
MySQL NDB Cluster 9.3 Release Notes

Abstract

This document contains release notes for the changes in each release of MySQL NDB Cluster that uses version 9.3 of the [NDB \(NDBCLUSTER\)](#) storage engine.

Each NDB Cluster 9.3 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 9.3](#).

For general information about features added in NDB Cluster 9.3, see [What is New in MySQL NDB Cluster 9.3](#). 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 9.3 documentation, see the [MySQL 9.3 Reference Manual](#), which includes an overview of features added in MySQL 9.3 that are not specific to NDB Cluster ([What Is New in MySQL 9.3](#)), and discussion of upgrade issues that you may encounter for upgrades from MySQL 9.2 to MySQL 9.3 ([Changes in MySQL 9.3](#)). For a complete list of all bug fixes and feature changes made in MySQL 9.3 that are not specific to [NDB](#), see [MySQL 9.3 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: 2025-08-12 (revision: 30436)

Table of Contents

Preface and Legal Notices	1
Changes in MySQL NDB Cluster 9.3.0 (2025-04-16, Innovation Release)	3
Index	6

Preface and Legal Notices

This document contains release notes for the changes in each release of MySQL NDB Cluster that uses version 9.3 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 9.3.0 (2025-04-16, Innovation Release)

MySQL NDB Cluster 9.3.0 is a new Innovation release of NDB Cluster, based on MySQL Server 9.3 and including features in version 9.3 of the [NDB](#) storage engine, as well as fixing recently discovered bugs in previous NDB Cluster releases.

Obtaining MySQL NDB Cluster 9.3. NDB Cluster 9.3 source code and binaries can be obtained from <https://dev.mysql.com/downloads/cluster/>.

For an overview of major changes made in NDB Cluster 9.3, see [What is New in MySQL NDB Cluster 9.3](#).

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 9 through MySQL 9.3.0 (see [Changes in MySQL 9.3.0 \(2025-04-15, Innovation Release\)](#)).

- [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

- **MySQL NDB ClusterJ:** For MySQL NDB Cluster 9.3.0 and later, Java 11 or above is required to build and run ClusterJ. (Bug #37213478)
- 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)
- The function [ndb_mgm_set_ignore_sigpipe\(\)](#) has been deprecated, and the MGM API no longer makes SIGPIPE to be ignored by all applications connected to a management node. SIGPIPE can be disabled on BSD platforms that support the socket option [SO_NOSIGPIPE](#), and for other platforms (and also all applications on all platforms), a signal handler can be used to ignore SIGPIPE. (Bug #35570336)

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 (OBR) 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 OBR process has completed for all buckets. (Bug #20648778)

- Removed the unused InnoDB and NDB handlerton `get_tablespace()` method. (Bug #109443, Bug #34916556)

Index

A

ApiFailureHandlingTimeout, 3

B

backup, 3

C

Cleaner, 3

compile-cluster, 3

I

include-stored-grants, 3

InvocationHandler, 3

L

libxml2, 3

LQH_TRANSCONF, 3

M

MGM API, 3

MySQL NDB ClusterJ, 3

N

NDB Client Programs, 3

NDB Cluster, 3

NDB Replication, 3

ndbinfo Information Database, 3

ndb_metadata_check, 3

ndb_schema_object, 3

Ndb_schema_participant_count, 3

O

OOBR, 3

out of buffer errors, 3

P

PURGE BINARY LOGS, 3

Q

QMGR, 3

R

REPORT BACKUPSTATUS, 3

S

signal dump, 3

T

testsuite, 3

TLS, 3

W

watchdog, 3