
MySQL NDB Cluster 8.4 Release Notes

Abstract

This document contains release notes for the changes in each release of MySQL NDB Cluster that uses version 8.4 of the [NDB \(NDBCLUSTER\)](#) storage engine.

Each NDB Cluster 8.4 release is based on a mainline MySQL Server release and a particular version of the [NDB](#) storage engine, as shown in the version string returned by executing `SELECT VERSION()` in the `mysql` client, or by executing the `ndb_mgm` client `SHOW` or `STATUS` command; for more information, see [MySQL NDB Cluster 8.4](#).

For general information about features added in NDB Cluster 8.4, see [What is New in MySQL NDB Cluster 8.4](#). For a complete list of all bug fixes and feature changes in MySQL NDB Cluster, please refer to the changelog section for each individual NDB Cluster release.

For additional MySQL 8.4 documentation, see the [MySQL 8.4 Reference Manual](#), which includes an overview of features added in MySQL 8.4 that are not specific to NDB Cluster ([What Is New in MySQL 8.4 since MySQL 8.0](#)), and discussion of upgrade issues that you may encounter for upgrades from MySQL 8.3 to MySQL 8.4 ([Changes in MySQL 8.4](#)). For a complete list of all bug fixes and feature changes made in MySQL 8.4 that are not specific to [NDB](#), see [MySQL 8.4 Release Notes](#).

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: 2025-12-10 (revision: 30746)

Table of Contents

Preface and Legal Notices	1
Changes in MySQL NDB Cluster 8.4.7 (2025-10-24, LTS Release)	3
Changes in MySQL NDB Cluster 8.4.6 (2025-07-23, LTS Release)	4
Changes in MySQL NDB Cluster 8.4.5 (2025-04-16, LTS Release)	5
Changes in MySQL NDB Cluster 8.4.4 (2025-01-22, LTS Release)	8
Changes in MySQL NDB Cluster 8.4.3 (2024-10-16, LTS Release)	11
Changes in MySQL NDB Cluster 8.4.2 (2024-07-23, LTS Release)	12
Changes in MySQL NDB Cluster 8.4.1 (2024-07-02, LTS Release)	13
Changes in MySQL NDB Cluster 8.4.0 (2024-04-30, LTS Release)	16
Release Series Changelogs: MySQL NDB Cluster 8.4	20
Changes in MySQL NDB Cluster 8.4.1 (2024-07-02, LTS Release)	20
Changes in MySQL NDB Cluster 8.4.0 (2024-04-30, LTS Release)	23
Index	26

Preface and Legal Notices

This document contains release notes for the changes in each release of MySQL NDB Cluster that uses version 8.4 of the [NDB](#) storage engine.

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 error-free. 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 NDB Cluster 8.4.7 (2025-10-24, LTS Release)

MySQL NDB Cluster 8.4.7 is a new LTS release of NDB 8.4, based on MySQL Server 8.4 and including features in version 8.4 of the [NDB](#) storage engine, as well as fixing recently discovered bugs in previous NDB Cluster releases.

Obtaining MySQL NDB Cluster 8.4. NDB Cluster 8.4 source code and binaries can be obtained from <https://dev.mysql.com/downloads/cluster/>.

For an overview of major changes made in NDB Cluster 8.4, see [What is New in MySQL NDB Cluster 8.4](#).

This release also incorporates all bug fixes and changes made in previous NDB Cluster releases, as well as all bug fixes and feature changes which were added in mainline MySQL 8.4 through MySQL 8.4.7 (see [Changes in MySQL 8.4.7 \(2025-10-21, LTS Release\)](#)).



Note

These release notes were created with the assistance of MySQL HeatWave GenAI.

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

Functionality Added or Changed

- MySQL NDB Cluster now generates log entries for redo logging issues, including [redo log buffer exhaustion](#), [redo log space exhaustion](#), and [file change problems](#), which

can cause transactions to be aborted, queued, or delayed. A secondary overload control mechanism monitors the rate at which the redo log part's Redo buffer is flushed to disk and estimates the time required to complete writing all data in the part's Redo buffer. (Bug #37903091)

- [LogLevelBackup](#) and [LogLevelSchema](#) Data node options were added in this release. (Bug #28004136)

Bugs Fixed

- If a data node restarted during a backup, the remaining data nodes could hang and the backup process stopped. The backup was marked as failed, and other nodes did not terminate their backup process, preventing further backups from being initiated. (Bug #38316252)
- It was not possible to restore backups with [BackupID](#) values greater than 2147483647, with the [ndb_restore](#) tool. Errors were returned similar to the following:

```
ndb_restore: [Warning] option 'backupid': signed value
3000000000 adjusted to 2147483647. Failed to find backup
2147483647.
```

(Bug #38260769)

- Concurrent execution of [TRUNCATE TABLE](#) statements and starting an Ndb backup could cause the statements to hang indefinitely, affecting cluster availability. Errors were returned similar to the following:

```
ERROR 1296 (HY000): Got error 761 'Unable to drop table as
backup is in progress' from NDBCLUSTER
```

(Bug #37486661)

- Setting log level options for data nodes in MySQL NDB Cluster, such as [LogLevelCheckpoint](#), [LogLevelStatistic](#), [LogLevelInfo](#), [LogLevelCongestion](#), [LogLevelError](#), and [LogLevelConnection](#) to a higher value does not provide additional log information beyond their default values. Errors were returned similar to the following:

```
Ndb kernel thread 0 is stuck in: Job Handling elapsed=100,
Watchdog: Warning overslept
```

(Bug #17847063)

- The MySQL NDB Cluster's example code for [MGMAPI_LOGEVENT](#) was intended to compile as plain C, but had a .cpp extension and included a C++ header. The example file has been renamed to [main.c](#) and the C++ header has been removed. (Bug #12678874)

Changes in MySQL NDB Cluster 8.4.6 (2025-07-23, LTS Release)

MySQL NDB Cluster 8.4.6 is a new LTS release of NDB 8.4, based on MySQL Server 8.4 and including features in version 8.4 of the [NDB](#) storage engine, as well as fixing recently discovered bugs in previous NDB Cluster releases.

Obtaining MySQL NDB Cluster 8.4. NDB Cluster 8.4 source code and binaries can be obtained from <https://dev.mysql.com/downloads/cluster/>.

For an overview of major changes made in NDB Cluster 8.4, see [What is New in MySQL NDB Cluster 8.4](#).

This release also incorporates all bug fixes and changes made in previous NDB Cluster releases, as well as all bug fixes and feature changes which were added in mainline MySQL 8.4 through MySQL 8.4.6 (see [Changes in MySQL 8.4.6 \(2025-07-22, LTS Release\)](#)).

**Note**

These release notes were created with the assistance of MySQL HeatWave GenAI.

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

Functionality Added or Changed

- Timestamps in [NDB](#) node logs can now be printed with microsecond resolution. Data nodes can enable this feature using the data node `--ndb-log-timestamps=UTC` option; management nodes can also do so using the `ndb_mgmd` option `--ndb-log-timestamps=UTC`. For backwards compatible behavior, you can set this option explicitly to `LEGACY`, which uses the system time zone and resolution in seconds, as in previous releases. In NDB 8.4, this is the default.

For SQL nodes, use the `--log-timestamps` option; be aware that this `mysqld` option does not support `LEGACY` as a value.

See [NDB Cluster Log Messages](#), for more information. (Bug #37924338)

Bugs Fixed

- **NDB Client Programs:** `ndb_restore`, when applying the log, allowed an infinite number of retries due to temporary errors, rather than limiting these to 11 as expected. (Bug #37883579)

References: This issue is a regression of: Bug #31546136.

- **NDB Client Programs:** When restoring from backup with `DefaultOperationRedoProblemAction=ABORT`, an error in data node handling of a redo log part overload condition resulted in an incorrect error code (626 in this case) being sent back to `ndb_restore`, which caused `ndb_restore` to exit prematurely since it did not expect such an error. (Bug #37687485)

References: This issue is a regression of: Bug #13219, Bug #13930, Bug #13980.

- Improved password handling for file system encryption.

Our thanks to Axel Svensson for the contribution. (Bug #37909595)

- `CREATE TABLESPACE`, when the specified logfile group did not exist, was rejected with error `723 No such table existed`, which did not correctly identify the source of the problem with the SQL statement. Now in such cases, the server issues Error `789 Logfile group not found`. (Bug #37802388)
- Warning 1296 `The temporary named table ... already exists` showed the table and schema names in the wrong order.

Our thanks to Axel Svensson for the contribution. (Bug #117918, Bug #37807409)

Changes in MySQL NDB Cluster 8.4.5 (2025-04-16, LTS Release)

MySQL NDB Cluster 8.4.5 is a new LTS release of NDB 8.4, based on MySQL Server 8.4 and including features in version 8.4 of the [NDB](#) storage engine, as well as fixing recently discovered bugs in previous NDB Cluster releases.

Obtaining MySQL NDB Cluster 8.4. NDB Cluster 8.4 source code and binaries can be obtained from <https://dev.mysql.com/downloads/cluster/>.

For an overview of major changes made in NDB Cluster 8.4, see [What is New in MySQL NDB Cluster 8.4](#).

This release also incorporates all bug fixes and changes made in previous NDB Cluster releases, as well as all bug fixes and feature changes which were added in mainline MySQL 8.4 through MySQL 8.4.5 (see [Changes in MySQL 8.4.5 \(2025-04-15, LTS Release\)](#)).

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

Compilation Notes

- The ability to build the source without NDB using the internal script `storage/ndb/compile-cluster` was adversely affected by work done in NDB 8.0.31 making the `ndbcluster` plugin part of its default build. (Bug #117215, Bug #37484376)

Functionality Added or Changed

- Added the `Ndb_schema_participant_count` status variable. This variable provides the count of MySQL servers which are currently participating in NDB Cluster schema change distribution. (Bug #37529202)

Bugs Fixed

- **NDB Replication:** Timestamps written to the binary log included fractional parts (microseconds) although the query being logged did not use any high-precision functionality. This problem was caused by not resetting the state indicating that fractional parts had been used.

We fix this by ensuring that the indicator for the use of microseconds in a given query after having run it is always reset. This avoids the possibility of a later query writing a timestamp which includes microseconds to the binary log when the query as executed did not use microseconds. (Bug #37112446)
- **ndbinfo Information Database:** Certain queries against `ndbinfo` tables were not handled correctly. (Bug #37372650)
- **NDB Client Programs:** With a data node in an unstarted state, such as immediately after executing `node_id RESTART -n` in the `ndb_mgm` client, issuing `ALL REPORT BACKUPSTATUS` in the client subsequently led to an unplanned shutdown of the cluster. (Bug #37505513)
- **MySQL NDB ClusterJ:** The ClusterJ log file only reported the configured, requested node ID for a cluster connection (which was often zero). With this fix, after a connection has been established, ClusterJ reports the actual assigned node ID in the log. (Bug #37556172)
- **MySQL NDB ClusterJ:** A potential circular reference from `NdbRecordSmartValueHandlerImpl` that can cause delays in garbage collection has been removed. (Bug #37361267)
- **MySQL NDB ClusterJ:** Setting the connection property `com.mysql.clusterj.byte.buffer.pool.sizes` to "512, 51200" caused ClusterJ application to fail with a fatal exception thrown by `java.nio.ByteBuffer`. (Bug #37188154)
- **MySQL NDB ClusterJ:** When using a debug build of ClusterJ to run any tests in the testsuite, it exited with the error "1 thread(s) did not exit." (Bug #36383937)
- **MySQL NDB ClusterJ:** Running a ClusterJ application with Java 10 resulted in `java.lang.ClassNotFoundException`, because the class `java.internal.ref.Cleaner` is not available in Java 10. With this fix, the `java.lang.ref.Cleaner` class is used instead for resource cleanup. (Bug #29931569)
- The bundled `libxml2` library has been upgraded to version 2.9.13. (Bug #37806165)

- API node failure is detected by one or more data nodes; data nodes detecting API node failure inform all other data nodes of the failure, eventually triggering API node failure handling on each data node.

Each data node handles API node failure independently; once all internal blocks have completed cleanup, the API node failure is considered handled, and, after a timed delay, the `QMGR` block allows the failed API node's node ID to be used for new connections.

`QMGR` monitors API node failure handling, periodically generating warning logs for API node failure handling that has not completed (approximately every 30 seconds). These logs indicate which blocks have yet to complete failure handling.

This enhancement improves logging in handling stalls particularly with regard to the `DBTC` block, which must roll back or commit and complete the API node's transactions, and release the associated `COMMIT` and `ACK` markers. In addition, the time to wait for API node failure handling is now configurable as the `ApiFailureHandlingTimeout` data node configuration parameter; after this number of seconds, handling is escalated to a data node restart. (Bug #37524092)

References: See also: Bug #37469364.

- When a data node hangs during shutdown reasons for this may include: I/O problems on the node, in which case the thread shutting down hangs while operating on error and trace files; or an error in the shutdown logic, where the thread shutting down raises a Unix signal, and causes a deadlock. When such issues occur, users might observe watchdog warnings in the logs, referring to the last signal processed; this could be misleading in cases where there was actually a (different) preceding cause which had triggered the shutdown.

To help pinpoint the origin of such problems if they occur, we have made the following improvements:

- Added a new watchdog state `shutting down`. This is set early enough in the error handling process that it causes all watchdog logging of shutdown stalls to attribute the delay to a shutdown delay (correctly) rather than problem in execution.
- We have also modified the watchdog mechanism to be aware of shutdown states, and use a more direct path—which is less likely to stall—to force the data node process to stop when needed.

(Bug #37518267)

- When restoring `NDB` tables from backup, it is now possible for `mysqld` to open such tables even if their indexes are not yet available. (Bug #37516858)
- Signal dump code run when handling an unplanned node shutdown sometimes exited unexpectedly when speculatively reading section IDs which might not be present. (Bug #37512526)
- The `LQH_TRANSCONF` signal printer did not validate its input length correctly, which could lead the node process to exit. (Bug #37512477)
- When restoring stored grants (using `ndb_restore --include-stored-grants`) from an `NDB` backup following an initial restart of the data nodes, the `ndb_sql_metadata` table was neither created nor restored. (Bug #37492169)
- Nothing was written to the cluster log to indicate when `PURGE BINARY LOGS` had finished waiting for the purge to complete. (Bug #37489870)
- `WITH_NDB_TLS_SEARCH_PATH` was not set when compiling NDB Cluster using `WITHOUT_SERVER`. (Bug #37398657)
- Now, when `ndb_metadata_check` is enabled, we synchronize both schema and tables in the same interval. (Bug #37382551)
- This fix addresses the following two issues:

1. When a resend could not be started due to a gap created by an out of buffer error in the event stream, forwarding event data for the bucket was not initiated. We fix this by ensure that the takeover process has been initiated before exiting the resend code.
2. An out of buffer error which occurred during an ongoing resend was not handled. In this case, we now interrupt the resend when such an error is raised.

(Bug #37349305)

- On certain rare occasions, when concurrent calls were made to `release()` and `get()`, instances of `ndb_schema_object` were doubly freed. (Bug #35793818)
- If an out-of-buffer-release (OOBR) process took an excessive amount of time, the reset was performed prematurely, before all buffers were released, thus interfering with concurrent seizing of new pages, beginning new `out_of_buffer` handling, or starting a resend.

We solve this issue by ensuring that resumption of event buffering takes place only after the OOBR process has completed for all buckets. (Bug #20648778)

Changes in MySQL NDB Cluster 8.4.4 (2025-01-22, LTS Release)

MySQL NDB Cluster 8.4.4 is a new LTS release of NDB 8.4, based on MySQL Server 8.4 and including features in version 8.4 of the `NDB` storage engine, as well as fixing recently discovered bugs in previous NDB Cluster releases.

Obtaining MySQL NDB Cluster 8.4. NDB Cluster 8.4 source code and binaries can be obtained from <https://dev.mysql.com/downloads/cluster/>.

For an overview of major changes made in NDB Cluster 8.4, see [What is New in MySQL NDB Cluster 8.4](#).

This release also incorporates all bug fixes and changes made in previous NDB Cluster releases, as well as all bug fixes and feature changes which were added in mainline MySQL 8.4 through MySQL 8.4.4 (see [Changes in MySQL 8.4.4 \(2025-01-21, LTS Release\)](#)).

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

Compilation Notes

- **macOS:** A `uint64_t` value used with `%zu` caused a `[-Wformat]` compiler warning on MacOS. (Bug #37174692)
- Removed a warning in `storage/ndb/src/common/util/cstrbuf.cpp`. (Bug #37049014)

Bugs Fixed

- **Microsoft Windows:** Successive iterations of the sequence `ndb_sign_keys --create-key` followed by `ndb_sign_keys --promote` were unsuccessful on Windows. (Bug #36951132)
- **NDB Disk Data:** `mysqld` did not use a disk scan for `NDB` tables with 256 disk columns or more. (Bug #37201922)
- **NDB Replication:** The replication applier normally retries temporary errors occurring while applying transactions. Such retry logic is not performed for transactions containing row events where the `STMT_END_F` flag is missing; instead, the statement is committed in an additional step while applying the subsequent `COMMIT` query event when there are still locked tables. Problems arose when committing this statement, because temporary errors were not handled properly. Replica skip error functionality was also affected in that it attempted to skip only the error that occurred when a transaction was committed a second time.

The binary log contains an epoch transaction with writes from multiple server IDs on the source. The replica then uses `IGNORE_SERVER_IDS` (<last_server_id_in_binlog>) to cause the `STMT_END_F` to be filtered away, thus committing the statement from the COMMIT query log event on the applier. Holding a lock on one of the rows to be updated by the applier triggered error handling, which caused replication to stop with an error, with no retries being performed.

We now handle such errors, logging all messages in diagnostics areas (as is already done for row log events) and then retrying the transaction. (Bug #37331118)

- **NDB Replication:** When a MySQL server performing binary logging connects to an NDB Cluster, it checks for existing binary logs; if it finds any, it writes an `Incident` event to a log file of its own so that any downstream replicas can detect the potential for lost events. Problems arose under some circumstances because it was possible for the timestamps of events logged in this file to be out of order; the `Incident` event was written following other events but had a smaller timestamp than these preceding events. We fix this issue by ensuring that a fresh timestamp is used prior to writing an incident to the binary log on startup rather than one which may have been obtained and held for some time previously. (Bug #37228735)
- **NDB Cluster APIs:** The `Ndb_cluster_connection` destructor calls `g_eventLogger::stopAsync()` in order to release the buffers used by the asynchronous logging mechanism as well as to stop the threads responsible for this logging. When the `g_eventLogger` object was deleted before the `Ndb_cluster_connection` destructor was called, the application terminated after trying to use a method on a null object. This could happen in either of two ways:
 - An API program deleted the logger object before deleting the `Ndb_cluster_connection`.
 - `ndb_end()` was called before the `Ndb_cluster_connection` was deleted.

We solve this issue by skipping the call to `stopAsync()` in the `Ndb_cluster_connection` destructor when `g_eventLogger` is `NULL`. This fix also adds a warning to inform API users that deleting `g_eventLogger` before calling the `Ndb_cluster_connection` destructor is incorrect usage.

For more information, see [API Initialization and Cleanup](#). (Bug #37300558)

- **NDB Cluster APIs:** Removed known causes of API node versus data node state misalignments, and improved the handling of state misalignments when detected. In one such case, separate handling of scan errors in the `NDB` kernel and those originating in API programs led to cleanup not being performed after some scans. Handling of `DBTC` and API state alignment errors has been improved by this set of fixes, as well as scan protocol timeout handling in `DBSPJ`; now, when such misalignments in state are detected, the involved API nodes are disconnected rather than the data node detecting it being forced to shut down. (Bug #20430083, Bug #22782511, Bug #23528433, Bug #28505289, Bug #36273474, Bug #36395384, Bug #36838756, Bug #37022773, Bug #37022901, Bug #37023549)

References: See also: Bug #22782511, Bug #23528433, Bug #36273474, Bug #36395384, Bug #36838756.

- **ndbinfo Information Database:** At table create and drop time, access of `ndbinfo` tables such as `operations_per_fragment` and `memory_per_fragment` sometimes examined data which was not valid.

To fix this, during scans of these `ndbinfo` tables, we ignore any fragments from tables in transient states at such times due to being created or dropped. (Bug #37140331)

- Work done previously to support opening `NDB` tables with missing indexes was intended to allow the features of the MySQL server to be used to solve problems in cases where indexes cannot be rebuilt due to unmet constraints. With missing indexes, some of the SQL handler functionality is unavailable—for example, the use of indexes to select rows for modification efficiently, or to identify duplicates when processing modifications, or to push joins relying on indexes. This could lead to the unplanned shutdown of an NDB Cluster SQL node.

In such cases, the server now simply returns an error. (Bug #37299071)

- Recent refactoring of the transporter layer added the reporting of the presence of socket shutdown errors, but not their nature. This led to confusion in the common case where a socket shutdown is requested, but the socket is already closed by the peer. To avoid such confusion, this logging has been removed. (Bug #37243135)

References: This issue is a regression of: Bug #35750771.

- It was not possible to create an [NDB](#) table with 256 or more [BLOB](#) columns when also specifying a reduced inline size, as in the following SQL statement:

```
CREATE TABLE t1 (
  pk INT PRIMARY KEY,
  b1 BLOB COMMENT 'NDB_COLUMN=BLOB_INLINE_SIZE=100',
  b2 BLOB COMMENT 'NDB_COLUMN=BLOB_INLINE_SIZE=100',
  ...
  b256 BLOB COMMENT 'NDB_COLUMN=BLOB_INLINE_SIZE=100'
) ENGINE=NDBCLUSTER;
```

(Bug #37201818)

- In some cases, the occurrence of node failures during shutdown led to the cluster becoming unrecoverable without manual intervention.

We fix this by modifying global checkpoint ID (GCI) information propagation ([CopyGCI](#) mechanism) to reject propagation of any set of GCI information which does not describe the ability to recover the cluster automatically as part of a system restart. (Bug #37163647)

References: See also: Bug #37162636.

- In some cases, node failures during an otherwise graceful shutdown could lead to a cluster becoming unrecoverable without manual intervention. This fix modifies the generic GCI info propagation mechanism ([CopyGCI](#)) to reject propagating any set of GCI information which does not describe the ability to recover a cluster automatically. (Bug #37162636)
- Improved variable names used in [start_resend\(\)](#), and enhanced related debug messages to users and developers with additional information. (Bug #37157987)
- In certain cases, a [COPY_FRAGREQ](#) signal did not honor a fragment scan lock. (Bug #37125935)
- In cases where [NDB](#) experienced an API protocol timeout when attempting to close a scan operation, it considered the [DBTC ApiConnectRecord](#) involved to be lost for further use, at least until the API disconnected and API failure handling within [DBTC](#) reclaimed the record.

This has been improved by having the API send a [TCRELEASEREQ](#) signal to [DBTC](#) in such cases, performing API failure handling for a single [ApiConnectRecord](#) within [DBTC](#). (Bug #37023661)

References: See also: Bug #36273474, Bug #36395384, Bug #37022773, Bug #37022901, Bug #37023549.

- For tables using the [NDB](#) storage engine, the column comment option [BLOB_INLINE_SIZE](#) was silently ignored for [TINYBLOB](#) columns, and (silently) defaulted to the hard-coded 256 byte value regardless of the size provided; this was misleading to users.

To fix this problem, we now specifically disallow [BLOB_INLINE_SIZE](#) on [TINYBLOB](#) columns altogether, and [NDB](#) now prints a warning saying that the column size is defaulting to 256 bytes. (Bug #36725332)

- Testing revealed that a fix for a previous issue which added a check of the [ApiConnectRecord](#) failure number against the system's current failure number did not initialize the [ApiConnectRecord](#) failure number in all cases. (Bug #36155195)

References: This issue is a regression of: Bug #36028828.

- `ndb_config` did not always handle very long file paths correctly.

Our thanks to Dirkjan Bussink for the contribution. (Bug #116748, Bug #37310680)

- Errors of unknown provenance were logged while assigning node IDs during cluster synchronization, leading to user doubt and concern. Logging of the data node `QMGR` block and the `ndb_mgmd` process relating to node ID allocation issues has therefore been improved, to supply more and better information about what is being reported in such cases. (Bug #116351, Bug #37189356)
- A multi-range scan sometimes lost its fragment lock for the second and subsequent ranges of the scan. (Bug #111932, Bug #35660890)

Changes in MySQL NDB Cluster 8.4.3 (2024-10-16, LTS Release)

MySQL NDB Cluster 8.4.3 is a new LTS release of NDB 8.4, based on MySQL Server 8.4 and including features in version 8.4 of the `NDB` storage engine, as well as fixing recently discovered bugs in previous NDB Cluster releases.

Obtaining MySQL NDB Cluster 8.4. NDB Cluster 8.4 source code and binaries can be obtained from <https://dev.mysql.com/downloads/cluster/>.

For an overview of major changes made in NDB Cluster 8.4, see [What is New in MySQL NDB Cluster 8.4](#).

This release also incorporates all bug fixes and changes made in previous NDB Cluster releases, as well as all bug fixes and feature changes which were added in mainline MySQL 8.4 through MySQL 8.4.3 (see [Changes in MySQL 8.4.3 \(2024-10-15, LTS Release\)](#)).

- [Deprecation and Removal Notes](#)
- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

Deprecation and Removal Notes

- **NDB Client Programs:** The `ndb_size.pl` utility is now deprecated and is no longer supported. You can expect it to be removed from a future version of the NDB Cluster distribution; for this reason, you should now modify any applications which depend on it accordingly. (WL #16456)
- Use of an `Ndb.cfg` file for setting the connection string for an NDB process was not well documented or supported. With this release, this file is now formally deprecated, and you should expect support for it to be removed in a future release of MySQL Cluster. (WL #15765)

Functionality Added or Changed

- The `ndbcluster` plugin subscribes to all changes that occur in `NDB` and writes them epoch by epoch to the binary log. Each epoch received from `NDB` consists of a large number of changes, all of which are written to the binary log transaction cache before flushing them to the binary log. Previously, it was possible to configure the cache size for all threads, which often led to improper resource allocation for a MySQL Server used for writing a binary log of changes for `NDB`.

To enable dimensioning and configuring the system properly, we introduce a new system variable `ndb_log_cache_size` which makes it possible to set the size of the transaction cache used by the `NDB` binary log injector, so that this size can be set separately for writing the binary log for NDB transactions and (using `binlog_cache_size`) for writing other transactions whose sizes are likely to be smaller. (Bug #36694848)

Bugs Fixed

- **NDB Cluster APIs:** Using `NdbRecord` and `OO_SETVALUE` from the NDB API to write the value of a `Varchar`, `Varbinary`, `Longvarchar`, or `Longvarbinary` column failed with error 829. (Bug #36989337)
- **MySQL NDB ClusterJ:** References to ClusterJPA and OpenJPA have been removed from the comments in the packaging files, as JPA code was already removed from ClusterJ some time ago. (Bug #36725675)
- **MySQL NDB ClusterJ:** `ReconnectTest` in the ClusterJ test suite failed sometimes due to a race condition. The test has been rewritten with proper synchronization. (Bug #28550140)
- Removed node management code from `TRIX` that was not actually used. (Bug #37006547)
- Submitting concurrent shutdown commands for individual nodes using `ndb_mgm SHUTDOWN node_id` or the MGM API sometimes had one or both of the following adverse results:
 - Cluster failure when all nodes in the same node group were stopped
 - Inability to recover when all nodes in the same node group were stopped, and the cluster had more than one node group

This was due to the fact that the (planned) shutdown of a single node assumed that only one such shutdown occurred at a time, but did not actually check this limitation.

We fix this so that concurrent single-node shutdown requests are serialized across the cluster, and any which would cause a cluster outage are rejected. (Bug #36943756)

References: See also: Bug #36839995.

- Shutdown of a data node late in a schema transaction updating index statistics caused the president node to shut down as well. (Bug #36886242)

References: See also: Bug #36877952.

- It was possible for duplicate events to be sent to user applications when a data node was shut down. (Bug #36750146)
- `BLOB_INLINE_SIZE=0` set within a column comment was not honored, and the default for the blob type was used instead (such as 256 bytes for `BLOB`).

See [NDB_COLUMN Options](#), for more information. (Bug #36724336)

- Issues arose when an attempt was made to use a SHM transporter's wakeup socket before it was ready, due in part to error-handling when setting up the SHM transporter, which did not close the socket correctly prior to making another attempt at setup. (Bug #36568752, Bug #36623058)
- An error in a `my.cnf` file could cause the management node to shut down unexpectedly. (Bug #36508565)
- A race condition sometimes occurred between the watchdog thread and the signal execution thread trying to start node failure handling in parallel. (Bug #35728261)

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

MySQL NDB Cluster 8.4.2 is a new LTS release of NDB 8.4, based on MySQL Server 8.4 and including features in version 8.4 of the [NDB](#) storage engine, as well as fixing recently discovered bugs in previous NDB Cluster releases.

Obtaining MySQL NDB Cluster 8.4. NDB Cluster 8.4 source code and binaries can be obtained from <https://dev.mysql.com/downloads/cluster/>.

For an overview of major changes made in NDB Cluster 8.4, see [What is New in MySQL NDB Cluster 8.4](#).

This release also incorporates all bug fixes and changes made in previous NDB Cluster releases, as well as all bug fixes and feature changes which were added in mainline MySQL 8.4 through MySQL 8.4.2 (see [Changes in MySQL 8.4.2 \(2024-07-23, LTS Release\)](#)).

This release contains no functional changes specific to MySQL NDB Cluster, and is published to align with and include changes made in MySQL Server 8.4.2.

Changes in MySQL NDB Cluster 8.4.1 (2024-07-02, LTS Release)

MySQL NDB Cluster 8.4.1 is a new LTS release of NDB 8.4, based on MySQL Server 8.4 and including features in version 8.4 of the [NDB](#) storage engine, as well as fixing recently discovered bugs in previous NDB Cluster releases.

Obtaining MySQL NDB Cluster 8.4. NDB Cluster 8.4 source code and binaries can be obtained from <https://dev.mysql.com/downloads/cluster/>.

For an overview of major changes made in NDB Cluster 8.4, see [What is New in MySQL NDB Cluster 8.4](#).

This release also incorporates all bug fixes and changes made in previous NDB Cluster releases, as well as all bug fixes and feature changes which were added in mainline MySQL 8.4 through MySQL 8.4.1 (see [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 Cluster 8.4.2 instead.

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

Functionality Added or Changed

- **Important Change:** Now, when the removal of a data node file or directory fails with a file does not exist ([ENOENT](#)) error, this is treated as a successful removal.
- **ndbinfo Information Database:** Added a [type](#) column to the [transporter_details](#) table in the [ndbinfo](#) information database. This column shows the type of connection used by the transporter, which is either of [TCP](#) or [SHM](#).
- **NDB Client Programs:** Added the [--CA-days](#) option to [ndb_sign_keys](#) to make it possible to specify a certificate's lifetime. (Bug #36549567)
- **NDB Client Programs:** When started, [ndbd](#) now produces a warning in the data node log like this one:

```
2024-05-28 13:32:16 [ndbd] WARNING -- Running ndbd with a single thread of
signal execution. For multi-threaded signal execution run the ndbmtd binary.
```

(Bug #36326896)

Bugs Fixed

- **NDB Replication:** When subscribing to changes in the [mysql.ndb_apply_status](#) table, different settings were used depending on whether [ndb_log_apply_status](#) was [ON](#) or [OFF](#). Since

`ndb_log_apply_status` can be changed at runtime and subscriptions are not recreated at that time, changing these settings at runtime did not have the desired effect.

The difference between enabling `ndb_log_apply_status` dynamically at runtime and doing so from the start of the MySQL process was in the format used when writing the `ndb_apply_status` updates to the binary log. When `ndb_log_apply_status` was enabled at runtime, writes were still done using the `UPDATE` format when `WRITE` was intended.

To fix this inconsistency and make the behavior more distinct, we now always use `WRITE` format in such cases; using the `WRITE` format also makes the binary log image slightly smaller and is thus preferred. In addition, the cleanup of old events has been improved, which improves the cleanup of failed attempts to create tables and events. (Bug #36453684)

- **NDB Replication:** The binary log index purge callback was skipped for the replica applier, which caused orphan rows to be left behind in the `ndb_binlog_index` table. (Bug #20573020, Bug #35847745, Bug #36378551, Bug #36420628, Bug #36423593, Bug #36485220, Bug #36492736)
- **NDB Cluster APIs:** It was possible to employ the following NDB API methods without them being used as `const`, although this alternative usage had long been deprecated (and was not actually documented):

- `Dictionary::listEvents()`
- `Dictionary::listIndexes()`
- `Dictionary::listObjects()`
- `NdbOperation::getNdbErrorLine()`

Now, each of these methods must always be invoked as `const`. (Bug #36165876)

- **NDB Client Programs:** `ndb_redo_log_reader` could not read data from encrypted files. (Bug #36313482)
- **NDB Client Programs:** `ndb_redo_log_reader` exited with `Record type = 0 not implemented` when reaching an unused page, all zero bytes, or a page which was only partially used (typically a page consisting of the page header only). (Bug #36313259)
- **NDB Client Programs:** `ndb_restore` did not restore a foreign key whose columns differed in order from those of the parent key.

Our thanks to Axel Svensson for the contribution. (Bug #114147, Bug #36345882)

- The destructor for `NDB_SCHEMA_OBJECT` makes several assertions about the state of the schema object, but the state was protected by a mutex, and the destructor did not acquire this mutex before testing the state.

We fix this by acquiring the mutex within the destructor. (Bug #36568964)

- `NDB` now writes a message to the MySQL server log before and after logging an incident in the binary log. (Bug #36548269)
- Removed a memory leak in `/util/NodeCertificate.cpp`. (Bug #36537931)
- Removed a memory leak from `src/ndbapi/NdbDictionaryImpl.cpp`. (Bug #36532102)
- The internal method `CertLifetime::set_set_cert_lifetime(X509 *cert)` should set the not-before and not-after times in the certificate to the same as those stored in the `CertLifetime` object, but instead it set the not-before time to the current time, and the not-after time to be of the same duration as the object. (Bug #36514834)
- Removed a possible use-after-free warning in `ConfigObject::copy_current()`. (Bug #36497108)

- When a thread acquires and releases the global schema lock required for schema changes and reads, the associated log message did not identify who performed the operation.

To fix this issue, we now do the following:

- Prepend the message in the log with the identification of the NDB Cluster component or user session responsible.
- Provide information about the related Performance Schema thread so that it can be traced.

(Bug #36446730)

References: See also: Bug #36446604.

- Metadata changes were not logged with their associated thread IDs. (Bug #36446604)

References: See also: Bug #36446730.

- When building NDB using `lld`, the build terminated prematurely with the error message `ld.lld: error: version script assignment of 'local' to symbol 'my_init' failed: symbol not defined` while attempting to link `libndbclient.so`. (Bug #36431274)
- TLS did not fail cleanly on systems which used OpenSSL 1.0, which is unsupported. Now in such cases, users get a clear error message advising that an upgrade to OpenSSL 1.1 or later is required to use TLS with NDB Cluster. (Bug #36426461)
- The included `libxml2` library was updated to version 2.9.13. (Bug #36417013)
- NDB Cluster's pushdown join functionality expects pushed conditions to filter exactly, so that no rows that do not match the condition must be returned, and all rows that do match the condition must be returned. When the condition contained a BINARY value compared to a BINARY column this was not always true; if the value was shorter than the column size, it could compare as equal to a column value despite having different lengths, if the condition was pushed down to NDB.

Now, when deciding whether a condition is pushable, we also make sure that the BINARY value length exactly matches the BINARY column's size. In addition, when binary string values were used in conditions with BINARY or VARBINARY columns, the actual length of a given string value was not used but rather an overestimate of its length. This is now changed; this should allow more conditions comparing short string values with VARBINARY columns to be pushed down than before this fix was made. (Bug #36390313, Bug #36513270)

References: See also: Bug #36399759, Bug #36400256. This issue is a regression of: Bug #36364619.

- Setting `AutomaticThreadConfig` and `NumCPUs` when running single-threaded data nodes (`ndbd`) sometimes led to unrecoverable errors. Now `ndbd` ignores settings for these parameters, which are intended to apply only to multi-threaded data nodes (`ndbmt`). (Bug #36388981)
- Improved the error message returned when trying to add a primary key to an NDBCLUSTER table using `ALGORITHM=INPLACE`. (Bug #36382071)

References: See also: Bug #30766579.

- The handling of the LQH operation pool which occurs as part of TC takeover skipped the last element in either of the underlying physical pools (static or dynamic). If this element was in use, holding an operation record for a transaction belonging to a transaction coordinator on the failed node, it was not returned, resulting in an incomplete takeover which sometimes left operations behind. Such operations interfered with subsequent transactions and the copying process (`CopyFrag`) used by the failed node to recover.

To fix this problem, we avoid skipping the final record while iterating through the LQH operation records during TC takeover. (Bug #36363119)

- The `libssh` library was updated to version 0.10.4. (Bug #36135621)
- When distribution awareness was not in use, the cluster tended to choose the same data node as the transaction coordinator repeatedly. (Bug #35840020, Bug #36554026)
- In certain cases, management nodes were unable to allocate node IDs to restarted data and SQL nodes. (Bug #35658072)
- Setting `ODirect` in the cluster's configuration caused excess logging when verifying that `ODirect` was actually settable for all paths. (Bug #34754817)
- In some cases, when trying to perform an online add index operation on an `NDB` table with no explicit primary key (see [Limitations of NDB online operations](#)), the resulting error message did not make the nature of the problem clear. (Bug #30766579)

References: See also: Bug #36382071.

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

MySQL NDB Cluster 8.4.0 is a new development release of NDB 8.4, based on MySQL Server 8.4 and including features in version 8.4 of the `NDB` storage engine, as well as fixing recently discovered bugs in previous NDB Cluster releases.

Obtaining MySQL NDB Cluster 8.4. NDB Cluster 8.4 source code and binaries can be obtained from <https://dev.mysql.com/downloads/cluster/>.

For an overview of major changes made in NDB Cluster 8.4, see [What is New in MySQL NDB Cluster 8.4](#).

This release also incorporates all bug fixes and changes made in previous NDB Cluster releases, as well as all bug fixes and feature changes which were added in mainline MySQL 8.4 through MySQL 8.4.0 (see [Changes in MySQL 8.4.0 \(2024-04-30, LTS Release\)](#)).

- [Deprecation and Removal Notes](#)
- [ndbinfo Information Database](#)
- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

Deprecation and Removal Notes

- **Packaging; Linux:** Removed the deprecated tool `/usr/bin/pathfix.py` from packages for Fedora 39. (Bug #35997178)
- The unused `INFORMATION_SCHEMA.TABLESPACES` table, deprecated in MySQL 8.0.22, has now been removed.

The Information Schema `FILES` table provides tablespace-related information for `NDB` tables. (WL #14065)

ndbinfo Information Database

- The `ndbinfo.transporter_details` table, introduced in NDB 8.0, provides information about individual transporters used in an NDB Cluster, rather than aggregate data as shown by the `transporters` table.

This release adds the following columns to `transporter_details`:

- `sendbuffer_used_bytes`: Number of bytes of signal data currently stored pending send using this transporter.

- `sendbuffer_max_used_bytes`: Historical maximum number of bytes of signal data stored pending send using this transporter. Reset when the transporter connects.
- `sendbuffer_alloc_bytes`: Number of bytes of send buffer currently allocated to store pending send bytes for this transporter. Send buffer memory is allocated in large blocks which may be sparsely used.
- `sendbuffer_max_alloc_bytes`: Historical maximum number of bytes of send buffer allocated to store pending send bytes for this transporter.

for more information, see [The ndbinfo transporter_details Table](#). . (WL #7662)

Functionality Added or Changed

- **Packaging:** Added support for Fedora 40 and Ubuntu 24.04.
- **NDB Replication:** Previously, when SQL nodes performing binary logging had `log_replica_updates=OFF`, replicated updates applied on a replica NDB cluster were still sent to the SQL nodes performing binary logging. Such updates, as well as any updates that do not trigger logging, are no longer sent, in order to decrease network traffic and resource consumption. (WL #15407)
- **ndbinfo Information Database:** Added the `transporter_details` table to the `ndbinfo` information database. This table is similar to the `transporters` table, but provides information about individual transporters rather than in the aggregate.

For more information, see [The ndbinfo transporter_details Table](#). (Bug #113163, Bug #36031560)

- **NDB Client Programs:** Added the `--verbose` option to the `ndb_waiter` test program to control the verbosity level of the output. (Bug #34547034)
- Improved logging related to purging of the binary log, including start and completions times, and whether it is the injector which has initiated the purge. (Bug #36176983)

Bugs Fixed

- **NDB Replication:** Replication of an `NDB` table stopped under the following conditions:
 - The table had no explicit primary key
 - The table contained `BIT` columns
 - A hash scan was used to find the rows to be updated or deleted

To fix this issue, we now make sure that the hash keys for the table match on the source and the replica. (Bug #34199339)

- **NDB Cluster APIs:** TLS connection errors were printed even though TLS was not specified for connections.

To fix this issue, following an ignored TLS error, we explicitly reset the error condition in the management handle to `NO_ERROR`. (Bug #36354973)

- **NDB Cluster APIs:** The `NdbEventOperation` methods `hasError()` and `clearError()`, long deprecated, are effectively disabled: `hasError()` now returns a constant 0, and `clearError()` does nothing. To determine an event type, use `getEventType2()` instead.
- **NDB Client Programs:** In some cases, it was not possible to load certificates generated using `ndb_sign_keys`. (Bug #36430004)
- **NDB Client Programs:** The following command-line options did not function correctly for the `ndb_redo_log_reader` utility program:

- `--mbyte`
- `--page`
- `--pageindex`

(Bug #36313427)

- **NDB Client Programs:** A certificate lifetime generated by `ndb_sign_keys` should consist of a fixed number of days, plus a random amount of extra time provided by the OpenSSL function `RAND_bytes()`, casting the result to a signed integer value. Because this value could sometimes be negative, this led to extra time being subtracted rather than added.

We eliminate this problem by using an unsigned integer type to hold the value obtained from `RAND_bytes()`. (Bug #36270629)

- **NDB Client Programs:** Invoking `ndb_mgmd` with the `--bind-address` option could in some cases cause the program to terminate unexpectedly. (Bug #36263410)
- **NDB Client Programs:** Some NDB utilities such `ndb_show_tables` leaked memory from API connections when TLS was required by the data nodes, and with valid certificates. (Bug #36170703)
- **NDB Client Programs:** Work begun in NDB 8.0.18 and 8.0.20 to remove the unnecessary text `NDBT_ProgramExit ...` from the output of NDB programs is completed in this release. This message should no longer appear in the release binaries of any such programs. (Bug #36169823)

References: See also: Bug #27096741.

- **NDB Client Programs:** The output from `ndb_waiter --ndb-tls-search-path` was not correctly formatted. (Bug #36132430)
- **NDB Client Programs:** On Windows hosts, `ndb_sign_keys` could not locate the `ssh` program. (Bug #36053948)
- **NDB Client Programs:** `ndb_sign_keys` did not handle the `--CA-tool` option correctly on Windows. (Bug #36053908)
- **NDB Client Programs:** The use of a strict 80-character limit for `clang-format` on the file `CommandInterpreter.cpp` broke the formatting of the interactive help text in the NDB management client. (Bug #36034395)
- **NDB Client Programs:** Trying to start `ndb_mgmd` with `--bind-address=localhost` failed with the error `Illegal bind address`, which was returned from the MGM API when attempting to parse the bind address to split it into host and port parts. `localhost` is now accepted as a valid address in such cases. (Bug #36005903)
- The included `libexpat` library was updated to version 2.5.0. (Bug #36324146)
- An implicit rollback generated when refusing to discover a table in an ongoing transaction caused the entire transaction to roll back. This could happen when a table definition changed while a transaction was active. We also checked at such times to see whether the table already existed in the data dictionary, which also meant that a subsequent read from same table within the same transaction would (wrongly) allow discovery.

Now in such cases, we skip checking whether or not a given table already exists in the data dictionary; instead, we now always refuse discovery of a table that is altered while a transaction is ongoing and return an error to the user. (Bug #36191370)
- When a backup was restored using `ndb_restore` with `--disable-indexes` and `--restore-privilege-tables`, the ordered index of the primary key was lost on the `mysql.ndb_sql_metadata` table, and could not be rebuilt even with `--rebuild-indexes`. (Bug #36157626)

- `NDB` maintains both a local and a global pool of free send buffers. When send buffers cannot be allocated from the local pool `NDB` allocates one from the global pool; likewise, buffers are freed and returned to the global pool when the local pool has too many free buffers. Both of these allocations require a mutex to be locked.

In order to reduce contention on this global mutex, we attempt to over-allocate buffers from the global pool when needed, keeping the excess buffers in the local pool, when releasing excess buffers to the global pool this was done only to the limit determined by `max_free`. After having released to the global pool, such that the `max_free` limit was met, it was likely that additional buffers would soon be released, once again exceeding `max_free`. This caused extra contention on the global pool mutex.

To address this issue, we now reduce the free buffers to 2/3 of the `max_free` limit in such cases. (Bug #36108639)

- `SSL_pending()` data from an SSL-enabled `NdbSocket` was not adequately checked for. (Bug #36076879)
- In certain cases, `ndb_mgmd` hung when attempting to sending a stop signal to `ndbmtid`. (Bug #36066725)
- Starting a replica to apply changes when `NDB` was not yet ready or had no yet started led to an unhelpful error message (`Fatal error: Failed to run 'applier_start' hook`). This happened when the replica started and the applier start hook waited for the number of seconds specified by `--ndb-wait-setup` for `NDB` to become ready; if it was not ready by then, the start hook reported the failure. Now in such cases, we let processing continue, instead, and allow the error to be returned from `NDB`, which better indicates its true source. (Bug #36054134)
- A `mysqld` process took much longer than expected to shut down when all data nodes were unreachable. (Bug #36052113)
- Negated the need for handling in the `NDB` binary log injector thread for a failure to instantiate an injector transaction by removing a potential point of failure in that operation. (Bug #36048889)
- It was possible in certain cases for the `TRPMAN` block to operate on transporters outside its own receive thread. (Bug #36028782)
- Removed a possible race condition between `start_clients_thread()` and `update_connections()`, due to both of these seeing the same transporter in the `DISCONNECTING` state. Now we make sure that disconnection is in fact completed before we set indicating that that the transporter has disconnected, so that `update_connections()` cannot close the `NdbSocket` before it has been completely shut down. (Bug #36009860)
- When a transporter was overloaded, the send thread did not yield to the CPU as expected, instead retrying the transporter repeatedly until reaching the hard-coded 200 microsecond timeout. (Bug #36004838)
- A MySQL server disconnected from schema distribution was unable to set up event operations because the table columns could not be found in the event. This could be made to happen by using `ndb_drop_table` or another means to drop a table directly from `NDB` that had been created using the MySQL server.

We fix this by making sure in such cases that we properly invalidate the `NDB` table definition from the dictionary cache. (Bug #35948153)

- The `ndb_sign_keys` utility's `--remote-openssl` option did not function as expected. (Bug #35853405)
- A replica could not apply a row change while handling a `Table definition changed` error. Now any such error is handled as a temporary error which can be retried multiple times. (Bug #35826145)

- Repeated incomplete attempts to perform a system restart in some cases left the cluster in a state from which it could not recover without restoring it from backup. (Bug #35801548)
- The event buffer used by the NDB API maintains an internal pool of free memory to reduce the interactions with the runtime and operating system, while allowing memory that is no longer needed to be returned for other uses. This free memory is subtracted from the total allocated memory to determine the memory in use which is reported and used for enforcing buffer limits and other purposes; this was represented using a 32-bit value, so that if it exceeded 4 GB, the value wrapped, and the amount of free memory appeared to be reduced. This had potentially adverse effects on event buffer memory release to the runtime and OS, free memory reporting, and memory limit handling.

This is fixed by using a 64-bit value to represent the amount of pooled free memory. (Bug #35483764)

References: See also: Bug #35655162, Bug #35663761.

- `START REPLICA`, `STOP REPLICA`, and `RESET REPLICA` statements are now written to `mysqld.log`. (Bug #35207235)
- NDB transporter handling in `mt.cpp` differentiated between neighbor transporters carrying signals between nodes in the same node group, and all other transporters. This sometimes led to issues with multiple transporters when a transporter connected nodes that were neighbors with nodes that were not. (Bug #33800633)
- Removed unnecessary warnings generated by transient disconnections of data nodes during restore operations. (Bug #33144487)
- During setup of utility tables, the schema event handler sometimes hung waiting for the global schema lock (GSL) to become available. This could happen when the physical tables had been dropped from the cluster, or when the connection was lost for some other reason. Now we use a try lock when attempting to acquire the GSL in such cases, thus causing another setup check attempt to be made at a later time if the global schema lock is not available. (Bug #32550019, Bug #35949017)
- API nodes did not record any information in the log relating to disconnects due to missed heartbeats from the data nodes. (Bug #29623286)

Release Series Changelogs: MySQL NDB Cluster 8.4

This section contains unified changelog information for the NDB Cluster 8.4 release series.

For changelogs covering individual MySQL NDB Cluster 8.4 releases, see [NDB Cluster Release Notes](#).

For general information about features added in MySQL NDB Cluster 8.4, see [What is New in NDB Cluster 8.4](#).

For an overview of features added in MySQL 8.4 that are not specific to NDB Cluster, see [What Is New in MySQL 8.4 since MySQL 8.0](#). For a complete list of all bug fixes and feature changes made in MySQL 8.4 that are not specific to NDB Cluster, see the MySQL 8.4 [Release Notes](#).

Changes in MySQL NDB Cluster 8.4.1 (2024-07-02, 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 Cluster 8.4.2 instead.

- [Functionality Added or Changed](#)

- [Bugs Fixed](#)

Functionality Added or Changed

- **Important Change:** Now, when the removal of a data node file or directory fails with a file does not exist ([ENOENT](#)) error, this is treated as a successful removal.
- **ndbinfo Information Database:** Added a [type](#) column to the [transporter_details](#) table in the [ndbinfo](#) information database. This column shows the type of connection used by the transporter, which is either of [TCP](#) or [SHM](#).
- **NDB Client Programs:** Added the [--CA-days](#) option to [ndb_sign_keys](#) to make it possible to specify a certificate's lifetime. (Bug #36549567)
- **NDB Client Programs:** When started, [ndbd](#) now produces a warning in the data node log like this one:

```
2024-05-28 13:32:16 [ndbd] WARNING -- Running ndbd with a single thread of
signal execution. For multi-threaded signal execution run the ndbmtd binary.
```

(Bug #36326896)

Bugs Fixed

- **NDB Cluster APIs:** It was possible to employ the following NDB API methods without them being used as [const](#), although this alternative usage had long been deprecated (and was not actually documented):
 - [Dictionary::listEvents\(\)](#)
 - [Dictionary::listIndexes\(\)](#)
 - [Dictionary::listObjects\(\)](#)
 - [NdbOperation::getNdbErrorLine\(\)](#)Now, each of these methods must always be invoked as [const](#). (Bug #36165876)
- **NDB Client Programs:** [ndb_redo_log_reader](#) could not read data from encrypted files. (Bug #36313482)
- **NDB Client Programs:** [ndb_redo_log_reader](#) exited with [Record type = 0 not implemented](#) when reaching an unused page, all zero bytes, or a page which was only partially used (typically a page consisting of the page header only). (Bug #36313259)
- **NDB Client Programs:** [ndb_restore](#) did not restore a foreign key whose columns differed in order from those of the parent key.

Our thanks to Axel Svensson for the contribution. (Bug #114147, Bug #36345882)

- The destructor for [NDB_SCHEMA_OBJECT](#) makes several assertions about the state of the schema object, but the state was protected by a mutex, and the destructor did not acquire this mutex before testing the state.

We fix this by acquiring the mutex within the destructor. (Bug #36568964)

- [NDB](#) now writes a message to the MySQL server log before and after logging an incident in the binary log. (Bug #36548269)
- Removed a memory leak in [/util/NodeCertificate.cpp](#). (Bug #36537931)
- Removed a memory leak from [src/ndbapi/NdbDictionaryImpl.cpp](#). (Bug #36532102)

- The internal method `CertLifetime::set_set_cert_lifetime(X509 *cert)` should set the not-before and not-after times in the certificate to the same as those stored in the `CertLifetime` object, but instead it set the not-before time to the current time, and the not-after time to be of the same duration as the object. (Bug #36514834)
- Removed a possible use-after-free warning in `ConfigObject::copy_current()`. (Bug #36497108)
- When a thread acquires and releases the global schema lock required for schema changes and reads, the associated log message did not identify who performed the operation.

To fix this issue, we now do the following:

- Prepend the message in the log with the identification of the NDB Cluster component or user session responsible.
- Provide information about the related Performance Schema thread so that it can be traced.

(Bug #36446730)

References: See also: Bug #36446604.

- Metadata changes were not logged with their associated thread IDs. (Bug #36446604)

References: See also: Bug #36446730.

- When building NDB using `lld`, the build terminated prematurely with the error message `ld.lld: error: version script assignment of 'local' to symbol 'my_init' failed: symbol not defined` while attempting to link `libndbclient.so`. (Bug #36431274)
- TLS did not fail cleanly on systems which used OpenSSL 1.0, which is unsupported. Now in such cases, users get a clear error message advising that an upgrade to OpenSSL 1.1 or later is required to use TLS with NDB Cluster. (Bug #36426461)
- NDB Cluster's pushdown join functionality expects pushed conditions to filter exactly, so that no rows that do not match the condition must be returned, and all rows that do match the condition must be returned. When the condition contained a BINARY value compared to a BINARY column this was not always true; if the value was shorter than the column size, it could compare as equal to a column value despite having different lengths, if the condition was pushed down to NDB.

Now, when deciding whether a condition is pushable, we also make sure that the BINARY value length exactly matches the BINARY column's size. In addition, when binary string values were used in conditions with BINARY or VARBINARY columns, the actual length of a given string value was not used but rather an overestimate of its length. This is now changed; this should allow more conditions comparing short string values with VARBINARY columns to be pushed down than before this fix was made. (Bug #36390313, Bug #36513270)

References: See also: Bug #36399759, Bug #36400256. This issue is a regression of: Bug #36364619.

- Setting `AutomaticThreadConfig` and `NumCPUs` when running single-threaded data nodes (`ndbd`) sometimes led to unrecoverable errors. Now `ndbd` ignores settings for these parameters, which are intended to apply only to multi-threaded data nodes (`ndbmt`). (Bug #36388981)
- Improved the error message returned when trying to add a primary key to an NDBCLUSTER table using `ALGORITHM=INPLACE`. (Bug #36382071)

References: See also: Bug #30766579.

- The handling of the LQH operation pool which occurs as part of TC takeover skipped the last element in either of the underlying physical pools (static or dynamic). If this element was in use, holding an operation record for a transaction belonging to a transaction coordinator on the failed node, it was

not returned, resulting in an incomplete takeover which sometimes left operations behind. Such operations interfered with subsequent transactions and the copying process ([CopyFrag](#)) used by the failed node to recover.

To fix this problem, we avoid skipping the final record while iterating through the [LQH](#) operation records during TC takeover. (Bug #36363119)

- When distribution awareness was not in use, the cluster tended to choose the same data node as the transaction coordinator repeatedly. (Bug #35840020, Bug #36554026)
- In certain cases, management nodes were unable to allocate node IDs to restarted data and SQL nodes. (Bug #35658072)
- Setting [ODirect](#) in the cluster's configuration caused excess logging when verifying that [ODirect](#) was actually settable for all paths. (Bug #34754817)
- In some cases, when trying to perform an online add index operation on an [NDB](#) table with no explicit primary key (see [Limitations of NDB online operations](#)), the resulting error message did not make the nature of the problem clear. (Bug #30766579)

References: See also: Bug #36382071.

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

- [Deprecation and Removal Notes](#)
- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

Deprecation and Removal Notes

- **Packaging; Linux:** Removed the deprecated tool `/usr/bin/pathfix.py` from packages for Fedora 39. (Bug #35997178)

Functionality Added or Changed

- **ndbinfo Information Database:** Added the `transporter_details` table to the `ndbinfo` information database. This table is similar to the `transporters` table, but provides information about individual transporters rather than in the aggregate.

For more information, see [The ndbinfo transporter_details Table](#). (Bug #113163, Bug #36031560)

- **NDB Client Programs:** Added the `--verbose` option to the `ndb_waiter` test program to control the verbosity level of the output. (Bug #34547034)
- Improved logging related to purging of the binary log, including start and completions times, and whether it is the injector which has initiated the purge. (Bug #36176983)

Bugs Fixed

- **NDB Cluster APIs:** TLS connection errors were printed even though TLS was not specified for connections.

To fix this issue, following an ignored TLS error, we explicitly reset the error condition in the management handle to `NO_ERROR`. (Bug #36354973)

- **NDB Cluster APIs:** The `NdbEventOperation` methods `hasError()` and `clearError()`, long deprecated, are effectively disabled: `hasError()` now returns a constant 0, and `clearError()` does nothing. To determine an event type, use `getEventType2()` instead.
- **NDB Client Programs:** In some cases, it was not possible to load certificates generated using `ndb_sign_keys`. (Bug #36430004)

- **NDB Client Programs:** The following command-line options did not function correctly for the `ndb_redo_log_reader` utility program:

- `--mbyte`
- `--page`
- `--pageindex`

(Bug #36313427)

- **NDB Client Programs:** A certificate lifetime generated by `ndb_sign_keys` should consist of a fixed number of days, plus a random amount of extra time provided by the OpenSSL function `RAND_bytes()`, casting the result to a signed integer value. Because this value could sometimes be negative, this led to extra time being subtracted rather than added.

We eliminate this problem by using an unsigned integer type to hold the value obtained from `RAND_bytes()`. (Bug #36270629)

- **NDB Client Programs:** Invoking `ndb_mgmd` with the `--bind-address` option could in some cases cause the program to terminate unexpectedly. (Bug #36263410)
- **NDB Client Programs:** Some NDB utilities such `ndb_show_tables` leaked memory from API connections when TLS was required by the data nodes, and with valid certificates. (Bug #36170703)
- **NDB Client Programs:** Work begun in NDB 8.0.18 and 8.0.20 to remove the unnecessary text `NDBT_ProgramExit ...` from the output of NDB programs is completed in this release. This message should no longer appear in the release binaries of any such programs. (Bug #36169823)

References: See also: Bug #27096741.

- **NDB Client Programs:** The output from `ndb_waiter --ndb-tls-search-path` was not correctly formatted. (Bug #36132430)
- **NDB Client Programs:** On Windows hosts, `ndb_sign_keys` could not locate the `ssh` program. (Bug #36053948)
- **NDB Client Programs:** `ndb_sign_keys` did not handle the `--CA-tool` option correctly on Windows. (Bug #36053908)
- **NDB Client Programs:** The use of a strict 80-character limit for `clang-format` on the file `CommandInterpreter.cpp` broke the formatting of the interactive help text in the NDB management client. (Bug #36034395)
- **NDB Client Programs:** Trying to start `ndb_mgmd` with `--bind-address=localhost` failed with the error `Illegal bind address`, which was returned from the MGM API when attempting to parse the bind address to split it into host and port parts. `localhost` is now accepted as a valid address in such cases. (Bug #36005903)
- An implicit rollback generated when refusing to discover a table in an ongoing transaction caused the entire transaction to roll back. This could happen when a table definition changed while a transaction was active. We also checked at such times to see whether the table already existed in the data dictionary, which also meant that a subsequent read from same table within the same transaction would (wrongly) allow discovery.

Now in such cases, we skip checking whether or not a given table already exists in the data dictionary; instead, we now always refuse discovery of a table that is altered while a transaction is ongoing and return an error to the user. (Bug #36191370)

- When a backup was restored using `ndb_restore` with `--disable-indexes` and `--restore-privilege-tables`, the ordered index of the primary key was lost on the

`mysql.ndb_sql_metadata` table, and could not be rebuilt even with `--rebuild-indexes`. (Bug #36157626)

- **NDB** maintains both a local and a global pool of free send buffers. When send buffers cannot be allocated from the local pool **NDB** allocates one from the global pool; likewise, buffers are freed and returned to the global pool when the local pool has too many free buffers. Both of these allocations require a mutex to be locked.

In order to reduce contention on this global mutex, we attempt to over-allocate buffers from the global pool when needed, keeping the excess buffers in the local pool, when releasing excess buffers to the global pool this was done only to the limit determined by `max_free`. After having released to the global pool, such that the `max_free` limit was met, it was likely that additional buffers would soon be released, once again exceeding `max_free`. This caused extra contention on the global pool mutex.

To address this issue, we now reduce the free buffers to 2/3 of the `max_free` limit in such cases. (Bug #36108639)

- `SSL_pending()` data from an SSL-enabled `NdbSocket` was not adequately checked for. (Bug #36076879)
- In certain cases, `ndb_mgmd` hung when attempting to sending a stop signal to `ndbmt.d`. (Bug #36066725)
- Starting a replica to apply changes when **NDB** was not yet ready or had no yet started led to an unhelpful error message (`Fatal error: Failed to run 'applier_start' hook`). This happened when the replica started and the applier start hook waited for the number of seconds specified by `--ndb-wait-setup` for **NDB** to become ready; if it was not ready by then, the start hook reported the failure. Now in such cases, we let processing continue, instead, and allow the error to be returned from **NDB**, which better indicates its true source. (Bug #36054134)
- A `mysqld` process took much longer than expected to shut down when all data nodes were unreachable. (Bug #36052113)
- Negated the need for handling in the **NDB** binary log injector thread for a failure to instantiate an injector transaction by removing a potential point of failure in that operation. (Bug #36048889)
- It was possible in certain cases for the `TRPMAN` block to operate on transporters outside its own receive thread. (Bug #36028782)
- Removed a possible race condition between `start_clients_thread()` and `update_connections()`, due to both of these seeing the same transporter in the `DISCONNECTING` state. Now we make sure that disconnection is in fact completed before we set indicating that that the transporter has disconnected, so that `update_connections()` cannot close the `NdbSocket` before it has been completely shut down. (Bug #36009860)
- When a transporter was overloaded, the send thread did not yield to the CPU as expected, instead retrying the transporter repeatedly until reaching the hard-coded 200 microsecond timeout. (Bug #36004838)
- A MySQL server disconnected from schema distribution was unable to set up event operations because the table columns could not be found in the event. This could be made to happen by using `ndb_drop_table` or another means to drop a table directly from **NDB** that had been created using the MySQL server.

We fix this by making sure in such cases that we properly invalidate the **NDB** table definition from the dictionary cache. (Bug #35948153)

- The `ndb_sign_keys` utility's `--remote-openssl` option did not function as expected. (Bug #35853405)

- A replica could not apply a row change while handling a `Table definition changed` error. Now any such error is handled as a temporary error which can be retried multiple times. (Bug #35826145)
- Repeated incomplete attempts to perform a system restart in some cases left the cluster in a state from which it could not recover without restoring it from backup. (Bug #35801548)
- The event buffer used by the NDB API maintains an internal pool of free memory to reduce the interactions with the runtime and operating system, while allowing memory that is no longer needed to be returned for other uses. This free memory is subtracted from the total allocated memory to determine the memory in use which is reported and used for enforcing buffer limits and other purposes; this was represented using a 32-bit value, so that if it exceeded 4 GB, the value wrapped, and the amount of free memory appeared to be reduced. This had potentially adverse effects on event buffer memory release to the runtime and OS, free memory reporting, and memory limit handling.

This is fixed by using a 64-bit value to represent the amount of pooled free memory. (Bug #35483764)

References: See also: Bug #35655162, Bug #35663761.

- `START REPLICA`, `STOP REPLICA`, and `RESET REPLICA` statements are now written to `mysqld.log`. (Bug #35207235)
- NDB transporter handling in `mt.cpp` differentiated between neighbor transporters carrying signals between nodes in the same node group, and all other transporters. This sometimes led to issues with multiple transporters when a transporter connected nodes that were neighbors with nodes that were not. (Bug #33800633)
- Removed unnecessary warnings generated by transient disconnections of data nodes during restore operations. (Bug #33144487)
- During setup of utility tables, the schema event handler sometimes hung waiting for the global schema lock (GSL) to become available. This could happen when the physical tables had been dropped from the cluster, or when the connection was lost for some other reason. Now we use a try lock when attempting to acquire the GSL in such cases, thus causing another setup check attempt to be made at a later time if the global schema lock is not available. (Bug #32550019, Bug #35949017)
- API nodes did not record any information in the log relating to disconnects due to missed heartbeats from the data nodes. (Bug #29623286)

Index

A

API timeouts, 8
ApiConnectRecord, 8
ApiFailureHandlingTimeout, 6
AutomaticThreadConfig, 13, 20

B

backup, 6
BINARY, 13, 20
binary log, 13
binary log injector, 16, 23
binary log rotation, 8
bind address, 16, 23
BIT, 16
BLOB_INLINE_SIZE, 8, 11

C

CA-days, 13, 20

- certificates, 13, 20
- changes
 - NDB Cluster, 20
- Cleaner, 6
- compile-cluster, 6
- compiling, 13, 20
- condition pushdown, 13, 20
- configuration, 11
- const, 13, 20
- CopyFrag, 13, 20
- CopyGCI, 8
- CREATE TABLESPACE, 5

D

- DBLQH, 13, 20
- DBSPJ, 8
- DBTC, 8, 13, 20
- disable-indexes, 16, 23

E

- encryption, 5
- error handling, 16, 23
- errors, 5
- EventBuffer, 16, 23
- expat, 16

F

- failure handling, 11
- free memory, 16, 23

G

- global schema lock, 13, 16, 20, 23

H

- hash scan, 16
- heartbeats, 16, 23
- help text, 16, 23

I

- Important Change, 13, 20
- include-stored-grants, 6
- index statistics, 11
- INFORMATION_SCHEMA, 16
- INPLACE, 13, 20
- InvocationHandler, 6

L

- libssh, 13
- libxml2, 6, 13
- Linux, 16, 23
- lld, 13, 20
- log level, 3
- logging, 5, 8, 13, 20
- logs, 16, 23
- log_replica, 16
- LQH_TRANSCONF, 6

M

macOS, 8
Microsoft Windows, 8
MySQL NDB ClusterJ, 6, 11
mysql.ndb_apply_status table, 13
mysqld, 16, 23

N

NDB Client Programs, 5, 6, 11, 13, 16, 20, 23
NDB Cluster, 3, 5, 6, 8, 11, 13, 16, 20, 23
NDB Cluster APIs, 8, 11, 13, 16, 20, 23
NDB Disk Data, 8
NDB Replication, 6, 8, 13, 16
ndb-wait-setup, 16, 23
ndbd, 13, 20
NdbDictionaryImpl.cpp, 13, 20
NdbEventOperation, 16, 23
ndbinfo Information Database, 6, 8, 13, 16, 20, 23
NdbRecord, 11
Ndb_cluster_connection, 8
ndb_config, 8
ndb_log_apply_status, 13
ndb_log_cache_size, 11
ndb_metadata_check, 6
ndb_mgmd, 16, 23
ndb_redo_log_reader, 13, 16, 20, 23
ndb_restore, 13, 20
ndb_schema_object, 6
NDB_SCHEMA_OBJECT, 13, 20
Ndb_schema_participant_count, 6
ndb_sign_keys, 13, 16, 20, 23
ndb_size.pl, 11
ndb_waiter, 16, 23
neighbor nodes, 16, 23
node IDs, 8
node shutdown, 11
NumCPUs, 13, 20

O

ODirect, 13, 20
online operations, 13, 20
OOBR, 6
OpenJPA, 11
openssl, 13, 20
out of buffer errors, 6

P

Packaging, 16, 23
PURGE BINARY LOGS, 6

Q

QMGR, 6

R

REPORT BACKUPSTATUS, 6
restarts, 13, 20

S

- scans, 8
- schema distribution, 16, 23
- send buffer, 16
- send threads, 16, 23
- SendBuffer, 16, 23
- SHM, 11
- signal dump, 6
- sockets, 16, 23
- SSL, 16, 23
- STMT_END_F, 8
- STOP, 16, 23
- system restarts, 8

T

- TABLESPACES table, 16
- TCRELEASEREQ, 8
- testsuite, 6
- TLS, 6, 13, 16, 20, 23
- transaction coordinator, 13, 20
- transactions, 16, 23
- transporters, 16, 23
- transporter_details, 13, 16, 20, 23
- transporter_details table, 16
- TRIX, 11
- TRPMAN, 16, 23

W

- warnings, 16, 23
- watchdog, 6

