Windows:** On Windows, if *MySQL Configurator* failed to find a valid my.ini or my.cnf file from the *MySQL Server Installations* page, then clicking the **Browse** button disabled the **Next** button even when the selected file was valid. (Bug #36395569)
- **Microsoft Windows:** On Windows, *MySQL Configurator* now only shows the removal steps if the associated MySQL Server was previously configured. (Bug #36395417)
- **Microsoft Windows:** On Windows, MySQL Configurator described an empty password as strong rather than weak. (Bug #35533726)

SQL Function and Operator Notes

• The output from DATABASE() was truncated when this function was used as part of a UNION query. (Bug #36871927)

JavaScript Programs

• Employing a user variable in a DML statement within a stored procedure that used the Statement handle interface sometimes led to an unplanned server exit when any statement was executed following execution of such a stored procedure. To address this issue, we now reset the memory used to the store the user variable and its value to the statement execution memory root after executing a sub-statement from the Statement handle interface. (Bug #36892945)

Performance Schema Notes

- Fixed issues relating to the OTLP exporter. (Bug #36792180, Bug #36783070)
- Running SELECT * FROM sys.innodb_lock_waits; on an instance which was under heavy load affected the performance of the server.

As of this release, SELECT * FROM sys.innodb_lock_waits; fetches only 2 locks for each wait, instead of scanning all locks twice for each wait.

As part of this fix, primary keys were added to DATA_LOCKS and DATA_LOCK_WAITS. (Bug #100537, Bug #31763497)

sys Schema Notes

• The performance of the innodb_lock_waits view is improved in this release. (Bug #36337708)

Thread Pool Notes

• Some concurrent connection attempts were not handled correctly. (Bug #36625082)

Functionality Added or Changed

- **Important Change:** For platforms on which OpenSSL libraries are bundled, the linked OpenSSL library for MySQL Server has been updated to version 3.0.15. For more information, see OpenSSL 3.0 Series Release Notes and OpenSSL Security Advisory [3rd September 2024]. (Bug #37021075)
- **Performance; Replication:** The data structure used in tracking binary log transaction dependencies has been changed from Tree to ankerl::unordered_dense::map, which uses approximately 60% less space, and which should thus contribute to better dependency tracking performance. (Bug #37008442, Bug #37529256)
- Added the -system-command option for the <code>mysql</code> client, which enables or disables the <code>system</code> client command.

This option is enabled by default. To disable it, use --system-command=OFF or --skip-systemcommand, which causes the system command to be rejected with an error. (Bug #36377685, WL #16482)

References: See also: Bug #36248967.

Bugs Fixed

• **Performance; InnoDB:** Several functions internal to InnoDB, which were defined as inline in MySQL 8.0.28, were found to be no longer inline in MySQL 8.0.33, due in part to refactoring which accompanied improvements made in MySQL 8.0.30 to improve the InnoDB adaptive hash index. This had an adverse effect on queries using joins on InnoDB tables. (Bug #111538, Bug #35531293)

References: This issue is a regression of: Bug #81814, Bug #16739204, Bug #23584861.

- **InnoDB:** InnoDB did not allow updating a **REDUNDANT** table column that was altered with **NULL** as the default value using the **INSTANT** algorithm. Instead, MySQL unexpectedly halted. (Bug #36840107)
- InnoDB: When redo log capacity was reduced with log_files_truncate, under rare circumstances the file.end_lsn equaled log_sys->write_lsn thus disallowing future redo log writes. (Bug #36730830)
- InnoDB: It was possible for the MySQL server to halt unexpectedly when executing a DELETE or UPDATE statement after a column was dropped using the INSTANT algorithm. (Bug #36723117)
- **InnoDB:** FTS index optimization would sometimes not function correctly with tokens equal in collation order but different in terms of bytes. (Bug #36652127)
- InnoDB: Replaced std::this_thread::sleep_for(std::chrono::seconds(0)) usage with std::this_thread::yield() to prevent it from being optimized out while compiled. (Bug #36522343)
- InnoDB: The upgrade process unexpectedly halted when the database contained a full-text index created in MySQL 5.7 with a user-created FTS_DOC_ID column as the primary key.

Our thanks to Huaxiong Song and the team at Alibaba for contributing to this fix. (Bug #36496164)

- InnoDB: Disabled the optimizer for indexes with range conditions on multiple-value virtual columns. (Bug #36341532)
- InnoDB: Redesigned the performance schema data_locks and data_lock_waits tables so that querying them does not require an exclusive global mutex on the transaction or lock system. It now iterates over buckets of hash tables that hold the locks to only latch the actively processed shard, when previously it iterated over the transactions. This also improves the iteration logic complexity in terms of speed and memory to decrease the impact of these queries on the rest of the system.

Note that the query result might show an incomplete list of transaction locks if it committed, started, or otherwise changed the set of owned locks in-between visiting two buckets. This differs from previous behavior which always showed a consistent snapshot of locks held by individual transactions, although two different transactions could have been presented at different moments. In other words, the new approach gives a consistent view of a single wait queue to show conflicting locks with a waiting lock because they are always in the same bucket, while the old approach could miss some of them because they belonged to other transactions. The old approach would always show all the other locks held by a reported transaction but could miss locks of other transactions even if they were conflicting. (Bug #36302624)

- **InnoDB:** On Windows, fixed a doublewrite buffer regression that slowed file access, and refactored FILE_FLAG_OVERLAPPED flag usage for opening files. (Bug #36259487)
- InnoDB: A schema mismatch was possible when a table was imported with a different sql_mode than when it was created, because strict mode influences the number of nullable columns in a table's secondary index. Now the secondary index trees are also checked for corruption after a tablespace is imported. (Bug #35799038)
- InnoDB: Given a table with an FTS index, when its tablespace is discarded the corresponding FTS tables are also dropped. When performing an operation that cleared out these tables, the code incorrectly assumed that the FTS tables must also exist if the FTS indexes were present in the table metadata. (Bug #35343458)
- InnoDB: An uninitialized buffer was being written to a temporary file when checking if the system supports a different sector size for FusionIO. This check is made when innodb_flush_method is set to O_DIRECT or O_DIRECT_NO_FSYNC. (Bug #115229, Bug #36705034)
- InnoDB: Table rebuild operations involving secondary indexes required greater file I/O for InnoDB temporary files as compared with MySQL 8.0.26, which degraded query performance. (Bug #114465, Bug #36444172)
- **InnoDB:** Adding indexes with the parallel index builder was much slower with the Performance Schema enabled due to issues in the Alter_stage class. (Bug #113505, Bug #36163502)
- **InnoDB:** Fixed an innodb.parallel_read_kill related unit test to better account for recent optimizer changes, and fixed an innodb.ddl_kill unit test and its associated macro usage.

Our thanks to Dmitry Lenev and the team at Percona for contributing to these fixes. (Bug #113002, Bug #115416, Bug #35992036, Bug #36764973)

References: See also: Bug #112767.

- Partitioning: ALTER TABLE did not always work correctly with partitioned tables. (Bug #36677952)
- **Group Replication:** Under certain conditions, all secondaries shut down unexpectedly following a brief period of network inactivity on the primary host. (Bug #35642087)

References: See also: Bug #32673639, Bug #34565855.

• Group Replication: Memory aggregated by threads as reported did not account for all memory/ sql/Gtid_set::Interval_chunk freed by other threads from what was allocated, leading to a incorrect ever-increasing consumption of resources by Group Replication thread thread/ group_rpl/THD_applier_module_receiver. Global memory as reported by

performance_schema.memory_summary_global_by_event_name was not affected by this
issue. (Bug #34819861)

- Group Replication: Removed a potential memory leak from xcom\network \xcom_network_provider_native_lib.cc. (Bug #115162, Bug #36673883)
- The server did not always handle connections correctly when running with both the thread pool and audit log plugins. (Bug #37039422)
- Packages for Debian-based systems are now built with WITH_ZLIB=system, and disable MYSQL_MAINTAINER_MODE for debug builds. (Bug #37038213)
- Updated the Kerberos library bundled with commercial builds to version 1.21.3. For more information, see the Kerberos 5 1.21.x Release Notes. (Bug #37034600)
- Removed a heap-use-after-free warning in regexp::Regexp_facade::~Regexp_facade(). (Bug #36867806)
- While dumping tablespaces, mysqldump did not properly escape certain SQL statements in its output. In addition, the dump now encloses the following identifiers within backticks: LOGFILE GROUP, TABLESPACE, and ENGINE. (Bug #36816986)
- Eliminated unnecessary copying in StoreLinkedImmutableStringFromTableBuffers(),
 improving the performance of some hash joins by 5 to 10 percent. (Bug #36805260)
- A previous fix for an issue in sql/sql_executor.cc checked for const_item() rather than const_for_execution(). (Bug #36804785)

References: This issue is a regression of: Bug #34951115.

• When executing an index range scan using IndexRangeScanIterator the record buffer was never set due to the fact that the data type used for IndexRangeScanIterator::m_expected_rows had been changed from double to boolean during refactoring. This unintended change has been reverted.

In addition, when the record buffer was enabled for index range scans a problem arose with multivalued indexes used for covering these scans. (This is possible only when they are used as part of index merge scans, since index merge scans force covering index scans.) The source of the problem was that the implementation of Field_typed_array::key_cmp() needs the value of the generated column for the indexed expression, and this column is not available in the multi-valued index, so the storage engine cannot safely evaluate the end range condition when filling the record buffer for a covering scan. To fix this, we now disable the record buffer when multi-valued indexes are used for covering index range scans. (Bug #36775910)

References: See also: Bug #36341532.

- Certain triggers and stored procedures were not handled correctly. (Bug #36775910)
- Some errors raised when referencing external tables were not handled correctly. (Bug #36758378)
- On Debian, stopped stripping libmysqlclient.a to increase LTO build compatibility. (Bug #36737581)
- The internal function my_convert_internal() sometimes returned a pointer rather than the intended value. (Bug #36684463)

References: This issue is a regression of: Bug #36479091.

- Removed a memory issue in the server component. (Bug #36600205)
- Fixed a memory leak in the mysql client. (Bug #36600203)

- The AES_ENCRYPT() function did not always return a valid result. (Bug #36593265)
- In some cases a join involving a very large number of rows and many tables was not processed correctly. (Bug #36562979)
- Preparation of table value constructors did not track the number of hidden columns added for expressions in an ORDER BY clause. This could cause problems later in the resolution process, possibly leading to an unplanned server exit.

We fix this by counting the number of hidden items in table value constructors in the same way as this is done in other query blocks. (Bug #36560156)

References: This issue is a regression of: Bug #35785452.

- DROP INDEX with the addition of a FULLTEXT index in the same transaction sometimes led to an unplanned server exit. (Bug #36559642)
- Some combinations of optimizer hints did not function correctly. (Bug #36492114)
- Following the fix for a previous issue, a const item that is part of a GROUP BY and is not found in the select list is not added as a hidden item to the field list, but this was not taken into consideration while checking for replacements for expressions involving ROLLUP operations while creating a temporary table. We fix this by checking whether the item that is not found in the field list is a const item, and if so, we use the same item as a replacement. (Bug #36444257)

References: This issue is a regression of: Bug #34951115.

 Removed assertions found during testing of the data masking functions gen_rnd_pan() and mask_ssn(). (Bug #36397869, Bug #36398272)

References: See also: Bug #36398221.

• Fixed an issue relating to FTS comparisons.

Our thanks to Shaohua Wang and the team at Alibaba for the contribution. (Bug #36210202)

- Some CREATE TABLE ... SELECT statements were not always processed correctly. (Bug #36138460)
- RelWithDebInfo builds are no longer compiled with the "-g1" flag, thus increasing the available debugging information and generated file sizes. (Bug #36111629)

References: See also: Bug #33664929.

- Corrected potential misbehavior when the server was run with --skip-grant-tables. (Bug #36043213)
- Fixed a regression in an earlier fix for a problem with references to columns from tables of outer query blocks in the ON condition of a join. (Bug #35854686)

References: This issue is a regression of: Bug #96946, Bug #30350696.

- In certain rare cases, a CREATE TABLE statement involving the creation of a table with foreign keys was not processed correctly. (Bug #35553557, Bug #36350852)
- Fixed an issue relating to SHOW INDEX and generated columns. (Bug #35497623)
- In some cases, table DDL in prepared statements was not handled correctly. (Bug #35221658)
- A left join with an impossible condition as part of an ON clause was not optimized as in MySQL 5.7, so that, in later versions, the query executed more quickly without the impossible condition than with it. An example of such a query, impossible condition included, is SELECT * FROM t1 JOIN t2 ON t1.cl=t2.cl AND 1=2. (Bug #34668756)

- The XML function ExtractValue() did not handle arguments containing accented characters correctly. (Bug #117778, Bug #36236440)
- Changed the SLES 15 target platform from openSUSE 15.5 to 15.6, which also means using GCC 13 instead of GCC 12 and building against the system's OpenSSL 3.x instead of OpenSSL 1.1.x. (Bug #115535, Bug #36934913)
- Fixed a typo in the Debian package description.

Our thanks to Henning Pöttker for the contribution. (Bug #115363, Bug #36749142)

• MYSQL_TYPE_JSON was missing from the description of binary resultsets in the MySQL source code documentation.

Our thanks to Daniël van Eeden for the contribution. (Bug #115360, Bug #36740656)

• With prefer_ordering_index=off, a query with no reference key reverted to scanning and sorting the full table even though it should have been possible to avoid the sort.

Our thanks to Daniel Nichter for the contribution. (Bug #113699, Bug #36213938)

• A query using a greater-than (>) or less-than (<) comparison with a multi-valued index executed much more slowly than the same query using an equality (=) comparison with the same index. (Bug #104897, Bug #33334911)

Changes in MySQL 8.4.2 (2024-07-23, LTS Release)

Bugs Fixed

• InnoDB: In some cases, following the creation of a very large number of tables (8001 or more), the server could not be restarted successfully. (Bug #36808732)

References: This issue is a regression of: Bug #33398681.

- InnoDB: Improved tablespace file scan performance at startup. (Bug #110402, Bug #35200385)
- **Group Replication:** Running a CREATE TABLE ... SELECT statement on a source coming from an asynchronous channel to Group Replication led to errors on the replica. (Bug #36784284)

Changes in MySQL 8.4.1 (2024-07-01, LTS Release)



Important

This release is no longer available for download. It was removed due to a critical issue that could stop the server from restarting following the creation of a very large number of tables (8001 or more). Please upgrade to MySQL 8.4.2 instead.

- Audit Log Notes
- Authentication Notes
- C API Notes
- Compilation Notes
- Component Notes
- Configuration Notes
- Connection Management Notes

- Data Dictionary Notes
- Performance Schema Notes
- Pluggable Authentication
- Thread Pool Notes
- X Plugin Notes
- Bugs Fixed

Audit Log Notes

- Audit log pruning did not function after removing or renaming a file from the audit log. Now pruning continues in such cases, but with a warning printed in the error log stating that it was not possible to delete the missing audit log file. (Bug #35902913)
- MySQL now calls plugin->deinit() with a valid plugin struct as an argument regardless of the plugin's type.

Our thanks to Martin Alderete for the contribution.

Authentication Notes

- Improved log messages to provide clear reasons for Access denied errors when using the authentication_ldap_sasl plugin without proxying. (Bug #35317691)
- MySQL LDAP SASL authentication, when used with the GSSAPI method to access an OpenLDAP server, was rejected with the MySQL server error Plugin authentication_ldap_sasl reported: 'LDAP authentication failed or group retrieval failed: LDAP error: Invalid DN syntax', because OpenLDAP did not recognize the root DN used. (Bug #32631511)

C API Notes

• C API applications stalled while receiving results for server side prepared statements.

Compilation Notes

- macOS: The Xcode version of zlib was removed from the default list of system libraries to use when configuring with -DWITH_SYSTEM_LIBS=ON. (Bug #36537593)
- **Microsoft Windows:** The BUILD_ALL target did not work when compiling on Windows. (Bug #36424619)
- **Microsoft Windows:** Excessive RAM usage led to disk swapping when compiling MySQL on Windows using Ninja. (Bug #36399256)
- Maintainer mode is now disabled when building the debug version of the server for .deb packages. (Bug #36619757)
- Upgraded the bundled googletest and googlemock sources to version 1.14.0. (Bug #36562482)
- Added a missing dependency on GenError. (Bug #36551721)
- When compiling on Fedora 38, grep -E is now used in place of egrep. (Bug #36507549)
- It is now possible on Linux systems to build MySQL using a bundled tcmalloc library that is provided with the source by specifying -DWITH_TCMALLOC=BUNDLED. This is supported on Linux only. (Bug #36313839)

- The bundled tcmalloc() is now used when building MySQL on Enterprise Linux 8. (Bug #114844, Bug #35674008)
- Removed warnings raised in sql/statement/ed_connection.cc when building on Ubuntu 23.04. (Bug #114436, Bug #36428465)
- Linux aarch64 platform binaries are now built using patchelf --page-size=65536 for compatibility with systems using either 4k or 64k for the page size. (Bug #114233, Bug #36393794)

Component Notes

• The values for component options set using the --loose prefix were not read when the component was installed. (Bug #28341329)

Configuration Notes

- Microsoft Windows: On Windows, *MySQL Configurator* was updated to support in-place upgrades as per Upgrade Paths. (Bug #36685422)
- Microsoft Windows: For MSI installations on Windows, MySQL Configurator now automatically upgrades MySQL 8.4 LTS installations without user intervention. (WL #16274)

Connection Management Notes

• The conn_delay/Waiting in connection_control plugin stage was not reset after a delay introduced by the connection control plugin which resulted in incorrect monitoring information. (Bug #35205358)

Data Dictionary Notes

• Attempting to upgrade a MyISAM table containing a mix of regular columns and generated columns from MySQL 5.7 to 8.0 or later led to table corruption. (Bug #105301, Bug #33503328)

Performance Schema Notes

- **Group Replication:** The following tables did not contain data on replication channels which did not have a configured hostname, such as Group Replication recovery channels:
 - REPLICATION_CONNECTION_STATUS
 - REPLICATION_CONNECTION_CONFIGURATION
 - REPLICATION_APPLIER_CONFIGURATION
 - REPLICATION_APPLIER_STATUS
 - REPLICATION_APPLIER_STATUS_BY_COORDINATOR
 - REPLICATION_APPLIER_STATUS_BY_WORKER

As of this release, these tables contain data for partially configured Group Replication channels. (Bug #36018242)

- Under certain conditions, a race condition could result in the amount of RAM used by TABLE_HANDLES increasing to a maximum of 9GB. (Bug #36170903)
- The PROCESSLIST_INFO column of THREADS was not updated when executing a prepared statement.

Thanks to Daniel Lenski and Amazon for the contribution. (Bug #104121, Bug #33057164)

Pluggable Authentication

• The deprecation warning issued when authenticating with the mysql_native_password plugin is
 now issued only once. (Bug #35792948)

Thread Pool Notes

- Connecting to a thread group that had no connection handler threads stalled. We fix this by making sure that connection handler threads terminate only if there is at least one connection thread left. (Bug #36550125)
- Previous refactoring incorrectly removed the connection locking performed when processing incoming connection requests, which ledto a race condition between the thread adding new connections and the connection handler thread processing them. This appeared to cause a situation in which connection requests might be ignored and not processed, so that the connection attempt appeared to hang.

We fix this by taking the connection before processing the queue, and releasing it before waking or creating new threads. (Bug #36548687)

• It was possible to set the thread_pool_longrun_trx_limit system variable to values outside its stated range.

In addition, settings for this variable were not reflected in the output of SHOW VARIABLES or SELECT. (Bug #36347102, Bug #36371145)

• SET PERSIST_ONLY did not work correctly with thread_pool_max_transactions_limit. (Bug #35019884)

X Plugin Notes

- The system variable caching_sha2_password_digest_rounds could not be set to a nondefault value using X Protocol. (Bug #36402455)
- An outdated link to the MySQL documentation in the mysql_function_names unit test source file has been updated.

Our thanks to Minha Jeong for the contribution. (Bug #113500, Bug #36137217)

Bugs Fixed

• InnoDB: MySQL unexpectedly halted on an UPDATE after an ALTER TABLE operation. (Bug #36571091)

References: This issue is a regression of: Bug #35183686.

• InnoDB: The log index size calculation now accounts for column order changes. (Bug #36526369)

References: This issue is a regression of: Bug #35183686.

- **InnoDB:** File system operations performed by InnoDB now consistently fsync the parent directory when performing directory altering tasks. (Bug #36174938)
- **InnoDB:** In debug builds, setting the innodb_interpreter_output debug variable would cause the server to unexpectedly halt. This is now a read-only variable. (Bug #36041032)
- InnoDB: Improved os_innodb_umask handling, and made it read-only. (Bug #35932118)

References: This issue is a regression of: Bug #29472125.

• InnoDB: For tables created with an index on a column that was too wide for the redundant row format (allowed before MySQL 5.7.35), an in-place upgrade silently imported the table but

it was not accessible, which interfered with making backups. Now all operations that involve using the invalid index are rejected with ER_INDEX_CORRUPT until the index is dropped. An ER_IB_INDEX_PART_TOO_LONG error is also reported in the error log. (Bug #35869747)

References: See also: Bug #34826861.

- InnoDB: An InnoDB assertion error referencing an invalid column index was triggered when the column index was valid. (Bug #34800754)
- InnoDB: With an empty XA transaction, shutting the server down after an XA START would cause the server to halt unexpectedly. (Bug #32416819)
- **InnoDB:** Shutting down the replication applier or binlog applier while processing an empty XA transaction caused the system to unexpectedly halt. (Bug #32416819)
- InnoDB: Removed unnecessary heap usage in the Validate_files::check() function.

Our thanks to Huaxiong Song for the contribution. (Bug #115041, Bug #36626203)

• InnoDB: If a partition table was read with innodb_parallel_read_threads=1, read performance greatly decreased from any table after 256 reads. InnoDB behaved as if it reached the maximum capacity of parallel read threads despite not using any.

Our thanks to Ke Yu for the contribution. (Bug #114154, Bug #36347408)

- InnoDB: The result from a spatial index containing a column with a spatial reference identifier (SRID) attribute was empty. In addition, using FORCE INDEX to force a covering index scan on a spatial index led to an assertion. (Bug #112676, Bug #114200, Bug #35894664, Bug #36361834)
- InnoDB: SELECT ... GROUP BY queries were at least twice as slow with the TempTable engine than the Memory engine. (Bug #107700, Bug #34338001)
- **Replication:** If a source contained a stored, generated column populated by a JSON function and binlog_row_image was set to MINIMAL, any subsequent update or deletion on the underlying column failed with the following error:

Invalid JSON text in argument 1 to function json_extract: 'The document is empty.'

The replica attempted to re-evaluate the generated column and failed with that error because the underlying column was unavailable. As of this release, stored, generated columns are not re-evaluated when the underlying columns are unavailable. (Bug #36515172)

• **Replication:** When running GTID-based replication with relay_log_space_limit enabled, a restart of the auto positioning protocol sometimes resulted in an infinite loop, leading to a deadlock in replication. This was because relay_log_space_limit was not honored, not only for transactions whose size exceed this limit, but when the replica could not purge previous logs as well.

To fix this issue, we make the following changes:

- The receiver respects relay_log_space_limit as set by the user, unless a transaction received by the receiver cannot fit into the purged relay log. Before queuing the received transaction, receiver now checks whether scheduling a full transaction is possible. If not, the receiver performs the following actions:
 - · Sets the flag indicating that receiver is waiting
 - Rotates the relay log
 - Waits until it is notified that relay log purge was executed and that the applier has purged all available relay logs; after this, the receiver may queue a transaction without checking the limit again

- Before moving to the next file, the coordinator checks whether the receiver is waiting for available relay log space. If so, the coordinator forcibly purges the applied logs, including the current relay log file. To purge the current relay log file safely, the coordinator must do the following:
 - · Synchronize all of its workers before moving to the next file
 - · Forcibly update group positions, which is necessary to allow current purging of the relay log
 - Update the variable read by the receiver which contains relay log filename to which applier was moved

These operations are allowed because we know that receiver waits at a transaction boundary and rotates the relay log before waiting.

(Bug #36507020)

- **Replication:** Worker jobs now contain information about the relay log file which initiated the transaction, instead of using the default defined by relay_log. (Bug #36395631)
- **Replication:** Handling an incident while transactions were being committed to the binary log caused MySQL to wait indefinitely. (Bug #35671897)
- Group Replication: Removed a memory leak from /xcom/gcs_xcom_networking.cc. (Bug #36532199)
- **Group Replication:** Under certain circumstances, if a primary's host experienced network inactivity of 20 seconds or more, the secondaries could stop unexpectedly. (Bug #36306144)
- **Group Replication:** Under certain circumstances, if garbage collection occurred just before a relay log rotation, it could cause the applier to stop applying new transactions on the secondary members.

This was caused by garbage collection incrementing the relay log's last_committed and sequence_number, creating a gap in the recorded sequence_number after the log rotation. The applier was unaffected if the gap occurred anywhere else in the relay log.

As of this release, only last_committed is updated during garbage collection. (Bug #36280130, Bug #36446250)

• **JSON:** Added missing checks for error handling to NULLIF(), COALESCE(), and the shift (>>) operator. (Bug #113668, Bug #35513196, Bug #36198403)

References: See also: Bug #31358416.

- **MySQL NDB ClusterJ:** Running the ClusterJ test suite resulted in an error message saying a number of threads did not exist. That was due to some wrong handling of threads and connections, which was corrected by this patch. (Bug #36086735)
- On macOS, reinstated the MySQL preference pane's ability to load a custom my.cnf configuration file. This includes loading it before initializing a new data directory. (Bug #36630493)
- On Fedora 40, there were conflicts when installing MySQL 8.4.0 community server on a system with native mysql-server packages present. (Bug #36575524)
- Averages of certain numbers were not always computed correctly. (Bug #36563773)
- The following files in strings contained incorrect license information:
 - mb_wc.h
 - ctype-uca.cc
 - ctype-ucs2.cc

- ctype-utf8.cc
- dtoa.cc
- strxmov.cc
- strxnmov.cc

(Bug #36506181)

- In certain unusual cases, the UpdateXML() function did not process all of its arguments correctly. (Bug #36479091)
- Explaining a query which used FORCE INDEX on a spatial index containing a column with SRID attributes led to an unplanned exit. (Bug #36418426)
- Added the ER_LOG_PARTITION_PREFIX_KEY_NOT_SUPPORTED error definition, which references functionality added in MySQL 8.4.0. (Bug #36350938)
- Updated BuildRequire rules to align with versions now required for CMake and Bison. (Bug #36343254)
- Some prepared statements were not reprepared correctly. (Bug #36267792, Bug #35712413)

References: This issue is a regression of: Bug #34929930.

- When incrementing the reference count for an expression, underlying expressions within this expression are not looked at. While removing an expression, after decrementing the reference count, even the underlying expressions were examined, which led to unintentional deletion of the underlying expressions. This issue manifested in ltem_ref::real_item() as well as in an assert in sql/ item.h. We fix this by not looking at the underlying expression unless the current expression contains the only remaining reference. (Bug #36204344, Bug #36356279)
- Under certain conditions, EXPLAIN FORMAT=JSON FOR CONNECTION sometimes led to an unplanned exit. (Bug #36189820)
- Some CREATE USER statements were not handled correctly. (Bug #36022885)
- In certain cases, a lateral join was not handled correctly. (Bug #35945239)

References: See also: Bug #107700, Bug #34338001. This issue is a regression of: Bug #32644631.

- For a SELECT with ORDER BY and LIMIT, the optimizer first chose a full table scan with a very expensive cost, then performed another check and used the perform_order_index type of path, but this was not reflected by the cost in the optimizer plan. (Bug #35930969)
- Client connections were not alway terminated correctly during shutdown. (Bug #35854919)
- Executing mysgldump on a replica would insert the FLUSH TABLES operation, an operation that writes to the binary log. Now FLUSH LOCAL TABLES is inserted instead to prevent GTID related issues during replication due to these binary log changes.

The workaround was to set the --source-data option to 1 or 2. (Bug #35665076)

References: This issue is a regression of: Bug #33630199.

- All internal ACL bitmask variables are now explicitly 32 bits (uint32_t). (Bug #35507223)
- It was not possible to add a functional index on FIND_IN_SET(). (Bug #35352161)
- Running two concurrent OPTIMIZE TABLE statements on the same table with fulltext indexes and innodb_optimize_fulltext_only enabled sometimes caused the server to exit. (Bug #34929814)

• The gen_range() function as implemented by the (deprecated) data masking plugin did not always return the correct result.

This issue affected the data masking plugin only, and did not affect the data masking component which supersedes it. (Bug #34163992)

• In some circumstances, such when DDL operations were performed on a very large number of tables, the error log was flooded with warnings from background histogram updates; the offending warning was concerning a failure to acquire metadata locks on a table.

To remedy this problem we now throttle messages written to the error log from background histogram update operations, the rate being capped at one message per minute, which should suffice for the user to identify potential problems with background histogram updates. In addition, we downgrade all error events that occur during background histogram updates from errors to warnings. (Bug #114845, Bug #36574298)

• On macOS, the DMG now installs the keyring component instead of installing keyring plugin functionality that was removed in MySQL Server 8.4.0.

The 8.4.0 workaround was to disable the "Keyring Data File" option in the preference pane or to manually remove the two keyring lines from the launchd plist. (Bug #114836, Bug #36577944)

• Fixed an erroneous comment in include/my_command.h.

Our thanks to Sho Nakazono for the contribution. (Bug #114507, Bug #36455468)

- It was possible for a deterministic stored function to return an incorrect result when the function used JOIN ON inside the return statement. If the query needed to be reprepared due to a table metadata caused by, for example, FLUSH TABLES between two executions, the ON clause was sometimes lost. (Bug #114235, Bug #36379879)
- Added the missing mysql-community-libs-compat package for the EL8 and EL9 platforms. (Bug #112949, Bug #35975348)

Changes in MySQL 8.4.0 (2024-04-30, LTS Release)

- Audit Log Notes
- C API Notes
- Character Set Support
- Compilation Notes
- Configuration Notes
- Deprecation and Removal Notes
- Firewall Notes
- INFORMATION_SCHEMA Notes
- Installation Notes
- Keyring Notes
- Optimizer Notes
- Performance Schema Notes
- Server Administration
- Thread Pool Notes

- Functionality Added or Changed
- Bugs Fixed

Audit Log Notes

- Invoking audit_api_message_emit_udf() with arguments of mixed types could lead to an unplanned shutdown of the server. (Bug #36301441)
- Audit log filtering by type, using error as the type, did not work correctly. (Bug #36142157)
- Following an unplanned shutdown and restart, the file that was in use by the server for writing at the time may be broken or otherwise unreadable. The Audit Log plugin log message indicating unreadability of the file was shown as an error; now instead this message is treated as a warning. (Bug #36118809)

C API Notes

- **Important Change:** The following MySQL C API functions, removed in MySQL 8.3, have been reimplemented and restored in MySQL 8.4.0:
 - mysql_kill(): In place of COM_PROCESS_KILL (removed), this function has been
 reimplemented such that it uses mysql_real_query() to execute KILL.
 - mysql_list_fields(): Restored as previously implemented, along with COM_FIELD_LIST.
 - mysql_list_processes(): Reimplemented using mysql_real_query() to execute SHOW PROCESSLIST, in place of COM_PROCESS_INFO (removed).
 - mysql_refresh(): Reimplemented using mysql_real_query() to execute FLUSH statements in place of COM_REFRESH, which was removed in MySQL 8.3.
 - mysql_reload()
 - mysql_shutdown(): Reimplemented using mysql_real_query() to execute a shutdown command rather than COM_SHUTDOWN, removed in MySQL 8.3.
 - mysql_ssl_set()
 - mysql_stmt_bind_param()

The functions just listed are supported for the lifetime of the MySQL 8.4 series. (WL #16221)

• **Microsoft Windows:** Third party DLL files on which MySQL plugins depend are located, when installed, in the same directory as the MySQL executables. The default Windows behavior is to look for dependences in the same directory as the current executable, which is not appropriate for clients using libmysql.dll outside of the installation directory.

We fix this by that changing the default behavior of MySQL clients so that the loader looks for dependencies in the directory of current module (the executable or libmysql.dll). In addition, since libsasl.dll expects to load all its required dependencies from a directory of its won, SASL plugins are now located in a dedicated subdirectory. (Bug #36006295)

Character Set Support

• When the character_set_server system variable was set using SET PERSIST or SET GLOBAL, it did not take effect for new client sessions or for a client establishing a connection to the server after the server was restarted. The only workaround was to set the corresponding command-line option when starting the server.

To fix this, we now make sure that, at the time of server restart, the configuration data is read in the correct order so that the variable setting takes effect as expected. (Bug #35529604)

Compilation Notes

- The libevent library has been removed. (Bug #36357190)
- Added the libcno library. (Bug #36357181)
- Some of the files in extra/libbacktrace contained incorrect licensing information, copyright information, or both. (Bug #36118772)
- Warnings for unused variables are turned into compilation errors when compiling with DMYSQL_MAINTAINER_MODE=1. To avoid this use –DMYSQL_MAINTAINER_MODE=0 to disable such errors. (Bug #113662, Bug #36198423)

Configuration Notes

- **Microsoft Windows:** On Windows, *MySQL Configurator* incorrectly altered the configuration settings after the **Back** and **Next** buttons were used. (Bug #36156577)
- **Microsoft Windows:** On Windows, *MySQL Configurator* no longer opens when removing a MySQL Server that was not configured. (Bug #35709927)
- **Microsoft Windows:** On Windows, MySQL Configurator stopped adding the default_authentication_plugin variable to the generated my.ini file, a variable removed in MySQL Server 8.4. It also removes it when upgrading an installation to MySQL 8.4. Note that the replacement variable authentication_policy is not set by *MySQL Configurator*. (WL #16137)
- **Microsoft Windows:** On Windows, MySQL Configurator no longer allows upgrading from MySQL 5.7 or earlier, when before it allowed the upgrade to execute after stating that it was not officially supported. (WL #16138)

Deprecation and Removal Notes

• Important Change: The deprecated mysql_native_password authentication plugin is now
disabled by default. It can be enabled by starting MySQL with the new --mysql-nativepassword=ON server option, or by adding mysql_native_password=ON to the [mysqld] section
of your MySQL configuration file.

For more information, see Native Pluggable Authentication. (Bug #36337893)

• **Partitioning:** Silent omission of columns with index prefixes as part of a table's partitioning key was deprecated in MySQL 8.0.21, and generated a warning. In this release, the use of any such columns in the proposed partitioning key is now expressly disallowed, and causes the CREATE TABLE or ALTER TABLE statement in which it occurs to be rejected with an error.

For more information, see Column index prefixes not supported for key partitioning, and KEY Partitioning. (WL #16054)

References: See also: Bug #31100205.

• **Replication:** Syntax for a number of features relating to MySQL Replication that was deprecated in previous versions of MySQL has now been removed. These features include aspects of SQL statement syntax as well as several system status variables in the MySQL server. These changes are detailed following.

SQL statements removed. The following SQL statements have been removed (replacements in brackets): START SLAVE (START REPLICA); STOP SLAVE (STOP REPLICA); SHOW SLAVE STATUS (SHOW REPLICA STATUS); SHOW SLAVE HOSTS (SHOW REPLICAS); RESET SLAVE (RESET REPLICA); CHANGE MASTER TO (CHANGE REPLICATION SOURCE TO); RESET MASTER (RESET BINARY LOGS AND GTIDS); SHOW MASTER STATUS (SHOW BINARY LOG STATUS); PURGE MASTER LOGS (PURGE BINARY LOGS); and SHOW MASTER LOGS (SHOW BINARY LOGS).

The statements just listed have also been removed from all MySQL test programs and files, and elsewhere, where used internally.

Statement options removed. The following options formerly supported by CHANGE REPLICATION SOURCE TO and START REPLICA have been removed and are no longer accepted by the server. They are listed here for each of these statements, with their replacements in brackets:

• CHANGE REPLICATION SOURCE TO options removed:

MASTER AUTO POSITION (SOURCE AUTO POSITION), MASTER HOST (SOURCE HOST), MASTER BIND (SOURCE BIND), MASTER USER (SOURCE USER), MASTER PASSWORD (SOURCE PASSWORD), MASTER PORT (SOURCE PORT), MASTER CONNECT RETRY (SOURCE CONNECT RETRY), MASTER RETRY COUNT (SOURCE RETRY COUNT), MASTER DELAY (SOURCE_DELAY), MASTER_SSL (SOURCE_SSL), MASTER_SSL_CA (SOURCE_SSL_CA), MASTER_SSL_CAPATH (SOURCE_SSL_CAPATH), MASTER_SSL_CIPHER (SOURCE_SSL_CIPHER), MASTER SSL CRL (SOURCE SSL CRL), MASTER SSL CRLPATH (SOURCE SSL CRLPATH), MASTER_SSL_KEY (SOURCE_SSL_KEY), MASTER_SSL_VERIFY_SERVER_CERT (SOURCE_SSL_VERIFY_SERVER_CERT), MASTER_TLS_VERSION (SOURCE_TLS_VERSION), MASTER TLS CIPHERSUITES (SOURCE TLS CIPHERSUITES), MASTER SSL CERT (SOURCE SSL CERT), MASTER PUBLIC KEY PATH (SOURCE PUBLIC KEY PATH), GET_MASTER_PUBLIC_KEY (GET_SOURCE_PUBLIC_KEY), MASTER_HEARTBEAT_PERIOD (SOURCE HEARTBEAT PERIOD), MASTER COMPRESSION ALGORITHMS (SOURCE COMPRESSION ALGORITHMS), MASTER ZSTD COMPRESSION LEVEL (SOURCE_ZSTD_COMPRESSION_LEVEL), MASTER_LOG_FILE (SOURCE_LOG_FILE), and MASTER_LOG_POS (SOURCE_LOG_POS).

• START REPLICA options removed: MASTER_LOG_FILE (SOURCE_LOG_FILE) and MASTER_LOG_POS (SOURCE_LOG_POS).

Status variables removed. Also as part of this work, the following system status variables have been removed from the server, and no longer appear in the output of statements such as SHOW STATUS. These variables are listed here, with their replacements in brackets: Com_slave_start (Com_replica_start); Com_slave_stop (Com_replica_stop); Com_show_slave_status (Com_show_replica_status); Com_show_slave_hosts (Com_show_replicas); Com_show_master_status (Com_show_binary_log_status); and Com_change_master (Com_change_replication_source).

See also Com_xxx Variables. (WL #15831, WL #16063, WL #16069, WL #16086, WL #16087, WL #16088, WL #16089, WL #16090)

• **Group Replication:** The group_replication_allow_local_lower_version_join system variable is now deprecated, and setting it raises a warning (ER_WARN_DEPRECATED_SYNTAX_NO_REPLACEMENT).

You should expect this variable to be removed in a future version of MySQL. Since there is no longer any reason to allow incompatible members to join a group, no replacement for this functionality is planned. (WL #16018)

- A number of server options and variables supported in previous versions of MySQL have been removed in this release. Attempting to set any of them in MySQL 8.4 raises an error. These options and variables are listed here:
 - binlog_transaction_dependency_tracking: Deprecated in MySQL 8.0.35 and MySQL 8.2.0.

There are no plans to replace this variable or its functionality, which has been made internal to the server: Now, when multithreaded replicas are in use, the source mysqld uses always

writesets to generate dependency information for the binary log; this has the same effect as setting binlog_transaction_dependency_tracking to WRITESET in previous versions of MySQL.

• group_replication_recovery_complete_at: Deprecated in MySQL 8.0.34.

Beginning with this release, the policy applied during the distributed recovery process is always to mark a new member online only after it has received, certified, and applied all transactions that took place before it joined the group; this is equivalent to setting group_replication_recovery_complete_at to TRANSACTIONS_APPLIED in previous versions of MySQL.

- avoid_temporal_upgrade and show_old_temporals: Both deprecated in MySQL 5.6. Each of these variables no longer had any effect, and has been removed. There are no plans to replace either variable.
- --no-dd-upgrade: Deprecated in MySQL 8.0.16, now removed. Use --upgrade=NONE instead.
- --old and --new: Both deprecated in MySQL 8.0.35 and MySQL 8.2.0, and now removed.
- --language: This option was deprecated in MySQL 5.5, and has now been removed.
- The --ssl, --skip-ssl, and --admin-ssl server options, as well as the have_ssl and have_openssl server system variables, were all deprecated in MySQL 8.0.26, and are all removed in this release. Use --tls-version and --admin-tls-version instead.
- default_authentication_plugin: Deprecated in MySQL 8.0.27, and now removed. Use authentication_policy instead.

You should also be aware that the syntax for setting the authentication_policy variable has changed; see its description in the Manual for more information.

(Bug #36337893, WL #9677, WL #13965, WL #15461, WL #15839, WL #16056, WL #16058, WL #16059, WL #16095)

- In cases where an aliased table was referenced in EXPLAIN output, the table name was qualified with a schema name, which was not necessary and could result in confusion. These schema qualifications have been removed from the output. (Bug #36053664)
- The unused INFORMATION_SCHEMA.TABLESPACES table, deprecated in MySQL 8.0.22, has now been removed.

For InnoDB tables, the Information Schema INNODB_TABLESPACES and INNODB_DATAFILES tables provide tablespace metadata. (WL #14065)

- LOW_PRIORITY used with LOCK TABLES ... WRITE had had no effect since MySQL 5.5, and was deprecated in MySQL 5.6. It is removed in this release; including LOW_PRIORITY in LOCK TABLES now causes a syntax error. (WL #16057)
- Support for use of the AUTO_INCREMENT modifier with FLOAT and DOUBLE columns was deprecated in MySQL 8.0, and is now removed. Attempting to use these together in CREATE TABLE and ALTER TABLE statements now causes an Incorrect column specifier for column error (ER_WRONG_FIELD_SPEC).



Important

Prior to upgrading to this release, you *must* alter any table having a FLOAT ... AUTO_INCREMENT OF DOUBLE ... AUTO_INCREMENT

column that it no longer uses either of these. Otherwise, the table cannot be upgraded.

(WL #13103)

- The mysql_ssl_rsa_setup utility, which was deprecated in MySQL 8.0.34, is removed in this release. For MySQL distributions compiled using OpenSSL, the MySQL server can perform automatic generation of missing SSL and RSA files at startup. For more inforamtion, Creating SSL and RSA Certificates and Keys using MySQL. (WL #16205)
- This release removes support for the ENGINE clause from the following SQL statements:
 - DROP TABLESPACE (all variants)
 - ALTER TABLESPACE ... DROP DATAFILE
 - All other variants of ALTER TABLESPACE, with the two exceptions listed here:
 - 1. ALTER TABLESPACE ... ADD DATAFILE ENGINE={NDB|NDBCLUSTER}
 - 2. ALTER UNDO TABLESPACE ... SET {ACTIVE | INACTIVE} ENGINE=INNODB

Other than in the exceptional cases listed previously, use of the ENGINE clause with ALTER TABLESPACE or DROP TABLESPACE causes the statement to be rejected with an error.

ENGINE clauses for the ALTER TABLESPACE and DROP TABLESPACE statements were deprecated in MySQL 8.0. (WL #16055)

- The SET_USER_ID privilege, deprecated in MySQL 8.2.0, has been removed in this release, and its use in GRANT statements now causes a syntax error. Use the SET_ANY_DEFINER and ALLOW_NONEXISTENT_DEFINER privileges instead. (WL #15875)
- Removed the mysql_upgrade utility, which was deprecated in MySQL 8.0.16. (WL #16096)
- Removed the deprecated <code>mysqlpump</code> utility along with its associated <code>lz4_decompress</code> and <code>zlib_decompress</code> helper utilities. Instead, use <code>mysqldump</code> or the MySQL Shell dump utilities. (WL #16096)
- The following plugins have been removed. They are noted in the list provided here, along with any server system variables, CMake options, and other features associated with them which have also been removed:
 - authentication_fido, authentication_fido_client: Use authentication_webauthn instead; see WebAuthn Pluggable Authentication.

The authentication_fido_rp_id server system variable, mysql client --fido-register-factor option, and the -DWITH_FIDO CMake option have also been removed.

• keyring_file: Use component_keyring_file instead; see Using the
 component_keyring_file File-Based Keyring Component.

The keyring_file_data system variable has also been removed. In addition, the CMake options -DINSTALL_MYSQLKEYRINGDIR and -DWITH_KEYRING_TEST have been removed.

• keyring_encrypted_file: Use component_keyring_encrypted_file instead; see Using the component_keyring_encrypted_file Encrypted File-Based Keyring Component.

The keyring_encrypted_file_data and keyring_encrypted_file_password server system variables have also been removed.

• keyring_oci: Use component_keyring_oci instead; see Using the Oracle Cloud Infrastructure Vault Keyring Component.

The following server system variables have also been removed: keyring_oci_ca_certificate, keyring_oci_compartment, keyring_oci_encryption_endpoint, keyring_oci_key_file, keyring_oci_key_fingerprint, keyring_oci_management_endpoint, keyring_oci_master_key, keyring_oci_secrets_endpoint, keyring_oci_tenancy, keyring_oci_user, keyring_oci_vaults_endpoint, and keyring_oci_virtual_vault.

• openssl_udf: Use the MySQL Enterprise Encryption component instead; see MySQL Enterprise Encryption.

(WL #15937, WL #15938, WL #15939, WL #15941, WL #16140)

- Support for weak encryption ciphers has been removed. This means that, when configuring encrypted connections, MySQL no longer allows specifying any cipher that does not meet the following conditions:
 - Proper TLS version (TLS v1.2 or TLSv1.3, as appropriate)
 - Forward secrecy
 - SHA2 in cipher, certificate, or both
 - AES in GCM or any other AEAD algorithms or modes

This has implications for setting the system variables ssl_cipher, admin_ssl_cipher,
tls_ciphersuites, and admin_tls_ciphersuites. See the descriptions of these variables for their permitted values.

You should be aware that libmysqlclient is not affected in this change, and continues to support ciphers that do not satisfy its conditions so that it can continue to connect to previous versions of MySQL. (WL #15801)

- The use of non-unique or partial keys as foreign keys is deprecated in MySQL. Beginning with this
 release, you must explicitly enable such nonstandard keys in one of the ways listed here:
 - Set restrict_fk_on_non_standard_key (added in this release) to OFF.
 - Start the server with the --skip-restrict-fk-on-non-standard-key option (also new in this release).

The restrict_fk_on_non_standard_key server system variable is ON by default. This means that any attempt to use a nonstandard key as a foreign key in a CREATE TABLE or ALTER TABLE statement is rejected with the error ER_FK_NO_INDEX_PARENT; setting it to ON allows such statements to run, but they raise ER_WARN_DEPRECATED_NON_STANDARD_KEY as a warning.

Upgrades to MySQL 8.4 releases from MySQL 8.0 are supported even if the old database contains one or more foreign keys referring to non-unique or partial keys. As part of the upgrade process, the server prints a list of warning messages with the names of those foreign keys referring to nonstandard keys.

See the description of restrict_fk_on_non_standard_key for more information. (WL #15699)

References: See also: Bug #30615520, Bug #97836.

Firewall Notes

 Following an upgrade, some MySQL Firewall stored procedures were not updated as expected. (Bug #36084930)

- Several enhancements have been made in the stored procedures provided by MySQL Enterprise Firewall. These improvements are listed here:
 - Stored procedures provided by MySQL Enterprise Firewall are now transactional. When an error
 occurs during execution of a firewall stored procedure, an error is reported, and all changes made
 by the stored procedure up to that point in time are rolled back.
 - Firewall stored procedures now avoid performing unnecessary combinations of DELETE plus INSERT statements, as well as those of INSERT IGNORE plus UPDATE operations, making them faster and more efficient.
 - User-based stored procedures and UDFs, previously deprecated, now raise a deprecation warning, such that calling either of sp_set_firewall_mode() or sp_reload_firewall_rules() now generates such a warning. See Firewall Account Profile Stored Procedures, as well as Migrating Account Profiles to Group Profiles, for more information.

(WL #15790)

INFORMATION_SCHEMA Notes

• Fixed a potential race condition in the **PROCESSLIST** table. (Bug #35509371)

Installation Notes

• As part of the installation process, a file in JSON format named mysql_upgrade_history is now created in the server's data directory, or updated if it already exists. Information contained in this file includes the following items, among others:

The day and time of the installation

The MySQL server version installed

The maturity level of the release (LTS or Innovation)

The mysql_upgrade_info file was deprecated in MySQL 8.0.17, and is no longer used; its presence is now checked for, and if this file is found, it is removed as part of the installation process. (WL #16039)

References: See also: Bug #95165, Bug #29702060.

Keyring Notes

• Migration from a keyring component to a keyring plugin is now supported. To perform such a migration, use the --keyring-migration-from-component server option introduced in this release, setting --keyring-migration-source to the name of the source component, and -- keyring-migration-destination the name of the target plugin.

See Key Migration Using a Migration Server, for more information. (WL #16017)

Optimizer Notes

• This release adds support for automatic updates of histograms. When this feature is enabled for a given histogram, the histogram is updated whenever ANALYZE TABLE is run on the parent table. Automatic recalculation of persistent statistics by InnoDB also updates the histogram when automatic updates are enabled.

Automatic histogram updates use the same number of buckets as the histogram was originally specified with, if any.

To enable automatic histogram updates, include the AUTO UPDATE option (added in this release) for the ANALYZE TABLE statement. To disable it, include MANUAL UPDATE instead. MANUAL UPDATE

(no automatic updates) is the default if neither option is specified. When upgrading to this release, existing histograms are treated as though they had been created using MANUAL UPDATE.

For more information, see Histogram Statistics Analysis. See also Configuring Persistent Optimizer Statistics Parameters. (Bug #36053241, WL #15786)

• The multi-range read (MRR) optimization did not perform as well as in previous releases. (Bug #113711, Bug #36220640)

Performance Schema Notes

• User variables assigned decimal values were rounded up in the user_variables_by_thread table. (Bug #35781732)

Server Administration

• **Important Change:** This release adds a privilege which is specific to the use of the FLUSH PRIVILEGES statements. Unlike the existing RELOAD privilege, the new FLUSH_PRIVILEGES privilege applies only to FLUSH PRIVILEGES statements. This privilege is global in scope, and is applicable to users and roles.

The RELOAD privilege continues to be supported in this capacity to provide backwards compatibility; users having this privilege can still execute FLUSH PRIVILEGES statements following an upgrade. As part of upgrading to a MySQL 8.4 release, a check is performed to see whether there are any users having the FLUSH_PRIVILEGES privilege; if there are none, users having the RELOAD privilege are automatically assigned the new privilege as well. (WL #16044)

• Important Change: This release adds a new OPTIMIZE_LOCAL_TABLE privilege. Users must have this privilege to execute OPTIMIZE LOCAL TABLE and OPTIMIZE NO_WRITE_TO_BINLOG TABLE statements.

When upgrading from a previous releases, users already having the <u>SYSTEM_USER</u> privilege are automatically granted the <u>OPTIMIZE_LOCAL_TABLE</u> privilege. (WL #15819)

Thread Pool Notes

• Important Change: Previously, when the limit defined by

thread_pool_max_transactions_limit was reached, new connections or transactions on existing connections sometimes appeared to hang until one or more of the existing transactions were completed. This release introduces a way to mitigate this issue in many cases by setting an upper limit thread_pool_longrun_trx_limit (added in this release) on the length of time during which the number of ongoing transactions is allowed to match the maximum number of thread pool transactions specified by thread_pool_max_transactions_limit; once this limit is reached, the upper limit on the number of transactions is suspended for the thread group.

When the number of long-running transactions decreases appreciably, thread_pool_max_transactions_limit can be (and is) enforced again. See the description of the thread_pool_longrun_trx_limit server system variable for more information about how this is determined. (WL #16132)

• The Performance Schema tp_connections thread pool plugin table contained no entries for connections that were in the admin group. (Bug #36296830)

Functionality Added or Changed

• Important Change; Group Replication: MySQL 8.0 performs special handling for group members whose version is 8.0.17 or earlier. This special handling is removed in the current release.

Users of MySQL 8.0 are encouraged to upgrade all instances to the latest 8.0 release prior to upgrading to MySQL 8.4. (Bug #36314222)

• Important Change; Group Replication: In-place downgrades of servers within groups are supported within the MySQL 8.4 LTS series. For example, a member of a group running MySQL 8.4.2 can be downgraded to MySQL 8.4.0.

Similarly, cross-version group membership is also supported within the 8.4 release series. For example, a server running MySQL 8.4.0 can join a group all of whose members currently run MySQL 8.4.2, as can a server running MySQL 8.4.3.

For more information, see Upgrading Group Replication. (Bug #35918034)

References: See also: Bug #35397276.

- Important Change; Group Replication: The default values of two server system variables relating to Group Replication have changed:
 - The default value of the group_replication_consistency system variable is now BEFORE_ON_PRIMARY_FAILOVER; previously, this was EVENTUAL.
 - The default value of the group_replication_exit_state_action system variable is now OFFLINE_MODE; previously, this was READ_ONLY.

For more information, see the descriptions of the variables listed, as well as Configuring Transaction Consistency Guarantees, and Responses to Failure Detection and Network Partitioning. (WL #15712, WL #15713)

- Important Change; Group Replication: When issued with group_replication_consistency set to BEFORE_ON_PRIMARY_FAILOVER, the MySQL KILL statement now ignores any consistency guarantees, with any interrupted transactions now being rolled back.
- **Important Change:** For platforms on which OpenSSL libraries are bundled, the linked OpenSSL library for MySQL Server has been updated to version 3.0.13. Issues fixed in OpenSSL version 3.0.13 are described at https://openssl-library.org/news/openssl-3.0-notes/. (Bug #36261675)
- Important Change: Upgrading from MySQL 5.7 to MySQL 8.4 is not supported; the code and behavior was updated to reflect this. Upgrade MySQL 5.7 to 8.0 before proceeding to 8.4. (WL #15924)
- InnoDB: Progress messages are now logged periodically during long-running rollbacks as informational note level error messages, initially as ER_IB_LONG_ROLLBACK_FULL (which appends transaction information) followed by successive ER_IB_LONG_ROLLBACK. (WL #15822)
- InnoDB: Changed the default values for the following InnoDB configuration options: innodb_adaptive_hash_index, innodb_buffer_pool_in_core_file, innodb_buffer_pool_instances, innodb_change_buffering, innodb_doublewrite_files, innodb_doublewrite_pages, innodb_flush_method, innodb_io_capacity, innodb_io_capacity_max, innodb_log_buffer_size, innodb_numa_interleave, innodb_page_cleaners, innodb_parallel_read_threads, innodb_purge_threads, innodb_read_io_threads, innodb_use_fdatasync, temptable_max_ram, temptable_max_mmap, and temptable_use_mmap.

The settings affected by the --innodb-dedicated-server startup option have also changed.

For a list of the new default values, see What Is New in MySQL 8.4 since MySQL 8.0. See also Enabling Automatic InnoDB Configuration for a Dedicated MySQL Server. (WL #16179)

- Packaging: Added support for Fedora 40 and Ubuntu 24.04.
- **Replication:** It is now possible to recover the relay log with any incomplete transactions removed. The relay log is now sanitized when the server is started with <u>--relay-log-recovery=OFF</u> (the default). This means that, on startup, the server removes all of the following items:
 - Incomplete transactions

- Relay log files containing incomplete transactions or parts of incomplete transactions only
- · References in the relay log index file to any relay log files removed

For more information, see the description of the relay_log_recovery server system variable.

• **Group Replication:** When a member rejoining a group has transactions to apply on the group_replication_applier channel from previous participation in the group, those transactions are applied when the member rejoins before connections to a donor during distributed recovery.

This backlog of transactions to apply can be monitored using the performance_schema.replication_applier_status_by_worker table, but there was no information about it in the error log, which could lead to the mistaken impression that the server was stalled.

Now in such cases, one of the messages Distributed recovery will wait until the transactions ... contained on the group_replication_applier channel are applied or Distributed recovery finished applying the transactions ... contained on the group_replication_applier channel is also written to the error log, as appropriate. (Bug #36229998)

- **Group Replication:** As of this release, distributed recovery using the clone plugin is permitted between different releases in the same LTS series. (Bug #35992145)
- **Group Replication:** MySQL Group Replication now supports preemptive certification information garbage collection when running in single-primary mode. This feature can be enabled using the group_replication_preemptive_garbage_collection system variable added in this release; when enabled, only the write sets for those transactions that have not yet been committed are kept, which can save time and memory consumption. group_replication_preemptive_garbage_collection_rows_threshold (also introduced in in this release) sets a lower bound on the number of certification rows needed to trigger preemptive garbage collection when the feature is enabled; the default value is 100000.

The value of group_replication_preemptive_garbage_collection can be changed only when Group Replication is not running, and has no effect on a group running in multi-primary mode. When this system variable is enabled, it is not possible to change between multi-primary mode and single-primary mode (see Changing the Group Mode). For help with obtaining information about memory consumed by the garbage collection process, see Monitoring Group Replication Memory Usage with Performance Schema Memory Instrumentation. (WL #15923)

- **Microsoft Windows:** MySQL Windows binary files (.exe and .dll files) now display additional information when their properties are viewed. (Bug #36379291)
- Clone plugin version requirements have been relaxed to allow cloning between different point releases in the same series. In other words, only the major and minor version numbers must match when previously the release number also had to match.

For example, cloning of MySQL 8.4.0 to (a future) MySQL 8.4.14, or from MySQL 8.4.14 to 8.4.0, is now supported.

For more information, see The Clone Plugin. (Bug #36293529, WL #15989)

• When using the iterator-based format for EXPLAIN FORMAT=JSON (that is, when explain_json_format_version is 2), the output now contains a query_type field identifying the type of statement (select, insert, delete, and so on). (Bug #36134568)

Bugs Fixed

- Important Change; Replication: The TRANSACTION_GTID_TAG privilege is now required to set the gtid_executed server system variable. (Bug #36201133)
- **Important Change:** The Robin Hood hashing library has been replaced with unordered_dense. (Bug #36158022)
- InnoDB; Microsoft Windows: Improved redo log performance on Windows by opening redo log files in overlapped mode. (Bug #36154818)

References: This issue is a regression of: Bug #12527.

- InnoDB: The log writer calls functions that temporarily release log.writer_mutex; when innodb_log_writer_threads=OFF, this potentially led to other threads writing to the redo log in between these times. (Bug #36425219)
- **InnoDB:** Some FTS operations on tables with FTS indexes led to inconsistent results. For example, if the server terminated while synchronizing the FTS cache or when synchronization occurred concurrently with another FTS operation.

Our thanks to Yin Peng and the Tencent team for the contribution. (Bug #36347647)

- InnoDB: When creating an index on a table containing data, valgrind occasionally reported reads of uninitialized memory from ddl::Builder::bulk_add_row(). (Bug #36342792)
- InnoDB: On Windows, keeping a file open without a shared write lock and attempting to acquire the fil_shard mutex caused a deadlock with another thread that had acquired the fil_shard mutex and was attempting to access the same file. (Bug #36159317)

References: See also: Bug #32808809.

- InnoDB: Fixed a potential redo log rotation issue that could emit a "Found existing redo log files, but at least one is missing" error during recovery. (Bug #36124625)
- InnoDB: Found and fixed an assertion failure related to full-text indexes. (Bug #35836581)
- InnoDB: Added a log buffer check to the fil_tablespace_redo_* functions for them to better handle corrupt redo logs. (Bug #35676721)
- InnoDB: Improved buffer handling during the tablespace deletion process, a situation that could have potentially caused an assertion failure. (Bug #35676106, Bug #36343647)
- InnoDB: The redo log would potentially not log a column order change with instant DDL, which could cause an incorrect log replay during recovery. (Bug #35183686)
- InnoDB: Results for SHOW ENGINE INNODB STATUS showed an invalid value (NULL) as the name of the first thread in the File IO section. (Bug #113819, Bug #36118112)

References: This issue is a regression of: Bug #34992157.

- InnoDB: With innodb_parallel_read_threads set to a value greater than 1, InnoDB unnecessarily disabled read-ahead heuristics which resulted in stalls when pages were not already in the buffer pool. (Bug #113482, Bug #36142806)
- InnoDB: Importing a tablespace had a hard limit of 128 characters for the imported column names, which did not properly account for variable-length encodings. It's now set to 64 * the maximum length of a multi-byte characters.

Our thanks to Lee Adria for the contribution. (Bug #113208, Bug #36047803)

• **InnoDB:** Running a query that used a unique hash index with the TempTable storage engine could take significantly more time compared to running the query with the MEMORY engine.

Our thanks to xiaoyang chen for the contribution. (Bug #113178, Bug #36037224, Bug #36224958)

- **InnoDB:** The redo log consumer could not advance if capacity was full and another thread was executing USER-related operations such as CREATE USER. This also blocked new connections, which potentially prevented the workaround solution of increasing innodb_redo_log_capacity size. (Bug #112608, Bug #36004840)
- **InnoDB:** In debug builds, there was an assertion failure in InnoDB's background when a transaction it wanted to acquire an MDL lock on was no longer active.

This fix is based on a patch from Genze Wu with Alibaba, thank you for the contribution. (Bug #112424, Bug #35835864)

References: This issue is a regression of: Bug #33700835.

- InnoDB: The MySQL truncate undo operation (purge thread) did not remove the undo_{space_number}_trunc.log file when attempting to truncate the undo tablespace. (Bug #112262, Bug #35784192)
- InnoDB: With innodb_parallel_read_threads set to a value greater than 1, InnoDB would unnecessarily request asynchronous reads which required more synchronization during I/O completion and created a bottleneck due to the limited number of available threads (innodb_read_io_threads) for handling I/O operations. Now this performs synchronous instead of asynchronous reads. (Bug #112137, Bug #35740866)
- **InnoDB:** A trx would unexpectedly halt after encountering an incorrect trx->in_innodb value.

Our thanks to Shaohua Wang for the contribution. (Bug #110652, Bug #35277407)

- InnoDB: Fixed performance issues related to querying the data_lock and data_lock_waits tables when thousands of read-only transactions were present. (Bug #109539, Bug #34951273)
- InnoDB: MySQL no longer ignores the optimizer hint to use a secondary index scan, which instead forced a clustered (parallel) index scan. (Bug #100597, Bug #112767, Bug #31791868, Bug #35952353)
- **Replication:** diagnostics.sql prevented upgrades to MySQL 8.4.0 from earlier versions of MySQL when restoring from data containing old replication terminology such as SHOW SLAVE STATUS. (Bug #36323066)
- **Replication:** Gtid_tagged_log_event encoded the correct value only when the original commit timestamp was equal to the immediate commit timestamp, instead of only when they were different. (Bug #36312880)
- **Replication:** In certain cases, the gtid_next server system variable accepted an invalid value, displayed an invalid value after setting it (even to a legal value), or both. (Bug #36308318)
- **Replication:** The replication receiver thread did not report errors when a replication channel was configured with an unknown network namespace. The receiver thread stopped when such an error occurred but no reason for the halt was shown or logged. (Bug #36054355)
- **Replication:** With binlog_format=ROW and gtid_mode=OFF, deadlocks were sometimes reported among workers contending for the auto-increment lock when applier concurrency was high on the replica. (Bug #35851009)
- **Replication:** In certain cases, the SQL thread terminated with error MY-001755 (ER_MTA_CANT_PARALLEL) when executed with the parallel applier. (Bug #35431274)
- **Replication:** Failure of XA COMMIT of a prepared transaction could result in transaction rollback. (Bug #33650776)

- **Replication:** The replication receiver thread stopped with an error if the replication source server sent a heartbeat event containing a binary log file position that was above the 4GB offset, due to the large size of the binary log file. A new heartbeat event (Heartbeat_log_event_v2, log event type 41) that handles the larger value correctly has been added for use in this situation. (Bug #29913991)
- **Replication:** When the server printed an ER_REPLICA_HEARTBEAT_FAILURE error message, it did not respect the length of the master log file name, leading to it print unrelated data. (Bug #29913928)
- **Group Replication:** Problems arose when members M1 and M2 were in a group, with M1 using u1 as its recovery user and M2 using u2 as its own recovery user, and both users u1 and u2 existing on M1 and M2 with all necessary privileges, and when a new member M3 joined the group using u2 as its recovery user. M3 knew only of user u2, but did not know of user u1, leading START GROUP REPLICATION on M2 to be rejected since M1 was unable to connect to M2. This also generated a new view_id listing the group members as M1 and M2, but M1 nevertheless continued trying to connect to M3, with M1 logging Error in establishing mysql connection and M3 logging Access denied errors for the connection attempts from M1.

By design, XCom stores the last three known configurations, including references to physical connections shared among all past and present configurations. This is done to facilitate quick reconnections by nodes rejoining the group, explicitly or implicitly, and that were already present in any of those configurations.

A side effect of this was that we might keep attempting to connect to a node that was currently not in the group. To solve this problem, we inhibit error logging if the node is not in the current configuration, in order to avoid false negatives which might lead a DBA or an operator to think mistakenly that there is a problem in the system. (Bug #36210988)

References: See also: Bug #32592027.

- Group Replication: Improved handling of GTID sets. (Bug #36093405)
- Group Replication: Two cases were found in which a member exited the group and moved to the ERROR state, but did not honor the action specified by group_replication_exit_state_action; these are listed here:
 - When an error occurred while enabling super_read_only
 - When member join recovery was not possible, due to missing binary logs and clone groups on group members

Example: When the value of group_replication_exit_state_action was OFFLINE_MODE and one of these events took place, offline mode was not enabled as expected. (Bug #36076308)

- Group Replication: After successfully setting a new primary, group_replication_set_as_primary() in some cases waited indefinitely for the operation to complete. (Bug #36059098)
- Group Replication: For errors affecting transactions with AFTER (ER_GRP_RPL_TRX_WAIT_FOR_GROUP_PREPARE_FAILED), the message that was written to the error log referenced a session ID instead of the UUID. (Bug #35953196)
- **Group Replication:** A group running group replication with a primary i1 and two secondaries i2 and i3 started to have intermittent issues because of high memory usage on the primary. The secondaries began reporting the primary as unreachable then reachable again, and the primary began reporting the secondaries as intermittently reachable then reachable as well. Following a period of such instability, the secondaries expelled the original primary (i1) and elected a new one (i2).

Under these conditions, queries against the

performance_schema.replication_group_members table on the former primary (i1) reported i1 as ONLINE and PRIMARY, i2 as ONLINE and SECONDARY, and i3 as ONLINE and

SECONDARY for an extended period of time (12 hours or more) until the mysqld process was restarted on i1.

The problems observed were found to have begun on the original primary (i1) when one of the secondaries was overloaded and began intermittently leaving and joining the group, its connections being dropped and recreated repeatedly on the primary server. During the reconnection process, the primary hung when trying to create the connection, thus blocking the single XCom thread. This was traced to the invocation of SSL_connect() on the XCom communication stack, which changed in MySQL 8.0.27 from asynchronous to synchronous form. When a node was overloaded, it might not respond to the SSL_connect() call, leaving the connecting end blocked indefinitely.

To fix this, we now connect in a way that is non-blocking, and that returns in case of a timeout, leaving the retry attempts to the caller—in this specific case, the XCom thread when trying to reconnect to another node. (Bug #34348094, Bug #36047891)

References: See also: Bug #37587252.

- **JSON:** JOIN and GROUP BY handled some JSON column values differently. (Bug #101048, Bug #31969607)
- **MySQL NDB ClusterJ:** The setLimits() method can now be chained to deletePersistentAll() to limit the number of items to delete. See the description of deletePersistentAll() for details. (Bug #36049906)
- The strings and strings_shared library files declared but did not supply the function mysql::collation::find_by_id(). (Bug #36353447)
- Raised the minimum required version of CMake to build MySQL from 3.5.1 to 3.14.6. (Bug #36338366)

References: See also: Bug #35553331.

- Configuration of the backtrace library was performed too early in the build process, and the library itself was built with an incomplete set of compiler flags, differing in both these respects from the rest of the server. (Bug #36292247)
- SET GLOBAL offline_mode=ON did not always perform correctly when issued under high loads. (Bug #36275182)

References: See also: Bug #36405894.

- Upgraded curl to version 8.6.0. (Bug #36267545)
- Added a new error message for the case when a timeout is detected in net_read_raw_loop()
 rather than in the thread pool code. This includes information about the conditions triggering
 the timeout. This is an error-level message if the timeout occurs earlier than indicated by
 wait_timeout. (Bug #36250895)

References: See also: Bug #34857147.

• mysqldump did not always interpret the server version correctly. (Bug #36248967)

References: See also: Bug #36405879.

- Condition pushdown to a view was rejected with a collation mismatch if the view was created with a different character set than the character set used when querying the view. (Bug #36246859)
- Improved the SQL grammar in sql/sql_yacc.yy by removing four shift-reduce conflicts which were not needed. (Bug #36221823)
- Use of the deprecated exec_program() command has been replaced by execute_process() to provide compatibility with CMake 3.28.1 and later. (Bug #36220656)

- The MLE component was added to the minimal RPM build. (Bug #36210740)
- Some queries using NULLIF() and EXCEPT raised an assertion in set_typelib(). (Bug #36151537)

References: See also: Bug #33045412.

• Certain queries raised an assertion in EstimateDeleteRowsCost(). (Bug #36130806)

References: This issue is a regression of: Bug #35590128.

- A query of the form SELECT 1 FROM t WHERE CAST(a AS UNSIGNED INTEGER) = 1 AND a = (SELECT 1 FROM t) led to an assertion in item_func.cc. (Bug #36128964)
- When selecting two empty strings that were combined with UNION as in SELECT '' AS a UNION SELECT '' AS b, the type of the resulting data was CHAR(0) instead of VARCHAR(0).

We fix this by removing an exception that was made for strings of length 0. (Bug #36112585)

- Upgraded the protobuf library to version 25.1. (Bug #36108397)
- For building Enterprise Linux RPMs, the build scripts now point to a newer strip command (under /opt/rh/gcc-toolset-12), and they now check that the corresponding dwz tool is available. Previously this was only implemented for EL8. (Bug #36090069)
- We now look for gcc-ar and gcc-ranlib when building on Oracle Linux with link-time optimization. (Bug #36089900)
- Use sa_sigaction rather than sh_handler for catching fatal signals, which allows the signal handler to output more information when handling SIGSEGV or SIGFPE signals. (Bug #36082110)
- The MySQL client was unable to authenticate with mysql_native_password to old MySQL Server versions that don't support pluggable authentication, such as MySQL 5.0.15. (Bug #36066161)
- Improved the messages written to the log during a server downgrade. (Bug #36053108)
- Keyring component error logging now supplies more information than previously when the component is unable to initialize. (Bug #36037172)
- Set _ITERATOR_DEBUG_LEVEL to 0 when compiling debug builds on Windows using Clang. (Bug #36032501)
- When performing a rollup on an ENUM or SET column, an assertion was raised in sql/ item_sum.cc during resolution when type information for neither of these types could be found. (Bug #36028294)

References: See also: Bug #33045412.

- When a Common Table Expression (CTE) contained an INTERSECT or EXCEPT set operation, the second use of the same CTE in a subsequent join returned a wrong result. (Bug #36002215)
- Killing a query, while it was evaluating an uncorrelated subquery containing a hash join during optimization, led to an assert in sql/sql_select.cc. (Bug #35991384)
- The server sometimes terminated unexpectedly in response to a specific query. (Bug #35957627)
- A rollup query with a window function such as COUNT() in the select list, which was also part of an ORDER BY, led to an unexpected shutdown of the server. (Bug #35947358)

References: This issue is a regression of: Bug #33069747.

- Improved view and trigger definer handling by view and table DDL. (Bug #35942937)
- The server did not always return metadata to the client correctly for certain queries. (Bug #35904044, Bug #36521376)

- Found and fixed an assertion failure at handler::ha_index_end() in handler.cc. (Bug #35877600)
- For a query such as SELECT DISTINCT t1.x, t2.x FROM t AS t1, t AS t2 WHERE t1.pk = t2.x, where t1.pk = t2.x and pk is the primary key, there is a functional dependency t2.x->t1.x. This means that some candidate plans grouped on {t2.x, t1.x} and others on {t1.x}, which were both valid but yielded different row estimates for two sets of fields, since this did not take functional dependencies into account.

Now we ensure that we perform a single calculation of the number of distinct rows, and use that number for all plan candidates. (Bug #35855573)

- When running queries against a table with a multi-value index, the server sometimes exited unexpectedly, often while executing a complex SELECT query which used this index. (Bug #35789759)
- Improved code in sql/item_subselect.cc. (Bug #35733778, Bug #35738531, Bug #35779012)
- Some aggregations of window functions were not handled correctly. (Bug #35560806)
- CREATE USER IF EXISTS was not always logged correctly. (Bug #35530823)
- The server did not disallow subqueries in partition expressions properly. These are invalid, and should cause a syntax error. (Bug #35476172)
- Events created within stored programs were not always handled correctly. (Bug #35395333, Bug #36402968, Bug #37918920)

References: This issue is a regression of: Bug #17809, Bug #11745618.

- Upgraded the minimum Boost version used to 1.84.0. (Bug #35259498)
- Some RANK() ... OVER() queries raised an assertion in sql/sql_executor.cc. (Bug #35228083)
- When successive ALTER TABLE ... ALGORITHM=COPY statements were issued within 10 seconds of one another, the n_rows value became 0. (Bug #35127747)
- Removed a memory leak observed while running authentication_kerberos under Valgrind. (Bug #34482788, Bug #36570929)
- A query using MAX(column) gave different results before and after an index was added to the column. (Bug #34057432)
- Some queries that used the LEAD() or LAG() window functions on a column of type SET or ENUM hit an assertion during resolution. The same assertion was hit in some queries using the LEAST() or GREATEST() function on a SET or ENUM column. (Bug #33045412)
- When adding a HAVING condition to a temporary table, it is expected that all the fields in the HAVING condition are already replaced with the temporary table fields, but for a query which had an expression involving the internal Item_row type in the HAVING clause, constant expressions were not getting cached, so that the HAVING clause still held references to the fields from the underlying tables. (Bug #30112096)
- In queries that materialized rows in a temporary table before performing hash join or streaming
 aggregation, data was sometimes copied twice from the temporary table to the join buffer or
 aggregation buffer. While this did not cause any wrong results, it led to inefficient use of buffer space
 with a possible negative impact on performance.

This was due to the internal WalkTablesUnderAccessPath() function visiting tables in MATERIALIZE access paths twice: first when it saw the MATERIALIZE access path itself, and then again when it visited the table_path member of the MATERIALIZE access path.

We fix this by not visiting the table when seeing the MATERIALIZE path, and doing so only when seeing the table_path below MATERIALIZE. (Bug #113647, Bug #36190386)

- Updated the URL used for downloading the Boost C++ libraries. (Bug #113576, Bug #36164514)
- In the debug server, an intersection comparing columns of different types sometimes triggered an assert in sql/item.cc. (Bug #113385, Bug #36094867)
- A transform could be semantically invalid when the selected item in the subquery tested for NULL; the left outer join with a grouped derived table might in such cases produce NULL while the original subquery might not. To prevent this from happening, we now bar such subqueries from being transformed. (Bug #113318, Bug #36070542)
- The fix for a previous issue, first addressed in MySQL 8.0.30, was incomplete.

Our thanks to Hao Lu for the contribution. (Bug #113174, Bug #36035044)

References: This issue is a regression of: Bug #110801, Bug #35328028.

• On s390x, we now compile the FMA test with -02 to avoid overoptimization.

Our thanks to Jonathan Albrecht for the contribution. (Bug #113096, Bug #36016140)

• Although s390x is a big-endian platform, the little-endian ICU data directory was used for compiling.

Our thanks to Jonathan Albrecht for the contribution. (Bug #113095, Bug #36016141)

- SET SESSION optimizer_switch = 'hash_set_operations=off' after preparing a statement led to an assertion in sql/sql_select.cc when trying to execute the same prepared statement. (Bug #112919, Bug #35970620)
- The server now uses ER_NO_REFERENCED_ROW_2 or ER_ROW_IS_REFERENCED_2 for foreign key errors whether error details are displayed, or not. In addition, we now display parent and child table details in error messages when the user has the proper grants. (Bug #112589, Bug #35868410)
- Incorrect results were sometimes obtained from a query that used a group by loose index scan. (Bug #112541, Bug #35854362)
- An assertion in sql/sql_derived.cc that checked whether a referenced item in an Item_ref object had consistent outer reference information failed when the reference was of type OUTER_REF. For objects of type Item_outer_ref, dependency information was set for the Item_outer_ref object and the original expression that this reference points to, but an intermediate reference object between the Item_outer_ref and the original expression did not contain this information. (Bug #112478, Bug #35846847)
- An assertion failed in debug builds when inserting data with a zero-length column, such as CHAR(0) or BINARY(0), into a table. Now, a less strict assertion more accurately fails only if it detects that a non-zero number of bytes copied from a source is identical to the target. (Bug #111450, Bug #35507763)
- MySQL did not build correctly using the musl version of libc.

Our thanks to Sam James for the contribution. (Bug #110808, Bug #35330950)

- Using a default string histogram on a TEXT column raised an assertion due to a collation mismatch when comparing histograms bucket values with the string returned by REVERSE(1). (Bug #110527, Bug #35227319)
- A VALUES statement in a correlated lateral or (other) dependent subquery yielded an incorrect result. (Bug #109252, Bug #110076, Bug #34852090, Bug #35087820)