

---

# MySQL NDB Cluster 7.3 Release Notes

## Abstract

This document contains release notes for the changes in each release of MySQL NDB Cluster that uses version 7.3 of the [NDB \(NDBCLUSTER\)](#) storage engine.

Each NDB Cluster 7.3 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 7.3 and NDB Cluster 7.4](#).

For general information about features added in NDB Cluster 7.3, see [What is New in MySQL NDB Cluster](#). For a complete list of all bug fixes and feature changes in MySQL Cluster, please refer to the changelog section for each individual NDB Cluster release.

For additional MySQL 5.6 documentation, see the [MySQL 5.6 Reference Manual](#), which includes an overview of features added in MySQL 5.6 that are not specific to NDB Cluster ([What Is New in MySQL 5.6](#)), and discussion of upgrade issues that you may encounter for upgrades from MySQL 5.5 to MySQL 5.6 ([Changes Affecting Upgrades to MySQL 5.6](#)). For a complete list of all bug fixes and feature changes made in MySQL 5.6 that are not specific to [NDB](#), see [MySQL 5.6 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 (<http://dev.mysql.com/downloads/>), the version has not yet been released.

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

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

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

Document generated on: 2018-06-22 (revision: 15419)

## Table of Contents

Preface and Legal Notices .....	2
Changes in MySQL NDB Cluster 7.3.22 (5.6.41-ndb-7.3.22) (Not yet released, General Availability) .....	4
Changes in MySQL NDB Cluster 7.3.21 (5.6.40-ndb-7.3.21) (2018-04-20, General Availability) .....	4
Changes in MySQL NDB Cluster 7.3.20 (5.6.39-ndb-7.3.20) (2018-01-17, General Availability) .....	4
Changes in MySQL NDB Cluster 7.3.19 (5.6.38-ndb-7.3.19) (2017-10-18, General Availability) .....	5
Changes in MySQL NDB Cluster 7.3.18 (5.6.37-ndb-7.3.18) (2017-07-18, General Availability) .....	6
Changes in MySQL NDB Cluster 7.3.17 (5.6.36-ndb-7.3.17) (2017-04-10, General Availability) .....	7
Changes in MySQL NDB Cluster 7.3.16 (5.6.35-ndb-7.3.16) (2017-01-17, General Availability) .....	8
Changes in MySQL NDB Cluster 7.3.15 (5.6.34-ndb-7.3.15) (2016-10-18, General Availability) .....	9
Changes in MySQL NDB Cluster 7.3.14 (5.6.31-ndb-7.3.14) (2016-07-18, General Availability) .....	10
Changes in MySQL NDB Cluster 7.3.13 (5.6.29-ndb-7.3.13) (2016-04-18, General Availability) .....	12
Changes in MySQL NDB Cluster 7.3.12 (5.6.28-ndb-7.3.12) (2016-01-19, General Availability) .....	13
Changes in MySQL NDB Cluster 7.3.11 (5.6.27-ndb-7.3.11) (2015-10-19, General Availability) .....	16
Changes in MySQL NDB Cluster 7.3.10 (5.6.25-ndb-7.3.10) (2015-07-13, General Availability) .....	18
Changes in MySQL NDB Cluster 7.3.9 (5.6.24-ndb-7.3.9) (2015-04-14, General Availability) .....	21

Changes in MySQL NDB Cluster 7.3.8 (5.6.22-ndb-7.3.8) (2015-01-21, General Availability) .....	24
Changes in MySQL NDB Cluster 7.3.7 (5.6.21-ndb-7.3.7) (2014-10-17, General Availability) .....	29
Changes in MySQL NDB Cluster 7.3.6 (5.6.19-ndb-7.3.6) (2014-07-11, General Availability) .....	32
Changes in MySQL NDB Cluster 7.3.5 (5.6.17-ndb-7.3.5) (2014-04-07, General Availability) .....	37
Changes in MySQL NDB Cluster 7.3.4 (5.6.15-ndb-7.3.4) (2014-02-06, General Availability) .....	41
Changes in MySQL NDB Cluster 7.3.3 (5.6.14-ndb-7.3.3) (2013-11-18, General Availability) .....	44
Changes in MySQL NDB Cluster 7.3.2 (5.6.11-ndb-7.3.2) (2013-06-18, General Availability) .....	48
Changes in MySQL NDB Cluster 7.3.1 (5.6.10-ndb-7.3.1) (2013-04-17, Development Milestone) .....	51
Release Series Changelogs: MySQL NDB Cluster 7.3 .....	54
Changes in MySQL NDB Cluster 7.3.19 (5.6.38-ndb-7.3.19) (2017-10-18, General Availability) .....	54
Changes in MySQL NDB Cluster 7.3.18 (5.6.37-ndb-7.3.18) (2017-07-18, General Availability) .....	55
Changes in MySQL NDB Cluster 7.3.17 (5.6.36-ndb-7.3.17) (2017-04-10, General Availability) .....	55
Changes in MySQL NDB Cluster 7.3.16 (5.6.35-ndb-7.3.16) (2017-01-17, General Availability) .....	56
Changes in MySQL NDB Cluster 7.3.15 (5.6.34-ndb-7.3.15) (2016-10-18, General Availability) .....	57
Changes in MySQL NDB Cluster 7.3.14 (5.6.31-ndb-7.3.14) (2016-07-18, General Availability) .....	57
Changes in MySQL NDB Cluster 7.3.13 (5.6.29-ndb-7.3.13) (2016-04-18, General Availability) .....	59
Changes in MySQL NDB Cluster 7.3.12 (5.6.28-ndb-7.3.12) (2016-01-19, General Availability) .....	60
Changes in MySQL NDB Cluster 7.3.11 (5.6.27-ndb-7.3.11) (2015-10-19, General Availability) .....	61
Changes in MySQL NDB Cluster 7.3.10 (5.6.25-ndb-7.3.10) (2015-07-13, General Availability) .....	63
Changes in MySQL NDB Cluster 7.3.9 (5.6.24-ndb-7.3.9) (2015-04-14, General Availability) .....	66
Changes in MySQL NDB Cluster 7.3.8 (5.6.22-ndb-7.3.8) (2015-01-21, General Availability) .....	68
Changes in MySQL NDB Cluster 7.3.7 (5.6.21-ndb-7.3.7) (2014-10-17, General Availability) .....	72
Changes in MySQL NDB Cluster 7.3.6 (5.6.19-ndb-7.3.6) (2014-07-11, General Availability) .....	75
Changes in MySQL NDB Cluster 7.3.5 (5.6.17-ndb-7.3.5) (2014-04-07, General Availability) .....	79
Changes in MySQL NDB Cluster 7.3.4 (5.6.15-ndb-7.3.4) (2014-02-06, General Availability) .....	82
Changes in MySQL NDB Cluster 7.3.3 (5.6.14-ndb-7.3.3) (2013-11-18, General Availability) .....	85
Changes in MySQL NDB Cluster 7.3.2 (5.6.11-ndb-7.3.2) (2013-06-18, General Availability) .....	89
Changes in MySQL NDB Cluster 7.3.1 (5.6.10-ndb-7.3.1) (2013-04-17, Development Milestone) ...	91
Index .....	93

## Preface and Legal Notices

This document contains release notes for the changes in each release of MySQL NDB Cluster that uses version 7.3 of the [NDB](#) storage engine.

## Legal Notices

Copyright © 1997, 2018, Oracle and/or its affiliates. All rights reserved.

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

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

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

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users

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

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

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

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

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

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

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 7.3.22 (5.6.41-ndb-7.3.22) (Not yet released, General Availability)

MySQL NDB Cluster 7.3.22 is a new release of NDB Cluster, based on MySQL Server 5.6 and including features from version 7.3 of the [NDB](#) storage engine, as well as fixing a number of recently discovered bugs in previous NDB Cluster releases.

**Obtaining MySQL NDB Cluster 7.3.** MySQL NDB Cluster 7.3 source code and binaries can be obtained from <http://dev.mysql.com/downloads/cluster/>.

For an overview of changes made in MySQL NDB Cluster 7.3, see [What is New in NDB Cluster 7.3](#).

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 5.6 through MySQL 5.6.41 (see [Changes in MySQL 5.6.41 \(Not yet released, General Availability\)](#)).

Version 5.6.41-ndb-7.3.22 has no changelog entries, or they have not been published because the product version has not been released.

## Changes in MySQL NDB Cluster 7.3.21 (5.6.40-ndb-7.3.21) (2018-04-20, General Availability)

MySQL NDB Cluster 7.3.21 is a new release of NDB Cluster, based on MySQL Server 5.6 and including features from version 7.3 of the [NDB](#) storage engine, as well as fixing a number of recently discovered bugs in previous NDB Cluster releases.

**Obtaining MySQL NDB Cluster 7.3.** MySQL NDB Cluster 7.3 source code and binaries can be obtained from <http://dev.mysql.com/downloads/cluster/>.

For an overview of changes made in MySQL NDB Cluster 7.3, see [What is New in NDB Cluster 7.3](#).

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 5.6 through MySQL 5.6.40 (see [Changes in MySQL 5.6.40 \(2018-04-19, General Availability\)](#)).

### Bugs Fixed

- **NDB Cluster APIs:** Incorrect results, usually an empty result set, were returned when `setBound()` was used to specify a `NULL` bound. This issue appears to have been caused by a problem in gcc, limited to cases using the old version of this method (which does not employ `NdbRecord`), and is fixed by rewriting the problematic internal logic in the old implementation. (Bug #89468, Bug #27461752)
- Queries using very large lists with `IN` were not handled correctly, which could lead to data node failures. (Bug #27397802)
- `ndb_restore --print_data --hex` did not print trailing 0s of `LONGVARBINARY` values. (Bug #65560, Bug #14198580)

## Changes in MySQL NDB Cluster 7.3.20 (5.6.39-ndb-7.3.20) (2018-01-17, General Availability)

MySQL NDB Cluster 7.3.20 is a new release of NDB Cluster, based on MySQL Server 5.6 and including features from version 7.3 of the [NDB](#) storage engine, as well as fixing a number of recently discovered bugs in previous NDB Cluster releases.

**Obtaining MySQL NDB Cluster 7.3.** MySQL NDB Cluster 7.3 source code and binaries can be obtained from <http://dev.mysql.com/downloads/cluster/>.

For an overview of changes made in MySQL NDB Cluster 7.3, see [What is New in NDB Cluster 7.3](#).

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 5.6 through MySQL 5.6.39 (see [Changes in MySQL 5.6.39 \(2018-01-15, General Availability\)](#)).

## Bugs Fixed

- **NDB Replication:** On an SQL node not being used for a replication channel with `sql_log_bin=0` it was possible after creating and populating an NDB table for a table map event to be written to the binary log for the created table with no corresponding row events. This led to problems when this log was later used by a slave cluster replicating from the mysqld where this table was created.

Fixed this by adding support for maintaining a cumulative `any_value` bitmap for global checkpoint event operations that represents bits set consistently for all rows of a specific table in a given epoch, and by adding a check to determine whether all operations (rows) for a specific table are all marked as `NOLOGGING`, to prevent the addition of this table to the `Table_map` held by the binlog injector.

As part of this fix, the NDB API adds a new `getNextEventOpInEpoch3()` method which provides information about any `AnyValue` received by making it possible to retrieve the cumulative `any_value` bitmap. (Bug #26333981)

- A query against the `INFORMATION_SCHEMA.FILES` table returned no results when it included an `ORDER BY` clause. (Bug #26877788)
- The `NDBFS` block's `OM_SYNC` flag is intended to make sure that all `FSWRITEREQ` signals used for a given file are synchronized, but was ignored by platforms that do not support `O_SYNC`, meaning that this feature did not behave properly on those platforms. Now the synchronization flag is used on those platforms that do not support `O_SYNC`. (Bug #76975, Bug #21049554)

## Changes in MySQL NDB Cluster 7.3.19 (5.6.38-ndb-7.3.19) (2017-10-18, General Availability)

MySQL NDB Cluster 7.3.19 is a new release of NDB Cluster, based on MySQL Server 5.6 and including features from version 7.3 of the `NDB` storage engine, as well as fixing a number of recently discovered bugs in previous NDB Cluster releases.

**Obtaining MySQL NDB Cluster 7.3.** MySQL NDB Cluster 7.3 source code and binaries can be obtained from <http://dev.mysql.com/downloads/cluster/>.

For an overview of changes made in MySQL NDB Cluster 7.3, see [What is New in NDB Cluster 7.3](#).

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 5.6 through MySQL 5.6.38 (see [Changes in MySQL 5.6.38 \(2017-10-16, General Availability\)](#)).

## Bugs Fixed

- Added `DUMP` code 7027 to facilitate testing of issues relating to local checkpoints. For more information, see [DUMP 7027](#). (Bug #26661468)

- A previous fix intended to improve logging of node failure handling in the transaction coordinator included logging of transactions that could occur in normal operation, which made the resulting logs needlessly verbose. Such normal transactions are no longer written to the log in such cases. (Bug #26568782)

References: This issue is a regression of: Bug #26364729.

- Some `DUMP` codes used for the `LGMAN` kernel block were incorrectly assigned numbers in the range used for codes belonging to `DBTUX`. These have now been assigned symbolic constants and numbers in the proper range (10001, 10002, and 10003). (Bug #26365433)
- Node failure handling in the `DBTC` kernel block consists of a number of tasks which execute concurrently, and all of which must complete before TC node failure handling is complete. This fix extends logging coverage to record when each task completes, and which tasks remain, includes the following improvements:
  - Handling interactions between GCP and node failure handling interactions, in which TC takeover causes GCP participant stall at the master TC to allow it to extend the current GCI with any transactions that were taken over; the stall can begin and end in different GCP protocol states. Logging coverage is extended to cover all scenarios. Debug logging is now more consistent and understandable to users.
  - Logging done by the `QMGR` block as it monitors duration of node failure handling duration is done more frequently. A warning log is now generated every 30 seconds (instead of 1 minute), and this now includes `DBDIH` block debug information (formerly this was written separately, and less often).
  - To reduce space used, `DBTC instance number:` is shortened to `DBTC number:`.
  - A new error code is added to assist testing.

(Bug #26364729)

- A potential hundredfold signal fan-out when sending a `START_FRAG_REQ` signal could lead to a node failure due to a `job buffer full` error in start phase 5 while trying to perform a local checkpoint during a restart. (Bug #86675, Bug #26263397)

References: See also: Bug #26288247, Bug #26279522.

## Changes in MySQL NDB Cluster 7.3.18 (5.6.37-ndb-7.3.18) (2017-07-18, General Availability)

MySQL NDB Cluster 7.3.18 is a new release of NDB Cluster, based on MySQL Server 5.6 and including features from version 7.3 of the `NDB` storage engine, as well as fixing a number of recently discovered bugs in previous NDB Cluster releases.

**Obtaining MySQL NDB Cluster 7.3.** MySQL NDB Cluster 7.3 source code and binaries can be obtained from <http://dev.mysql.com/downloads/cluster/>.

For an overview of changes made in MySQL NDB Cluster 7.3, see [What is New in NDB Cluster 7.3](#).

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 5.6 through MySQL 5.6.37 (see [Changes in MySQL 5.6.37 \(2017-07-17, General Availability\)](#)).

- [Functionality Added or Changed](#)

- [Bugs Fixed](#)

## Functionality Added or Changed

- **Important Change; MySQL NDB ClusterJ:** The ClusterJPA plugin for OpenJPA is no longer supported by NDB Cluster, and has been removed from the distribution. (Bug #23563810)

## Bugs Fixed

- **NDB Replication:** Added a check to stop an NDB replication slave when configuration as a multithreaded slave is detected (for example, if `slave_parallel_workers` is set to a nonzero value). (Bug #21074209)
- **MySQL NDB ClusterJ:** Compilation for the tests for NDB JTie failed. It was due to how null references were handled, which has been corrected by this fix. (Bug #26080804)
- Backup `.log` files contained log entries for one or more extra fragments, due to an issue with filtering out changes logged by other nodes in the same node group. This resulted in a larger `.log` file and thus use of more resources than necessary; it could also cause problems when restoring, since backups from different nodes could interfere with one another while the log was being applied. (Bug #25891014)
- Error 240 is raised when there is a mismatch between foreign key trigger columns and the values supplied to them during trigger execution, but had no error message indicating the source of the problem. (Bug #23141739)

References: See also: Bug #23068914, Bug #85857.

- `ALTER TABLE ... MAX_ROWS=0` can now be performed only by using a copying `ALTER TABLE` statement. Resetting `MAX_ROWS` to 0 can no longer be performed using `ALGORITHM=INPLACE` or the `ONLINE` keyword. (Bug #21960004)
- When compiling the NDB kernel with `gcc` version 6.0.0 or later, it is now built using `-flifetime-dse=1`. (Bug #85381, Bug #25690926)

## Changes in MySQL NDB Cluster 7.3.17 (5.6.36-ndb-7.3.17) (2017-04-10, General Availability)

MySQL NDB Cluster 7.3.17 is a new release of NDB Cluster, based on MySQL Server 5.6 and including features from version 7.3 of the NDB storage engine, as well as fixing a number of recently discovered bugs in previous NDB Cluster releases.

**Obtaining MySQL NDB Cluster 7.3.** MySQL NDB Cluster 7.3 source code and binaries can be obtained from <http://dev.mysql.com/downloads/cluster/>.

For an overview of changes made in MySQL NDB Cluster 7.3, see [What is New in NDB Cluster 7.3](#).

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 5.6 through MySQL 5.6.36 (see [Changes in MySQL 5.6.36 \(2017-04-10, General Availability\)](#)).

## Bugs Fixed

- **NDB Disk Data:** Stale data from NDB Disk Data tables that had been dropped could potentially be included in backups due to the fact that disk scans were enabled for these. To prevent this possibility,

disk scans are now disabled—as are other types of scans—when taking a backup. (Bug #84422, Bug #25353234)

- **NDB Disk Data:** In some cases, setting dynamic in-memory columns of an NDB Disk Data table to `NULL` was not handled correctly. (Bug #79253, Bug #22195588)

## Changes in MySQL NDB Cluster 7.3.16 (5.6.35-ndb-7.3.16) (2017-01-17, General Availability)

MySQL NDB Cluster 7.3.16 is a new release of NDB Cluster, based on MySQL Server 5.6 and including features from version 7.3 of the `NDB` storage engine, as well as fixing a number of recently discovered bugs in previous NDB Cluster releases.

**Obtaining MySQL NDB Cluster 7.3.** MySQL NDB Cluster 7.3 source code and binaries can be obtained from <http://dev.mysql.com/downloads/cluster/>.

For an overview of changes made in MySQL NDB Cluster 7.3, see [What is New in NDB Cluster 7.3](#).

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 5.6 through MySQL 5.6.35 (see [Changes in MySQL 5.6.35 \(2016-12-12, General Availability\)](#)).

### Bugs Fixed

- `ndb_restore` did not restore tables having more than 341 columns correctly. This was due to the fact that the buffer used to hold table metadata read from `.ctl` files was of insufficient size, so that only part of the table descriptor could be read from it in such cases. This issue is fixed by increasing the size of the buffer used by `ndb_restore` for file reads. (Bug #25182956)

References: See also: Bug #25302901.

- The `rand()` function was used to produce a unique table ID and table version needed to identify a schema operation distributed between multiple SQL nodes, relying on the assumption that `rand()` would never produce the same numbers on two different instances of `mysqld`. It was later determined that this is not the case, and that in fact it is very likely for the same random numbers to be produced on all SQL nodes.

This fix removes the usage of `rand()` for producing a unique table ID or version, and instead uses a sequence in combination with the node ID of the coordinator. This guarantees uniqueness until the counter for the sequence wraps, which should be sufficient for this purpose.

The effects of this duplication could be observed as timeouts in the log (for example `NDB create db: waiting max 119 sec for distributing`) when restarting multiple `mysqld` processes simultaneously or nearly so, or when issuing the same `CREATE DATABASE` or `DROP DATABASE` statement on multiple SQL nodes. (Bug #24926009)

- Long message buffer exhaustion when firing immediate triggers could result in row ID leaks; this could later result in persistent `RowId already allocated` errors (`NDB Error 899`). (Bug #23723110)

References: See also: Bug #19506859, Bug #13927679.

- when a parent `NDB` table in a foreign key relationship was updated, the update cascaded to a child table as expected, but the change was not cascaded to a child table of this child table (that is, to a grandchild of the original parent). This can be illustrated using the tables generated by the following `CREATE TABLE` statements:



```

CREATE TABLE parent(
  id INT PRIMARY KEY AUTO_INCREMENT,
  coll INT UNIQUE,
  col2 INT
) ENGINE NDB;

CREATE TABLE child(
  ref1 INT UNIQUE,
  FOREIGN KEY fk1(ref1)
    REFERENCES parent(coll) ON UPDATE CASCADE
) ENGINE NDB;

CREATE TABLE grandchild(
  ref2 INT,
  FOREIGN KEY fk2(ref2)
    REFERENCES child(ref1) ON UPDATE CASCADE
) ENGINE NDB;

```

Table `child` is a child of table `parent`; table `grandchild` is a child of table `child`, and a grandchild of `parent`. In this scenario, a change to column `coll` of `parent` cascaded to `ref1` in table `child`, but it was not always propagated in turn to `ref2` in table `grandchild`. (Bug #83743, Bug #25063506)

## Changes in MySQL NDB Cluster 7.3.15 (5.6.34-ndb-7.3.15) (2016-10-18, General Availability)

MySQL NDB Cluster 7.3.15 is a new release of NDB Cluster, based on MySQL Server 5.6 and including features from version 7.3 of the NDB storage engine, as well as fixing a number of recently discovered bugs in previous NDB Cluster releases.

**Obtaining MySQL NDB Cluster 7.3.** MySQL NDB Cluster 7.3 source code and binaries can be obtained from <http://dev.mysql.com/downloads/cluster/>.

For an overview of changes made in MySQL NDB Cluster 7.3, see [What is New in NDB Cluster 7.3](#).

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 5.6 through MySQL 5.6.34 (see [Changes in MySQL 5.6.34 \(2016-10-12, General Availability\)](#)).

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

### Functionality Added or Changed

- **MySQL NDB ClusterJ:** To help applications handle database errors better, a number of new features have been added to the `ClusterJDatastoreException` class:
  - A new method, `getCode()`, returns `code` from the `NdbError` object.
  - A new method, `getMysqlCode()`, returns `mysql_code` from the `NdbError` object.
  - A new subclass, `ClusterJDatastoreException.Classification`, gives users the ability to decode the result from `getClassification()`. The method `Classification.toString()` gives the name of the error classification as listed in [NDB Error Classifications](#).

(Bug #22353594)

## Bugs Fixed

- **NDB Cluster APIs:** Reuse of transaction IDs could occur when `Ndb` objects were created and deleted concurrently. As part of this fix, the NDB API methods `lock_ndb_objects()` and `unlock_ndb_objects` are now declared as `const`. (Bug #23709232)
- Removed an invalid assertion to the effect that all cascading child scans are closed at the time API connection records are released following an abort of the main transaction. The assertion was invalid because closing of scans in such cases is by design asynchronous with respect to the main transaction, which means that subscans may well take some time to close after the main transaction is closed. (Bug #23709284)
- A number of potential buffer overflow issues were found and fixed in the `NDB` codebase. (Bug #23152979)
- When a data node has insufficient redo buffer during a system restart, it does not participate in the restart until after the other nodes have started. After this, it performs a takeover of its fragments from the nodes in its node group that have already started; during this time, the cluster is already running and user activity is possible, including DML and DDL operations.

During a system restart, table creation is handled differently in the `DIH` kernel block than normally, as this creation actually consists of reloading table definition data from disk on the master node. Thus, `DIH` assumed that any table creation that occurred before all nodes had restarted must be related to the restart and thus always on the master node. However, during the takeover, table creation can occur on non-master nodes due to user activity; when this happened, the cluster underwent a forced shutdown.

Now an extra check is made during system restarts to detect in such cases whether the executing node is the master node, and use that information to determine whether the table creation is part of the restart proper, or is taking place during a subsequent takeover. (Bug #23028418)

- When restoring a backup taken from a database containing tables that had foreign keys, `ndb_restore` disabled the foreign keys for data, but not for the logs. (Bug #83155, Bug #24736950)
- Several object constructors and similar functions in the `NDB` codebase did not always perform sanity checks when creating new instances. These checks are now performed under such circumstances. (Bug #77408, Bug #21286722)
- An internal call to `malloc()` was not checked for `NULL`. The function call was replaced with a direct write. (Bug #77375, Bug #21271194)

## Changes in MySQL NDB Cluster 7.3.14 (5.6.31-ndb-7.3.14) (2016-07-18, General Availability)

MySQL NDB Cluster 7.3.14 is a new release of NDB Cluster, based on MySQL Server 5.6 and including features from version 7.3 of the `NDB` storage engine, as well as fixing a number of recently discovered bugs in previous NDB Cluster releases.

**Obtaining MySQL NDB Cluster 7.3.** MySQL NDB Cluster 7.3 source code and binaries can be obtained from <http://dev.mysql.com/downloads/cluster/>.

For an overview of changes made in MySQL NDB Cluster 7.3, see [What is New in NDB Cluster 7.3](#).

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 5.6 through MySQL 5.6.31 (see [Changes in MySQL 5.6.31 \(2016-06-02, General Availability\)](#)).

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

## Functionality Added or Changed

- **MySQL NDB ClusterJ:** To make it easier for ClusterJ to handle fatal errors that require the `SessionFactory` to be closed, a new public method in the `SessionFactory` interface, `getConnectionPoolSessionCounts()`, has been created. When it returns zeros for all pooled connections, it means all sessions have been closed, at which point the `SessionFactory` can be closed and reopened. See [Error Handling and Reconnection](#) for more detail. (Bug #22353594)

## Bugs Fixed

- **Incompatible Change:** When the data nodes are only partially connected to the API nodes, a node used for a pushdown join may get its request from a transaction coordinator on a different node, without (yet) being connected to the API node itself. In such cases, the `NodeInfo` object for the requesting API node contained no valid info about the software version of the API node, which caused the `DBSPJ` block to assume (incorrectly) when aborting to assume that the API node used NDB version 7.2.4 or earlier, requiring the use of a backward compatibility mode to be used during query abort which sent a node failure error instead of the real error causing the abort.

Now, whenever this situation occurs, it is assumed that, if the NDB software version is not yet available, the API node version is greater than 7.2.4. (Bug #23049170)

- **NDB Cluster APIs:** Deletion of Ndb objects used a disproportionately high amount of CPU. (Bug #22986823)
- Reserved send buffer for the loopback transporter, introduced in MySQL NDB Cluster 7.4.8 and used by API and management nodes for administrative signals, was calculated incorrectly. (Bug #23093656, Bug #22016081)

References: This issue is a regression of: Bug #21664515.

- During a node restart, re-creation of internal triggers used for verifying the referential integrity of foreign keys was not reliable, because it was possible that not all distributed TC and LDM instances agreed on all trigger identities. To fix this problem, an extra step is added to the node restart sequence, during which the trigger identities are determined by querying the current master node. (Bug #23068914)

References: See also: Bug #23221573.

- Following the forced shutdown of one of the 2 data nodes in a cluster where `NoOfReplicas=2`, the other data node shut down as well, due to arbitration failure. (Bug #23006431)
- `ClusterMgr` is an internal component of NDB API and `ndb_mgmd` processes, part of `TransporterFacade`—which in turn is a wrapper around the transporter registry—and shared with data nodes. This component is responsible for a number of tasks including connection setup requests; sending and monitoring of heartbeats; provision of node state information; handling of cluster disconnects and reconnects; and forwarding of cluster state indicators. `ClusterMgr` maintains a count of live nodes which is incremented on receiving a report of a node having connected (`reportConnected()` method call), and decremented on receiving a report that a node has disconnected (`reportDisconnected()`) from `TransporterRegistry`. This count is checked within `reportDisconnected()` to verify that it is greater than zero.

The issue addressed here arose when node connections were very brief due to send buffer exhaustion (among other potential causes) and the check just described failed. This occurred

because, when a node did not fully connect, it was still possible for the connection attempt to trigger a `reportDisconnected()` call in spite of the fact that the connection had not yet been reported to `ClusterMgr`; thus, the pairing of `reportConnected()` and `reportDisconnected()` calls was not guaranteed, which could cause the count of connected nodes to be set to zero even though there remained nodes that were still in fact connected, causing node crashes with debug builds of MySQL NDB Cluster, and potential errors or other adverse effects with release builds.

To fix this issue, `ClusterMgr::reportDisconnected()` now verifies that a disconnected node had actually finished connecting completely before checking and decrementing the number of connected nodes. (Bug #21683144, Bug #22016081)

References: See also: Bug #21664515, Bug #21651400.

- To reduce the possibility that a node's loopback transporter becomes disconnected from the transporter registry by `reportError()` due to send buffer exhaustion (implemented by the fix for Bug #21651400), a portion of the send buffer is now reserved for the use of this transporter. (Bug #21664515, Bug #22016081)

References: See also: Bug #21651400, Bug #21683144.

- The loopback transporter is similar to the TCP transporter, but is used by a node to send signals to itself as part of many internal operations. Like the TCP transporter, it could be disconnected due to certain conditions including send buffer exhaustion, but this could result in blocking of `TransporterFacade` and so cause multiple issues within an `ndb_mgmd` or API node process. To prevent this, a node whose loopback transporter becomes disconnected is now simply shut down, rather than allowing the node process to hang. (Bug #21651400, Bug #22016081)

References: See also: Bug #21683144, Bug #21664515.

## Changes in MySQL NDB Cluster 7.3.13 (5.6.29-ndb-7.3.13) (2016-04-18, General Availability)

MySQL NDB Cluster 7.3.13 is a new release of NDB Cluster, based on MySQL Server 5.6 and including features from version 7.3 of the `NDB` storage engine, as well as fixing a number of recently discovered bugs in previous NDB Cluster releases.

**Obtaining MySQL NDB Cluster 7.3.** MySQL NDB Cluster 7.3 source code and binaries can be obtained from <http://dev.mysql.com/downloads/cluster/>.

For an overview of changes made in MySQL NDB Cluster 7.3, see [What is New in NDB Cluster 7.3](#).

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 5.6 through MySQL 5.6.29 (see [Changes in MySQL 5.6.29 \(2016-02-05, General Availability\)](#)).

### Bugs Fixed

- **NDB Cluster APIs:** Executing a transaction with an `NdbIndexOperation` based on an obsolete unique index caused the data node process to fail. Now the index is checked in such cases, and if it cannot be used the transaction fails with an appropriate error. (Bug #79494, Bug #22299443)
- During node failure handling, the request structure used to drive the cleanup operation was not maintained correctly when the request was executed. This led to inconsistencies that were harmless during normal operation, but these could lead to assertion failures during node failure handling, with subsequent failure of additional nodes. (Bug #22643129)

- The previous fix for a lack of mutex protection for the internal `TransporterFacade::deliver_signal()` function was found to be incomplete in some cases. (Bug #22615274)

References: This issue is a regression of: Bug #77225, Bug #21185585.

- When setup of the binary log as an atomic operation on one SQL node failed, this could trigger a state in other SQL nodes in which they appeared to detect the SQL node participating in schema change distribution, whereas it had not yet completed binary log setup. This could in turn cause a deadlock on the global metadata lock when the SQL node still retrying binary log setup needed this lock, while another mysqld had taken the lock for itself as part of a schema change operation. In such cases, the second SQL node waited for the first one to act on its schema distribution changes, which it was not yet able to do. (Bug #22494024)
- Duplicate key errors could occur when `ndb_restore` was run on a backup containing a unique index. This was due to the fact that, during restoration of data, the database can pass through one or more inconsistent states prior to completion, such an inconsistent state possibly having duplicate values for a column which has a unique index. (If the restoration of data is preceded by a run with `--disable-indexes` and followed by one with `--rebuild-indexes`, these errors are avoided.)

Added a check for unique indexes in the backup which is performed only when restoring data, and which does not process tables that have explicitly been excluded. For each unique index found, a warning is now printed. (Bug #22329365)

- Restoration of metadata with `ndb_restore -m` occasionally failed with the error message `Failed to create index...` when creating a unique index. While diagnosing this problem, it was found that the internal error `PREPARE_SEIZE_ERROR` (a temporary error) was reported as an unknown error. Now in such cases, `ndb_restore` retries the creation of the unique index, and `PREPARE_SEIZE_ERROR` is reported as NDB Error 748 `Busy during read of event table`. (Bug #21178339)

References: See also: Bug #22989944.

- Optimization of signal sending by buffering and sending them periodically, or when the buffer became full, could cause `SUB_GCP_COMPLETE_ACK` signals to be excessively delayed. Such signals are sent for each node and epoch, with a minimum interval of `TimeBetweenEpochs`; if they are not received in time, the `SUMA` buffers can overflow as a result. The overflow caused API nodes to be disconnected, leading to current transactions being aborted due to node failure. This condition made it difficult for long transactions (such as altering a very large table), to be completed. Now in such cases, the `ACK` signal is sent without being delayed. (Bug #18753341)

## Changes in MySQL NDB Cluster 7.3.12 (5.6.28-ndb-7.3.12) (2016-01-19, General Availability)

MySQL NDB Cluster 7.3.12 is a new release of NDB Cluster, based on MySQL Server 5.6 and including features from version 7.3 of the NDB storage engine, as well as fixing a number of recently discovered bugs in previous NDB Cluster releases.

**Obtaining MySQL NDB Cluster 7.3.** MySQL NDB Cluster 7.3 source code and binaries can be obtained from <http://dev.mysql.com/downloads/cluster/>.

For an overview of changes made in MySQL NDB Cluster 7.3, see [What is New in NDB Cluster 7.3](#).

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 5.6 through MySQL 5.6.28 (see [Changes in MySQL 5.6.28 \(2015-12-07, General Availability\)](#)).

## Bugs Fixed

- **Important Change:** A fix made in MySQL NDB Cluster 7.3.11 and MySQL NDB Cluster 7.4.8 caused `ndb_restore` to perform unique key checks even when operating in modes which do not restore data, such as when using the program's `--restore_epoch` or `--print_data` option.

That change in behavior caused existing valid backup routines to fail; to keep this issue from affecting this and future releases, the previous fix has been reverted. This means that the requirement added in those versions that `ndb_restore` be run `--disable-indexes` or `--rebuild-indexes` when used on tables containing unique indexes is also lifted. (Bug #22345748)

References: See also: Bug #22329365. Reverted patches: Bug #57782, Bug #11764893.

- **Important Note:** If an NDB table having a foreign key was dropped while one of the data nodes was stopped, the data node later failed when trying to restart. (Bug #18554390)
- **NDB Replication:** While the binary log injector thread was handling failure events, it was possible for all NDB tables to be left indefinitely in read-only mode. This was due to a race condition between the binary log injector thread and the utility thread handling events on the `ndb_schema` table, and to the fact that, when handling failure events, the binary log injector thread places all NDB tables in read-only mode until all such events are handled and the thread restarts itself.

When the binary log inject thread receives a group of one or more failure events, it drops all other existing event operations and expects no more events from the utility thread until it has handled all of the failure events and then restarted itself. However, it was possible for the utility thread to continue attempting binary log setup while the injector thread was handling failures and thus attempting to create the schema distribution tables as well as event subscriptions on these tables. If the creation of these tables and event subscriptions occurred during this time, the binary log injector thread's expectation that there were no further event operations was never met; thus, the injector thread never restarted, and NDB tables remained in read-only as described previously.

To fix this problem, the `Ndb` object that handles schema events is now definitely dropped once the `ndb_schema` table drop event is handled, so that the utility thread cannot create any new events until after the injector thread has restarted, at which time, a new `Ndb` object for handling schema events is created. (Bug #17674771, Bug #19537961, Bug #22204186, Bug #22361695)

- **NDB Cluster APIs:** The binary log injector did not work correctly with `TE_INCONSISTENT` event type handling by `Ndb::nextEvent()`. (Bug #22135541)

References: See also: Bug #20646496.

- **NDB Cluster APIs:** `Ndb::pollEvents()` and `pollEvents2()` were slow to receive events, being dependent on other client threads or blocks to perform polling of transporters on their behalf. This fix allows a client thread to perform its own transporter polling when it has to wait in either of these methods.

Introduction of transporter polling also revealed a problem with missing mutex protection in the `ndbcluster_binlog` handler, which has been added as part of this fix. (Bug #79311, Bug #20957068, Bug #22224571)

- In debug builds, a `WAIT_EVENT` while polling caused excessive logging to stdout. (Bug #22203672)
- When executing a schema operation such as `CREATE TABLE` on a MySQL NDB Cluster with multiple SQL nodes, it was possible for the SQL node on which the operation was performed to time out while waiting for an acknowledgement from the others. This could occur when different SQL nodes had

different settings for `--ndb-log-updated-only`, `--ndb-log-update-as-write`, or other `mysqld` options effecting binary logging by NDB.

This happened due to the fact that, in order to distribute schema changes between them, all SQL nodes subscribe to changes in the `ndb_schema` system table, and that all SQL nodes are made aware of each others subscriptions by subscribing to `TE_SUBSCRIBE` and `TE_UNSUBSCRIBE` events. The names of events to subscribe to are constructed from the table names, adding `REPL$` or `REPLF$` as a prefix. `REPLF$` is used when full binary logging is specified for the table. The issue described previously arose because different values for the options mentioned could lead to different events being subscribed to by different SQL nodes, meaning that all SQL nodes were not necessarily aware of each other, so that the code that handled waiting for schema distribution to complete did not work as designed.

To fix this issue, MySQL NDB Cluster now treats the `ndb_schema` table as a special case and enforces full binary logging at all times for this table, independent of any settings for `mysqld` binary logging options. (Bug #22174287, Bug #79188)

- Attempting to create an NDB table having greater than the maximum supported combined width for all `BIT` columns (4096) caused data node failure when these columns were defined with `COLUMN_FORMAT DYNAMIC`. (Bug #21889267)
- Creating a table with the maximum supported number of columns (512) all using `COLUMN_FORMAT DYNAMIC` led to data node failures. (Bug #21863798)
- Using `ndb_mgm STOP -f` to force a node shutdown even when it triggered a complete shutdown of the cluster, it was possible to lose data when a sufficient number of nodes were shut down, triggering a cluster shutdown, and the timing was such that `SUMA` handovers had been made to nodes already in the process of shutting down. (Bug #17772138)
- The internal `NdbEventBuffer::set_total_buckets()` method calculated the number of remaining buckets incorrectly. This caused any incomplete epoch to be prematurely completed when the `SUB_START_CONF` signal arrived out of order. Any events belonging to this epoch arriving later were then ignored, and so effectively lost, which resulted in schema changes not being distributed correctly among SQL nodes. (Bug #79635, Bug #22363510)
- Compilation of MySQL NDB Cluster failed on SUSE Linux Enterprise Server 12. (Bug #79429, Bug #22292329)
- Schema events were appended to the binary log out of order relative to non-schema events. This was caused by the fact that the binary log injector did not properly handle the case where schema events and non-schema events were from different epochs.

This fix modifies the handling of events from the two schema and non-schema event streams such that events are now always handled one epoch at a time, starting with events from the oldest available epoch, without regard to the event stream in which they occur. (Bug #79077, Bug #22135584, Bug #20456664)

- NDB failed during a node restart due to the status of the current local checkpoint being set but not as active, even though it could have other states under such conditions. (Bug #78780, Bug #21973758)
- The value set for `spintime` by the `ThreadConfig` parameter was not calculated correctly, causing the spin to continue for longer than actually specified. (Bug #78525, Bug #21886476)

## Changes in MySQL NDB Cluster 7.3.11 (5.6.27-ndb-7.3.11) (2015-10-19, General Availability)

MySQL NDB Cluster 7.3.11 is a new release of NDB Cluster, based on MySQL Server 5.6 and including features from version 7.3 of the [NDB](#) storage engine, as well as fixing a number of recently discovered bugs in previous NDB Cluster releases.

**Obtaining MySQL NDB Cluster 7.3.** MySQL NDB Cluster 7.3 source code and binaries can be obtained from <http://dev.mysql.com/downloads/cluster/>.

For an overview of changes made in MySQL NDB Cluster 7.3, see [What is New in NDB Cluster 7.3](#).

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 5.6 through MySQL 5.6.27 (see [Changes in MySQL 5.6.27 \(2015-09-30, General Availability\)](#)).

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

### Functionality Added or Changed

- **Important Change; NDB Replication:** Added the `create_old_temporals` server system variable to compliment the system variables `avoid_temporal_upgrade` and `show_old_temporals` introduced in MySQL 5.6.24 and available in MySQL NDB Cluster beginning with NDB 7.3.9 and NDB 7.4.6. Enabling `create_old_temporals` causes `mysqld` to use the storage format employed prior to MySQL 5.6.4 when creating any `DATE`, `DATETIME`, or `TIMESTAMP` column—that is, the column is created without any support for fractional seconds. `create_old_temporals` is disabled by default. The system variable is read-only; to enable the use of pre-5.6.4 temporal types, set the equivalent option (`--create-old-temporals`) on the command line, or in an option file read by the MySQL server.

`create_old_temporals` is available only in MySQL NDB Cluster; it is not supported in the standard MySQL 5.6 server. It is intended to facilitate upgrades from MySQL NDB Cluster 7.2 to MySQL NDB Cluster 7.3 and 7.4, after which table columns of the affected types can be upgraded to the new storage format. `create_old_temporals` is deprecated and scheduled for removal in a future version of MySQL NDB Cluster.

`avoid_temporal_upgrade` must also be enabled for this feature to work properly. You should also enable `show_old_temporals` as well. For more information, see the descriptions of these variables. For more about the changes in MySQL's temporal types, see [Date and Time Type Storage Requirements](#). (Bug #20701918)

References: See also: Bug #21492598, Bug #72997, Bug #18985760.

### Bugs Fixed

- **Important Change:** When `ndb_restore` was run without `--disable-indexes` or `--rebuild-indexes` on a table having a unique index, it was possible for rows to be restored in an order that resulted in duplicate values, causing it to fail with duplicate key errors. Running `ndb_restore` on such a table now requires using at least one of these options; failing to do so now results in an error. (Bug #57782, Bug #11764893)

References: See also: Bug #22329365, Bug #22345748.

- **NDB Cluster APIs:** While executing `dropEvent()`, if the coordinator `DBDICT` failed after the subscription manager (`SUMA` block) had removed all subscriptions but before the coordinator had deleted



the event from the system table, the dropped event remained in the table, causing any subsequent drop or create event with the same name to fail with NDB error 1419 `Subscription already dropped` or error 746 `Event name already exists`. This occurred even when calling `dropEvent()` with a nonzero force argument.

Now in such cases, error 1419 is ignored, and `DBDICT` deletes the event from the table. (Bug #21554676)

- **NDB Cluster APIs:** The internal value representing the latest global checkpoint was not always updated when a completed epoch of event buffers was inserted into the event queue. This caused subsequent calls to `Ndb::pollEvents()` and `pollEvents2()` to fail when trying to obtain the correct GCI for the events available in the event buffers. This could also result in later calls to `nextEvent()` or `nextEvent2()` seeing events that had not yet been discovered. (Bug #78129, Bug #21651536)
- Backup block states were reported incorrectly during backups. (Bug #21360188)

References: See also: Bug #20204854, Bug #21372136.

- When a data node is known to have been alive by other nodes in the cluster at a given global checkpoint, but its `sysfile` reports a lower GCI, the higher GCI is used to determine which global checkpoint the data node can recreate. This caused problems when the data node being started had a clean file system (GCI = 0), or when it was more than more global checkpoint behind the other nodes.

Now in such cases a higher GCI known by other nodes is used only when it is at most one GCI ahead. (Bug #19633824)

References: See also: Bug #20334650, Bug #21899993. This issue is a regression of: Bug #29167.

- When restoring a specific database or databases with the `--include-databases` or `--exclude-databases` option, `ndb_restore` attempted to apply foreign keys on tables in databases which were not among those being restored. (Bug #18560951)
- After restoring the database schema from backup using `ndb_restore`, auto-discovery of restored tables in transactions having multiple statements did not work correctly, resulting in `Deadlock found when trying to get lock; try restarting transaction` errors.

This issue was encountered both in the `mysql` client, as well as when such transactions were executed by application programs using Connector/J and possibly other MySQL APIs.

Prior to upgrading, this issue can be worked around by executing `SELECT TABLE_NAME, TABLE_SCHEMA FROM INFORMATION_SCHEMA.TABLES WHERE ENGINE = 'NDBCLUSTER'` on all SQL nodes following the restore operation, before executing any other statements. (Bug #18075170)

- `ndb_desc` used with the `--extra-partition-info` and `--blob-info` options failed when run against a table containing one or more `TINYBLOB` columns. (Bug #14695968)
- Trying to create an NDB table with a composite foreign key referencing a composite primary key of the parent table failed when one of the columns in the composite foreign key was the table's primary key and in addition this column also had a unique key. (Bug #78150, Bug #21664899)
- When attempting to enable index statistics, creation of the required system tables, events and event subscriptions often fails when multiple `mysqld` processes using index statistics are started concurrently in conjunction with starting, restarting, or stopping the cluster, or with node failure handling. This is normally recoverable, since the affected `mysqld` process or processes can (and do) retry these operations shortly thereafter. For this reason, such failures are no longer logged as warnings, but merely as informational events. (Bug #77760, Bug #21462846)

- Adding a unique key to an [NDB](#) table failed when the table already had a foreign key. Prior to upgrading, you can work around this issue by creating the unique key first, then adding the foreign key afterwards, using a separate `ALTER TABLE` statement. (Bug #77457, Bug #20309828)

## Changes in MySQL NDB Cluster 7.3.10 (5.6.25-ndb-7.3.10) (2015-07-13, General Availability)

MySQL NDB Cluster 7.3.10 is a new release of NDB Cluster, based on MySQL Server 5.6 and including features from version 7.3 of the [NDB](#) storage engine, as well as fixing a number of recently discovered bugs in previous NDB Cluster releases.

**Obtaining MySQL NDB Cluster 7.3.** MySQL NDB Cluster 7.3 source code and binaries can be obtained from <http://dev.mysql.com/downloads/cluster/>.

For an overview of changes made in MySQL NDB Cluster 7.3, see [What is New in NDB Cluster 7.3](#).

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 5.6 through MySQL 5.6.25 (see [Changes in MySQL 5.6.25 \(2015-05-29, General Availability\)](#)).

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

### Functionality Added or Changed

- **MySQL NDB ClusterJ:** Under high workload, it was possible to overload the direct memory used to back domain objects, because direct memory is not garbage collected in the same manner as objects allocated on the heap. Two strategies have been added to the ClusterJ implementation: first, direct memory is now pooled, so that when the domain object is garbage collected, the direct memory can be reused by another domain object. Additionally, a new user-level method, `release(instance)`, has been added to the `Session` interface, which allows users to release the direct memory before the corresponding domain object is garbage collected. See the description for `release(T)` for more information. (Bug #20504741)
- A number of improvements, listed here, have been made with regard to handling issues that could arise when an overload arose due to a great number of inserts being performed during a local checkpoint (LCP):
  - Failures sometimes occurred during restart processing when trying to execute the undo log, due to a problem with finding the end of the log. This happened when there remained unwritten pages at the end of the first undo file when writing to the second undo file, which caused the execution of undo logs in reverse order and so execute old or even nonexistent log records.

This is fixed by ensuring that execution of the undo log begins with the proper end of the log, and, if started earlier, that any unwritten or faulty pages are ignored.

- It was possible to fail during an LCP, or when performing a `COPY_FRAGREQ`, due to running out of operation records. We fix this by making sure that LCPs and `COPY_FRAG` use resources reserved for operation records, as was already the case with scan records. In addition, old code for ACC operations that was no longer required but that could lead to failures was removed.
- When an LCP was performed while loading a table, it was possible to hit a livelock during LCP scans, due to the fact that each record that was inserted into new pages after the LCP had started had its `LCP_SKIP` flag set. Such records were discarded as intended by the LCP scan, but when inserts occurred faster than the LCP scan could discard records, the scan appeared to hang. As part of this

issue, the scan failed to report any progress to the LCP watchdog, which after 70 seconds of livelock killed the process. This issue was observed when performing on the order of 250000 inserts per second over an extended period of time (120 seconds or more), using a single LDM.

This part of the fix makes a number of changes, listed here:

- We now ensure that pages created after the LCP has started are not included in LCP scans; we also ensure that no records inserted into those pages have their `LCP_SKIP` flag set.
- Handling of the scan protocol is changed such that a certain amount of progress is made by the LCP regardless of load; we now report progress to the LCP watchdog so that we avoid failure in the event that an LCP is making progress but not writing any records.
- We now take steps to guarantee that LCP scans proceed more quickly than inserts can occur, by ensuring that scans are prioritized this scanning activity, and thus, that the LCP is in fact (eventually) completed.
- In addition, scanning is made more efficient, by prefetching tuples; this helps avoid stalls while fetching memory in the CPU.
- Row checksums for preventing data corruption now include the tuple header bits.

(Bug #76373, Bug #20727343, Bug #76741, Bug #69994, Bug #20903880, Bug #76742, Bug #20904721, Bug #76883, Bug #20980229)

## Bugs Fixed

- **Important Change; NDB Cluster APIs:** Added the method `Ndb::isExpectingHigherQueuedEpochs()` to the NDB API to detect when additional, newer event epochs were detected by `pollEvents2()`.

The behavior of `Ndb::pollEvents()` has also been modified such that it now returns `NDB_FAILURE_GCI` (equal to `~(Uint64) 0`) when a cluster failure has been detected. (Bug #18753887)

- **NDB Cluster APIs:** Added the `Column::getSizeInBytesForRecord()` method, which returns the size required for a column by an `NdbRecord`, depending on the column's type (text/blob, or other). (Bug #21067283)
- **NDB Cluster APIs:** Creation and destruction of `Ndb_cluster_connection` objects by multiple threads could make use of the same application lock, which in some cases led to failures in the global dictionary cache. To alleviate this problem, the creation and destruction of several internal NDB API objects have been serialized. (Bug #20636124)
- **NDB Cluster APIs:** A number of timeouts were not handled correctly in the NDB API. (Bug #20617891)
- **NDB Cluster APIs:** When an `Ndb` object created prior to a failure of the cluster was reused, the event queue of this object could still contain data node events originating from before the failure. These events could reference "old" epochs (from before the failure occurred), which in turn could violate the assumption made by the `nextEvent()` method that epoch numbers always increase. This issue is addressed by explicitly clearing the event queue in such cases. (Bug #18411034)

References: See also: Bug #20888668.

- **MySQL NDB ClusterJ:** When used with Java 1.7 or higher, ClusterJ might cause the Java VM to crash when querying tables with BLOB columns, because `NdbDictionary::createRecord` calculates the wrong size needed for the record. Subsequently, when ClusterJ called

`NdbScanOperation::nextRecordCopyOut`, the data overran the allocated buffer space. With this fix, ClusterJ checks the size calculated by `NdbDictionary::createRecord` and uses the value for the buffer size, if it is larger than the value ClusterJ itself calculates. (Bug #20695155)

- After restoring the database metadata (but not any data) by running `ndb_restore --restore_meta` (or `-m`), SQL nodes would hang while trying to `SELECT` from a table in the database to which the metadata was restored. In such cases the attempt to query the table now fails as expected, since the table does not actually exist until `ndb_restore` is executed with `--restore_data (-r)`. (Bug #21184102)

References: See also: Bug #16890703.

- When a great many threads opened and closed blocks in the NDB API in rapid succession, the internal `close_clnt()` function synchronizing the closing of the blocks waited an insufficiently long time for a self-signal indicating potential additional signals needing to be processed. This led to excessive CPU usage by `ndb_mgmd`, and prevented other threads from opening or closing other blocks. This issue is fixed by changing the function polling call to wait on a specific condition to be woken up (that is, when a signal has in fact been executed). (Bug #21141495)
- Previously, multiple send threads could be invoked for handling sends to the same node; these threads then competed for the same send lock. While the send lock blocked the additional send threads, work threads could be passed to other nodes.

This issue is fixed by ensuring that new send threads are not activated while there is already an active send thread assigned to the same node. In addition, a node already having an active send thread assigned to it is no longer visible to other, already active, send threads; that is, such a node is longer added to the node list when a send thread is currently assigned to it. (Bug #20954804, Bug #76821)

- Queueing of pending operations when the redo log was overloaded (`DefaultOperationRedoProblemAction` API node configuration parameter) could lead to timeouts when data nodes ran out of redo log space (`P_TAIL_PROBLEM` errors). Now when the redo log is full, the node aborts requests instead of queuing them. (Bug #20782580)

References: See also: Bug #20481140.

- NDB statistics queries could be delayed by the error delay set for `ndb_index_stat_option` (default 60 seconds) when the index that was queried had been marked with internal error. The same underlying issue could also cause `ANALYZE TABLE` to hang when executed against an NDB table having multiple indexes where an internal error occurred on one or more but not all indexes.

Now in such cases, any existing statistics are returned immediately, without waiting for any additional statistics to be discovered. (Bug #20553313, Bug #20707694, Bug #76325)

- The multithreaded scheduler sends to remote nodes either directly from each worker thread or from dedicated send threadsL, depending on the cluster's configuration. This send might transmit all, part, or none of the available data from the send buffers. While there remained pending send data, the worker or send threads continued trying to send in a loop. The actual size of the data sent in the most recent attempt to perform a send is now tracked, and used to detect lack of send progress by the send or worker threads. When no progress has been made, and there is no other work outstanding, the scheduler takes a 1 millisecond pause to free up the CPU for use by other threads. (Bug #18390321)

References: See also: Bug #20929176, Bug #20954804.

- In some cases, attempting to restore a table that was previously backed up failed with a `File Not Found` error due to a missing table fragment file. This occurred as a result of the NDB kernel `BACKUP` block receiving a `Busy` error while trying to obtain the table description, due to other traffic from external clients, and not retrying the operation.

The fix for this issue creates two separate queues for such requests—one for internal clients such as the `BACKUP` block or `ndb_restore`, and one for external clients such as API nodes—and prioritizing the internal queue.

Note that it has always been the case that external client applications using the NDB API (including MySQL applications running against an SQL node) are expected to handle `Busy` errors by retrying transactions at a later time; this expectation is *not* changed by the fix for this issue. (Bug #17878183)

References: See also: Bug #17916243.

- In some cases, the `DBDICT` block failed to handle repeated `GET_TABINFOREQ` signals after the first one, leading to possible node failures and restarts. This could be observed after setting a sufficiently high value for `MaxNoOfExecutionThreads` and low value for `LcpScanProgressTimeout`. (Bug #77433, Bug #21297221)
- Client lookup for delivery of API signals to the correct client by the internal `TransporterFacade::deliver_signal()` function had no mutex protection, which could cause issues such as timeouts encountered during testing, when other clients connected to the same `TransporterFacade`. (Bug #77225, Bug #21185585)
- It was possible to end up with a lock on the send buffer mutex when send buffers became a limiting resource, due either to insufficient send buffer resource configuration, problems with slow or failing communications such that all send buffers became exhausted, or slow receivers failing to consume what was sent. In this situation worker threads failed to allocate send buffer memory for signals, and attempted to force a send in order to free up space, while at the same time the send thread was busy trying to send to the same node or nodes. All of these threads competed for taking the send buffer mutex, which resulted in the lock already described, reported by the watchdog as `Stuck in Send`. This fix is made in two parts, listed here:
  1. The send thread no longer holds the global send thread mutex while getting the send buffer mutex; it now releases the global mutex prior to locking the send buffer mutex. This keeps worker threads from getting stuck in send in such cases.
  2. Locking of the send buffer mutex done by the send threads now uses a try-lock. If the try-lock fails, the node to make the send to is reinserted at the end of the list of send nodes in order to be retried later. This removes the `Stuck in Send` condition for the send threads.

(Bug #77081, Bug #21109605)

## Changes in MySQL NDB Cluster 7.3.9 (5.6.24-ndb-7.3.9) (2015-04-14, General Availability)

MySQL NDB Cluster 7.3.9 is a new release of NDB Cluster, based on MySQL Server 5.6 and including features from version 7.3 of the `NDB` storage engine, as well as fixing a number of recently discovered bugs in previous NDB Cluster releases.

**Obtaining MySQL NDB Cluster 7.3.** MySQL NDB Cluster 7.3 source code and binaries can be obtained from <http://dev.mysql.com/downloads/cluster/>.

For an overview of changes made in MySQL NDB Cluster 7.3, see [What is New in NDB Cluster 7.3](#).

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 5.6 through MySQL 5.6.24 (see [Changes in MySQL 5.6.24 \(2015-04-06, General Availability\)](#)).

## Bugs Fixed

- **Important Change:** The maximum failure time calculation used to ensure that normal node failure handling mechanisms are given time to handle survivable cluster failures (before global checkpoint watchdog mechanisms start to kill nodes due to GCP delays) was excessively conservative, and neglected to consider that there can be at most  $\text{number\_of\_data\_nodes} / \text{NoOfReplicas}$  node failures before the cluster can no longer survive. Now the value of `NoOfReplicas` is properly taken into account when performing this calculation.

This fix adds the `TimeBetweenGlobalCheckpointsTimeout` data node configuration parameter, which makes the minimum timeout between global checkpoints settable by the user. This timeout was previously fixed internally at 120000 milliseconds, which is now the default value for this parameter. (Bug #20069617, Bug #20069624)

References: See also: Bug #19858151, Bug #20128256, Bug #20135976.

- **NDB Cluster APIs:** When a transaction is started from a cluster connection, `Table` and `Index` schema objects may be passed to this transaction for use. If these schema objects have been acquired from a different connection (`Ndb_cluster_connection` object), they can be deleted at any point by the deletion or disconnection of the owning connection. This can leave a connection with invalid schema objects, which causes an NDB API application to fail when these are dereferenced.

To avoid this problem, if your application uses multiple connections, you can now set a check to detect sharing of schema objects between connections when passing a schema object to a transaction, using the `NdbTransaction::setSchemaObjectOwnerChecks()` method added in this release. When this check is enabled, the schema objects having the same names are acquired from the connection and compared to the schema objects passed to the transaction. Failure to match causes the application to fail with an error. (Bug #19785977)

- **NDB Cluster APIs:** The increase in the default number of hashmap buckets (`DefaultHashMapSize` API node configuration parameter) from 240 to 3480 in MySQL NDB Cluster 7.2.11 increased the size of the internal `DictHashMapInfo::HashMap` type considerably. This type was allocated on the stack in some `getTable()` calls which could lead to stack overflow issues for NDB API users.

To avoid this problem, the hashmap is now dynamically allocated from the heap. (Bug #19306793)

- **NDB Cluster APIs:** A scan operation, whether it is a single table scan or a query scan used by a pushed join, stores the result set in a buffer. This maximum size of this buffer is calculated and preallocated before the scan operation is started. This buffer may consume a considerable amount of memory; in some cases we observed a 2 GB buffer footprint in tests that executed 100 parallel scans with 2 single-threaded (`ndbd`) data nodes. This memory consumption was found to scale linearly with additional fragments.

A number of root causes, listed here, were discovered that led to this problem:

- Result rows were unpacked to full `NdbRecord` format before they were stored in the buffer. If only some but not all columns of a table were selected, the buffer contained empty space (essentially wasted).
- Due to the buffer format being unpacked, `VARCHAR` and `VARBINARY` columns always had to be allocated for the maximum size defined for such columns.
- `BatchByteSize` and `MaxScanBatchSize` values were not taken into consideration as a limiting factor when calculating the maximum buffer size.

These issues became more evident in NDB 7.2 and later MySQL NDB Cluster release series. This was due to the fact buffer size is scaled by `BatchSize`, and that the default value for this parameter was increased fourfold (from 64 to 256) beginning with MySQL NDB Cluster 7.2.1.

This fix causes result rows to be buffered using the packed format instead of the unpacked format; a buffered scan result row is now not unpacked until it becomes the current row. In addition, `BatchByteSize` and `MaxScanBatchSize` are now used as limiting factors when calculating the required buffer size.

Also as part of this fix, refactoring has been done to separate handling of buffered (packed) from handling of unbuffered result sets, and to remove code that had been unused since NDB 7.0 or earlier. The `NdbRecord` class declaration has also been cleaned up by removing a number of unused or redundant member variables. (Bug #73781, Bug #75599, Bug #19631350, Bug #20408733)

- It was found during testing that problems could arise when the node registered as the arbitrator disconnected or failed during the arbitration process.

In this situation, the node requesting arbitration could never receive a positive acknowledgement from the registered arbitrator; this node also lacked a stable set of members and could not initiate selection of a new arbitrator.

Now in such cases, when the arbitrator fails or loses contact during arbitration, the requesting node immediately fails rather than waiting to time out. (Bug #20538179)

- The values of the `Ndb_last_commit_epoch_server` and `Ndb_last_commit_epoch_session` status variables were incorrectly reported on some platforms. To correct this problem, these values are now stored internally as `long long`, rather than `long`. (Bug #20372169)
- When a data node fails or is being restarted, the remaining nodes in the same nodegroup resend to subscribers any data which they determine has not already been sent by the failed node. Normally, when a data node (actually, the `SUMA` kernel block) has sent all data belonging to an epoch for which it is responsible, it sends a `SUB_GCP_COMPLETE_REP` signal, together with a count, to all subscribers, each of which responds with a `SUB_GCP_COMPLETE_ACK`. When `SUMA` receives this acknowledgment from all subscribers, it reports this to the other nodes in the same nodegroup so that they know that there is no need to resend this data in case of a subsequent node failure. If a node failed before all subscribers sent this acknowledgement but before all the other nodes in the same nodegroup received it from the failing node, data for some epochs could be sent (and reported as complete) twice, which could lead to an unplanned shutdown.

The fix for this issue adds to the count reported by `SUB_GCP_COMPLETE_ACK` a list of identifiers which the receiver can use to keep track of which buckets are completed and to ignore any duplicate reported for an already completed bucket. (Bug #17579998)

- When performing a restart, it was sometimes possible to find a log end marker which had been written by a previous restart, and that should have been invalidated. Now when searching for the last page to invalidate, the same search algorithm is used as when searching for the last page of the log to read. (Bug #76207, Bug #20665205)
- When reading and copying transporter short signal data, it was possible for the data to be copied back to the same signal with overlapping memory. (Bug #75930, Bug #20553247)
- When a bulk delete operation was committed early to avoid an additional round trip, while also returning the number of affected rows, but failed with a timeout error, an SQL node performed no verification that the transaction was in the Committed state. (Bug #74494, Bug #20092754)

References: See also: Bug #19873609.

- An `ALTER TABLE` statement containing comments and a partitioning option against an NDB table caused the SQL node on which it was executed to fail. (Bug #74022, Bug #19667566)

## Changes in MySQL NDB Cluster 7.3.8 (5.6.22-ndb-7.3.8) (2015-01-21, General Availability)

MySQL NDB Cluster 7.3.8 is a new release of NDB Cluster, based on MySQL Server 5.6 and including features from version 7.3 of the NDB storage engine, as well as fixing a number of recently discovered bugs in previous NDB Cluster releases.

**Obtaining MySQL NDB Cluster 7.3.** MySQL NDB Cluster 7.3 source code and binaries can be obtained from <http://dev.mysql.com/downloads/cluster/>.

For an overview of changes made in MySQL NDB Cluster 7.3, see [What is New in NDB Cluster 7.3](#).

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 5.6 through MySQL 5.6.22 (see [Changes in MySQL 5.6.22 \(2014-12-01, General Availability\)](#)).

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

### Functionality Added or Changed

- **Performance:** Recent improvements made to the multithreaded scheduler were intended to optimize the cache behavior of its internal data structures, with members of these structures placed such that those local to a given thread do not overflow into a cache line which can be accessed by another thread. Where required, extra padding bytes are inserted to isolate cache lines owned (or shared) by other threads, thus avoiding invalidation of the entire cache line if another thread writes into a cache line not entirely owned by itself. This optimization improved MT Scheduler performance by several percent.

It has since been found that the optimization just described depends on the global instance of struct `thr_repository` starting at a cache line aligned base address as well as the compiler not rearranging or adding extra padding to the scheduler struct; it was also found that these prerequisites were not guaranteed (or even checked). Thus this cache line optimization has previously worked only when `g_thr_repository` (that is, the global instance) ended up being cache line aligned only by accident. In addition, on 64-bit platforms, the compiler added extra padding words in struct `thr_safe_pool` such that attempts to pad it to a cache line aligned size failed.

The current fix ensures that `g_thr_repository` is constructed on a cache line aligned address, and the constructors modified so as to verify cacheline aligned addresses where these are assumed by design.

Results from internal testing show improvements in MT Scheduler read performance of up to 10% in some cases, following these changes. (Bug #18352514)

- **NDB Cluster APIs:** Two new example programs, demonstrating reads and writes of `CHAR`, `VARCHAR`, and `VARBINARY` column values, have been added to `storage/ndb/ndbapi-examples` in the MySQL NDB Cluster source tree. For more information about these programs, including source code listings, see [NDB API Simple Array Example](#), and [NDB API Simple Array Example Using Adapter](#).
- **MySQL NDB ClusterJ:** A new property, `PROPERTY_CLUSTER_CONNECT_TIMEOUT_MGM`, can now be used to set the network connection timeout. (Bug #19913569)



## Bugs Fixed

- **NDB Disk Data:** An update on many rows of a large Disk Data table could in some rare cases lead to node failure. In the event that such problems are observed with very large transactions on Disk Data tables you can now increase the number of page entries allocated for disk page buffer memory by raising the value of the `DiskPageBufferEntries` data node configuration parameter added in this release. (Bug #19958804)
- **NDB Disk Data:** In some cases, during `DICT` master takeover, the new master could crash while attempting to roll forward an ongoing schema transaction. (Bug #19875663, Bug #74510)
- **NDB Disk Data:** When a node acting as a `DICT` master fails, the arbitrator selects another node to take over in place of the failed node. During the takeover procedure, which includes cleaning up any schema transactions which are still open when the master failed, the disposition of the uncommitted schema transaction is decided. Normally this transaction be rolled back, but if it has completed a sufficient portion of a commit request, the new master finishes processing the commit. Until the fate of the transaction has been decided, no new `TRANS_END_REQ` messages from clients can be processed. In addition, since multiple concurrent schema transactions are not supported, takeover cleanup must be completed before any new transactions can be started.

A similar restriction applies to any schema operations which are performed in the scope of an open schema transaction. The counter used to coordinate schema operation across all nodes is employed both during takeover processing and when executing any non-local schema operations. This means that starting a schema operation while its schema transaction is in the takeover phase causes this counter to be overwritten by concurrent uses, with unpredictable results.

The scenarios just described were handled previously using a pseudo-random delay when recovering from a node failure. Now we check before the new master has rolled forward or backwards any schema transactions remaining after the failure of the previous master and avoid starting new schema transactions or performing operations using old transactions until takeover processing has cleaned up after the abandoned transaction. (Bug #19874809, Bug #74503)

- **NDB Disk Data:** When a node acting as `DICT` master fails, it is still possible to request that any open schema transaction be either committed or aborted by sending this request to the new `DICT` master. In this event, the new master takes over the schema transaction and reports back on whether the commit or abort request succeeded. In certain cases, it was possible for the new master to be misidentified—that is, the request was sent to the wrong node, which responded with an error that was interpreted by the client application as an aborted schema transaction, even in cases where the transaction could have been successfully committed, had the correct node been contacted. (Bug #74521, Bug #19880747)
- **NDB Replication:** When an `NDB` client thread made a request to flush the binary log using statements such as `FLUSH BINARY LOGS` or `SHOW BINLOG EVENTS`, this caused not only the most recent changes made by this client to be flushed, but all recent changes made by all other clients to be flushed as well, even though this was not needed. This behavior caused unnecessary waiting for the statement to execute, which could lead to timeouts and other issues with replication. Now such statements flush the most recent database changes made by the requesting thread only.

As part of this fix, the status variables `Ndb_last_commit_epoch_server`, `Ndb_last_commit_epoch_session`, and `Ndb_slave_max_replicated_epoch`, originally implemented in MySQL NDB Cluster 7.4, are also now available in MySQL NDB Cluster 7.3. For descriptions of these variables, see [NDB Cluster Status Variables](#); for further information, see [NDB Cluster Replication Conflict Resolution](#). (Bug #19793475)

- **NDB Replication:** It was possible using wildcards to set up conflict resolution for an exceptions table (that is, a table named using the suffix `$EX`), which should not be allowed. Now when a replication conflict function is defined using wildcard expressions, these are checked for possible matches so

that, in the event that the function would cover an exceptions table, it is not set up for this table. (Bug #19267720)

- **NDB Cluster APIs:** It was possible to delete an `Ndb_cluster_connection` object while there remained instances of `Ndb` using references to it. Now the `Ndb_cluster_connection` destructor waits for all related `Ndb` objects to be released before completing. (Bug #19999242)

References: See also: Bug #19846392.

- **NDB Cluster APIs:** The buffer allocated by an `NdbScanOperation` for receiving scanned rows was not released until the `NdbTransaction` owning the scan operation was closed. This could lead to excessive memory usage in an application where multiple scans were created within the same transaction, even if these scans were closed at the end of their lifecycle, unless `NdbScanOperation::close()` was invoked with the `releaseOp` argument equal to `true`. Now the buffer is released whenever the cursor navigating the result set is closed with `NdbScanOperation::close()`, regardless of the value of this argument. (Bug #75128, Bug #20166585)
- **MySQL NDB ClusterJ:** The following errors were logged at the `SEVERE` level; they are now logged at the `NORMAL` level, as they should be:
  - Duplicate primary key
  - Duplicate unique key
  - Foreign key constraint error: key does not exist
  - Foreign key constraint error: key exists

(Bug #20045455)

- **MySQL NDB ClusterJ:** The `com.mysql.clusterj.tie` class gave off a logging message at the `INFO` logging level for every single query, which was unnecessary and was affecting the performance of applications that used ClusterJ. (Bug #20017292)
- **MySQL NDB ClusterJ:** ClusterJ reported a segmentation violation when an application closed a session factory while some sessions were still active. This was because MySQL NDB Cluster allowed an `Ndb_cluster_connection` object be to deleted while some `Ndb` instances were still active, which might result in the usage of null pointers by ClusterJ. This fix stops that happening by preventing ClusterJ from closing a session factory when any of its sessions are still active. (Bug #19846392)

References: See also: Bug #19999242.

- The global checkpoint commit and save protocols can be delayed by various causes, including slow disk I/O. The `DIH` master node monitors the progress of both of these protocols, and can enforce a maximum lag time during which the protocols are stalled by killing the node responsible for the lag when it reaches this maximum. This `DIH` master GCP monitor mechanism did not perform its task more than once per master node; that is, it failed to continue monitoring after detecting and handling a GCP stop. (Bug #20128256)

References: See also: Bug #19858151, Bug #20069617, Bug #20062754.

- When running `mysql_upgrade` on a MySQL NDB Cluster SQL node, the expected drop of the `performance_schema` database on this node was instead performed on all SQL nodes connected to the cluster. (Bug #20032861)
- A number of problems relating to the fired triggers pool have been fixed, including the following issues:

- When the fired triggers pool was exhausted, `NDB` returned Error 218 (`Out of LongMessageBuffer`). A new error code 221 is added to cover this case.
- An additional, separate case in which Error 218 was wrongly reported now returns the correct error.
- Setting low values for `MaxNoOfFiredTriggers` led to an error when no memory was allocated if there was only one hash bucket.
- An aborted transaction now releases any fired trigger records it held. Previously, these records were held until its `ApiConnectRecord` was reused by another transaction.
- In addition, for the `Fired Triggers` pool in the internal `ndbinfo.ndb$poools` table, the high value always equalled the total, due to the fact that all records were momentarily seized when initializing them. Now the high value shows the maximum following completion of initialization.

(Bug #19976428)

- Online reorganization when using `ndbmt` data nodes and with binary logging by `mysqld` enabled could sometimes lead to failures in the `TRIX` and `DBLQH` kernel blocks, or in silent data corruption. (Bug #19903481)

References: See also: Bug #19912988.

- The local checkpoint scan fragment watchdog and the global checkpoint monitor can each exclude a node when it is too slow when participating in their respective protocols. This exclusion was implemented by simply asking the failing node to shut down, which in case this was delayed (for whatever reason) could prolong the duration of the GCP or LCP stall for other, unaffected nodes.

To minimize this time, an isolation mechanism has been added to both protocols whereby any other live nodes forcibly disconnect the failing node after a predetermined amount of time. This allows the failing node the opportunity to shut down gracefully (after logging debugging and other information) if possible, but limits the time that other nodes must wait for this to occur. Now, once the remaining live nodes have processed the disconnection of any failing nodes, they can commence failure handling and restart the related protocol or protocol, even if the failed node takes an excessively long time to shut down. (Bug #19858151)

References: See also: Bug #20128256, Bug #20069617, Bug #20062754.

- A watchdog failure resulted from a hang while freeing a disk page in `TUP_COMMITREQ`, due to use of an uninitialized block variable. (Bug #19815044, Bug #74380)
- Multiple threads crashing led to multiple sets of trace files being printed and possibly to deadlocks. (Bug #19724313)
- When a client retried against a new master a schema transaction that failed previously against the previous master while the latter was restarting, the lock obtained by this transaction on the new master prevented the previous master from progressing past start phase 3 until the client was terminated, and resources held by it were cleaned up. (Bug #19712569, Bug #74154)
- When using the `NDB` storage engine, the maximum possible length of a database or table name is 63 characters, but this limit was not always strictly enforced. This meant that a statement using a name having 64 characters such `CREATE DATABASE`, `DROP DATABASE`, or `ALTER TABLE RENAME` could cause the SQL node on which it was executed to fail. Now such statements fail with an appropriate error message. (Bug #19550973)

- When a new data node started, API nodes were allowed to attempt to register themselves with the data node for executing transactions before the data node was ready. This forced the API node to wait an extra heartbeat interval before trying again.

To address this issue, a number of `HA_ERR_NO_CONNECTION` errors (Error 4009) that could be issued during this time have been changed to `Cluster temporarily unavailable` errors (Error 4035), which should allow API nodes to use new data nodes more quickly than before. As part of this fix, some errors which were incorrectly categorised have been moved into the correct categories, and some errors which are no longer used have been removed. (Bug #19524096, Bug #73758)

- When executing very large pushdown joins involving one or more indexes each defined over several columns, it was possible in some cases for the `DBSPJ` block (see [The DBSPJ Block](#)) in the `NDB` kernel to generate `SCAN_FRAGREQ` signals that were excessively large. This caused data nodes to fail when these could not be handled correctly, due to a hard limit in the kernel on the size of such signals (32K). This fix bypasses that limitation by breaking up `SCAN_FRAGREQ` data that is too large for one such signal, and sending the `SCAN_FRAGREQ` as a chunked or fragmented signal instead. (Bug #19390895)
- `ndb_index_stat` sometimes failed when used against a table containing unique indexes. (Bug #18715165)
- Queries against tables containing a `CHAR(0)` columns failed with `ERROR 1296 (HY000): Got error 4547 'RecordSpecification has overlapping offsets' from NDBCLUSTER`. (Bug #14798022)
- In the `NDB` kernel, it was possible for a `TransporterFacade` object to reset a buffer while the data contained by the buffer was being sent, which could lead to a race condition. (Bug #75041, Bug #20112981)
- `mysql_upgrade` failed to drop and recreate the `ndbinfo` database and its tables as expected. (Bug #74863, Bug #20031425)
- Due to a lack of memory barriers, MySQL NDB Cluster programs such as `ndbmt.d` did not compile on `POWER` platforms. (Bug #74782, Bug #20007248)
- In some cases, when run against a table having an `AFTER DELETE` trigger, a `DELETE` statement that matched no rows still caused the trigger to execute. (Bug #74751, Bug #19992856)
- A basic requirement of the `NDB` storage engine's design is that the transporter registry not attempt to receive data (`TransporterRegistry::performReceive()`) from and update the connection status (`TransporterRegistry::update_connections()`) of the same set of transporters concurrently, due to the fact that the updates perform final cleanup and reinitialization of buffers used when receiving data. Changing the contents of these buffers while reading or writing to them could lead to "garbage" or inconsistent signals being read or written.

During the course of work done previously to improve the implementation of the transporter facade, a mutex intended to protect against the concurrent use of the `performReceive()` and `update_connections()` methods on the same transporter was inadvertently removed. This fix adds a watchdog check for concurrent usage. In addition, `update_connections()` and `performReceive()` calls are now serialized together while polling the transporters. (Bug #74011, Bug #19661543)

- `ndb_restore` failed while restoring a table which contained both a built-in conversion on the primary key and a staging conversion on a `TEXT` column.

During staging, a `BLOB` table is created with a primary key column of the target type. However, a conversion function was not provided to convert the primary key values before loading them into the

staging blob table, which resulted in corrupted primary key values in the staging `BLOB` table. While moving data from the staging table to the target table, the `BLOB` read failed because it could not find the primary key in the `BLOB` table.

Now all `BLOB` tables are checked to see whether there are conversions on primary keys of their main tables. This check is done after all the main tables are processed, so that conversion functions and parameters have already been set for the main tables. Any conversion functions and parameters used for the primary key in the main table are now duplicated in the `BLOB` table. (Bug #73966, Bug #19642978)

- Corrupted messages to data nodes sometimes went undetected, causing a bad signal to be delivered to a block which aborted the data node. This failure in combination with disconnecting nodes could in turn cause the entire cluster to shut down.

To keep this from happening, additional checks are now made when unpacking signals received over TCP, including checks for byte order, compression flag (which must not be used), and the length of the next message in the receive buffer (if there is one).

Whenever two consecutive unpacked messages fail the checks just described, the current message is assumed to be corrupted. In this case, the transporter is marked as having bad data and no more unpacking of messages occurs until the transporter is reconnected. In addition, an entry is written to the cluster log containing the error as well as a hex dump of the corrupted message. (Bug #73843, Bug #19582925)

- `ndb_restore --print_data` truncated `TEXT` and `BLOB` column values to 240 bytes rather than 256 bytes. (Bug #65467, Bug #14571512)
- Transporter send buffers were not updated properly following a failed send. (Bug #45043, Bug #20113145)

## Changes in MySQL NDB Cluster 7.3.7 (5.6.21-ndb-7.3.7) (2014-10-17, General Availability)

MySQL NDB Cluster 7.3.7 is a new release of NDB Cluster, based on MySQL Server 5.6 and including features from version 7.3 of the `NDB` storage engine, as well as fixing a number of recently discovered bugs in previous NDB Cluster releases.

**Obtaining MySQL NDB Cluster 7.3.** MySQL NDB Cluster 7.3 source code and binaries can be obtained from <http://dev.mysql.com/downloads/cluster/>.

For an overview of changes made in MySQL NDB Cluster 7.3, see [What is New in NDB Cluster 7.3](#).

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 5.6 through MySQL 5.6.21 (see [Changes in MySQL 5.6.21 \(2014-09-23, General Availability\)](#)).

- [Bundled SSL Update \(Commercial Releases\)](#)
- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

### Bundled SSL Update (Commercial Releases)

- Starting with this release, commercial distributions of MySQL NDB Cluster 7.3 are built using OpenSSL 1.0.1i.

## Functionality Added or Changed

- After adding new data nodes to the configuration file of a MySQL NDB Cluster having many API nodes, but prior to starting any of the data node processes, API nodes tried to connect to these “missing” data nodes several times per second, placing extra loads on management nodes and the network. To reduce unnecessary traffic caused in this way, it is now possible to control the amount of time that an API node waits between attempts to connect to data nodes which fail to respond; this is implemented in two new API node configuration parameters `StartConnectBackoffMaxTime` and `ConnectBackoffMaxTime`.

Time elapsed during node connection attempts is not taken into account when applying these parameters, both of which are given in milliseconds with approximately 100 ms resolution. As long as the API node is not connected to any data nodes as described previously, the value of the `StartConnectBackoffMaxTime` parameter is applied; otherwise, `ConnectBackoffMaxTime` is used.

In a MySQL NDB Cluster with many unstarted data nodes, the values of these parameters can be raised to circumvent connection attempts to data nodes which have not yet begun to function in the cluster, as well as moderate high traffic to management nodes.

For more information about the behavior of these parameters, see [Defining SQL and Other API Nodes in an NDB Cluster](#). (Bug #17257842)

- Added the `--exclude-missing-tables` option for `ndb_restore`. When enabled, the option causes tables present in the backup but not in the target database to be ignored. (Bug #57566, Bug #11764704)

## Bugs Fixed

- **NDB Replication:** The fix for Bug #18770469 in the MySQL Server made changes in the transactional behavior of the temporary conversion tables used when replicating between tables with different schemas. These changes as implemented are not compatible with NDB, and thus the fix for this bug has been reverted in MySQL NDB Cluster. (Bug #19692387)

References: See also: Bug #19704825. Reverted patches: Bug #18770469.

- **NDB Cluster APIs:** The fix for Bug #16723708 stopped the `ndb_logevent_get_next()` function from casting a log event's `ndb_mgm_event_category` to an `enum` type, but this change interfered with existing applications, and so the function's original behavior is now reinstated. A new MGM API function exhibiting the corrected behavior `ndb_logevent_get_next2()` has been added in this release to take the place of the reverted function, for use in applications that do not require backward compatibility. In all other respects apart from this, the new function is identical with its predecessor. (Bug #18354165)

References: Reverted patches: Bug #16723708.

- **NDB Cluster APIs:** NDB API scans leaked `Ndb_cluster_connection` objects after `nextResult()` was called when an operation resulted in an error. This leak locked up the corresponding connection objects in the `DBTC` kernel block until the connection was closed. (Bug #17730825, Bug #20170731)
- **MySQL NDB ClusterJ:** Retrieval of values from `BLOB` and `TEXT` columns by ClusterJ column accessor methods was not handled correctly. (Bug #18419468, Bug #19028487)
- When assembling error messages of the form `Incorrect state for node n state: node_state`, written when the transporter failed to connect, the node state was used in place of the node ID in a number of instances, which resulted in errors of this type for which the node state was reported incorrectly. (Bug #19559313, Bug #73801)

- In some cases, transporter receive buffers were reset by one thread while being read by another. This happened when a race condition occurred between a thread receiving data and another thread initiating disconnect of the transporter (disconnection clears this buffer). Concurrency logic has now been implemented to keep this race from taking place. (Bug #19552283, Bug #73790)
- The failure of a data node could in some situations cause a set of API nodes to fail as well due to the sending of a `CLOSE_COMREQ` signal that was sometimes not completely initialized. (Bug #19513967)
- A more detailed error report is printed in the event of a critical failure in one of the NDB internal `sendSignal*()` methods, prior to crashing the process, as was already implemented for `sendSignal()`, but was missing from the more specialized `sendSignalNoRelease()` method. Having a crash of this type correctly reported can help with identifying configuration hardware issues in some cases. (Bug #19414511)

References: See also: Bug #19390895.

- `ndb_restore` failed to restore the cluster's metadata when there were more than approximately 17 K data objects. (Bug #19202654)
- The fix for a previous issue with the handling of multiple node failures required determining the number of TC instances the failed node was running, then taking them over. The mechanism to determine this number sometimes provided an invalid result which caused the number of TC instances in the failed node to be set to an excessively high value. This in turn caused redundant takeover attempts, which wasted time and had a negative impact on the processing of other node failures and of global checkpoints. (Bug #19193927)

References: This issue is a regression of: Bug #18069334.

- Parallel transactions performing reads immediately preceding a delete on the same tuple could cause the NDB kernel to crash. This was more likely to occur when separate TC threads were specified using the `ThreadConfig` configuration parameter. (Bug #19031389)
- Attribute promotion between different `TEXT` types (any of `TINYTEXT`, `TEXT`, `MEDIUMTEXT`, and `LONGTEXT`) by `ndb_restore` was not handled properly in some cases. In addition, `TEXT` values are now truncated according to the limits set by `mysqld` (for example, values converted to `TINYTEXT` from another type are truncated to 256 bytes). In the case of columns using a multibyte character set, the value is truncated to the end of the last well-formed character.

Also as a result of this fix, conversion to a `TEXT` column of any size that uses a different character set from the original is now disallowed. (Bug #18875137)

- To assist with diagnostic issues where many watchdog warnings are raised, it is now possible to activate (or deactivate) a killer watchdog using `DUMP 2610` in the `ndb_mgm` client. When set, this shuts down the data node on which the next watchdog warning occurs, providing a trace log. See `DUMP 2610`, for more information. (Bug #18703922)
- The NDB optimized node recovery mechanism attempts to transfer only relevant page changes to a starting node in order to speed the recovery process; this is done by having the starting node indicate the index of the last global checkpoint (GCI) in which it participated, so that the node that was already running copies only data for rows which have changed since that GCI. Every row has a GCI metacolumn which facilitates this; for a deleted row, the slot formerly storing this row's data contains a GCI value, and for deleted pages, every row on the missing page is considered changed and thus needs to be sent.

When these changes are received by the starting node, this node performs a lookup for the page and index to determine what they contain. This lookup could cause a real underlying page to be mapped against the logical page ID, even when this page contained no data.

One way in which this issue could manifest itself occurred after cluster `DataMemory` usage approached maximum, and deletion of many rows followed by a rolling restart of the data nodes was performed with the expectation that this would free memory, but in fact it was possible in this scenario for memory not to be freed and in some cases for memory usage actually to increase to its maximum.

This fix solves these issues by ensuring that a real physical page is mapped to a logical ID during node recovery only when this page contains actual data which needs to be stored. (Bug #18683398, Bug #18731008)

- When a data node sent a `MISSING_DATA` signal due to a buffer overflow and no event data had yet been sent for the current epoch, the dummy event list created to handle this inconsistency was not deleted after the information in the dummy event list was transferred to the completed list. (Bug #18410939)
- Incorrect calculation of the next autoincrement value following a manual insertion towards the end of a cached range could result in duplicate values sometimes being used. This issue could manifest itself when using certain combinations of values for `auto_increment_increment`, `auto_increment_offset`, and `ndb_autoincrement_prefetch_sz`.

This issue has been fixed by modifying the calculation to make sure that the next value from the cache as computed by `NDB` is of the form `auto_increment_offset + (N * auto_increment_increment)`. This avoids any rounding up by the MySQL Server of the returned value, which could result in duplicate entries when the rounded-up value fell outside the range of values cached by `NDB`. (Bug #17893872)

- `ndb_show_tables --help` output contained misleading information about the `--database (-d)` option. In addition, the long form of the option (`--database`) did not work properly. (Bug #17703874)
- Using the `--help` option with `ndb_print_file` caused the program to segfault. (Bug #17069285)
- For multithreaded data nodes, some threads do communicate often, with the result that very old signals can remain at the top of the signal buffers. When performing a thread trace, the signal dumper calculated the latest signal ID from what it found in the signal buffers, which meant that these old signals could be erroneously counted as the newest ones. Now the signal ID counter is kept as part of the thread state, and it is this value that is used when dumping signals for trace files. (Bug #73842, Bug #19582807)

## Changes in MySQL NDB Cluster 7.3.6 (5.6.19-ndb-7.3.6) (2014-07-11, General Availability)

MySQL NDB Cluster 7.3.6 is a new release of NDB Cluster, based on MySQL Server 5.6 and including features from version 7.3 of the `NDB` storage engine, as well as fixing a number of recently discovered bugs in previous NDB Cluster releases.

**Obtaining MySQL NDB Cluster 7.3.** MySQL NDB Cluster 7.3 source code and binaries can be obtained from <http://dev.mysql.com/downloads/cluster/>.

For an overview of changes made in MySQL NDB Cluster 7.3, see [What is New in NDB Cluster 7.3](#).

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 5.6 through MySQL 5.6.19 (see [Changes in MySQL 5.6.19 \(2014-05-30, General Availability\)](#)).

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



## Functionality Added or Changed

- **NDB Cluster APIs:** Added as an aid to debugging the ability to specify a human-readable name for a given `Ndb` object and later to retrieve it. These operations are implemented, respectively, as the `setNdbObjectName()` and `getNdbObjectName()` methods.

To make tracing of event handling between a user application and NDB easier, you can use the reference (from `getReference()` followed by the name (if provided) in printouts; the reference ties together the application `Ndb` object, the event buffer, and the NDB storage engine's `SUMA` block. (Bug #18419907)

## Bugs Fixed

- **NDB Disk Data:** Setting the undo buffer size used by `InitialLogFileGroup` to a value greater than that set by `SharedGlobalMemory` prevented data nodes from starting; the data nodes failed with Error 1504 `Out of logbuffer memory`. While the failure itself is expected behavior, the error message did not provide sufficient information to diagnose the actual source of the problem; now in such cases, a more specific error message `Out of logbuffer memory (specify smaller undo_buffer_size or increase SharedGlobalMemory)` is supplied. (Bug #11762867, Bug #55515)
- **NDB Replication:** When using `NDB$EPOCH_TRANS`, conflicts between `DELETE` operations were handled like conflicts between updates, with the primary rejecting the transaction and dependents, and realigning the secondary. This meant that their behavior with regard to subsequent operations on any affected row or rows depended on whether they were in the same epoch or a different one: within the same epoch, they were considered conflicting events; in different epochs, they were not considered in conflict.

This fix brings the handling of conflicts between deletes by `NDB$EPOCH_TRANS` with that performed when using `NDB$EPOCH` for conflict detection and resolution, and extends testing with `NDB$EPOCH` and `NDB$EPOCH_TRANS` to include “delete-delete” conflicts, and encapsulate the expected result, with transactional conflict handling modified so that a conflict between `DELETE` operations *alone* is not sufficient to cause a transaction to be considered in conflict. (Bug #18459944)

- **NDB Replication:** The implicit `FLUSH LOGS` performed as part of binary log rotation could deadlock and time out with a sufficiently small value for `max_binlog_size`. (Bug #18274220, Bug #19895502, Bug #20009154)

References: See also: Bug #18845822, Bug #19793475, Bug #16884594.

- **NDB Cluster APIs:** When two tables had different foreign keys with the same name, `ndb_restore` considered this a name conflict and failed to restore the schema. As a result of this fix, a slash character (`/`) is now expressly disallowed in foreign key names, and the naming format `parent_id/child_id/fk_name` is now enforced by the NDB API. (Bug #18824753)
- **NDB Cluster APIs:** When an NDB data node indicates a buffer overflow via an empty epoch, the event buffer places an inconsistent data event in the event queue. When this was consumed, it was not removed from the event queue as expected, causing subsequent `nextEvent()` calls to return 0. This caused event consumption to stall because the inconsistency remained flagged forever, while event data accumulated in the queue.

Event data belonging to an empty inconsistent epoch can be found either at the beginning or somewhere in the middle. `pollEvents()` returns 0 for the first case. This fix handles the second case: calling `nextEvent()` call dequeues the inconsistent event before it returns. In order to benefit from this fix, user applications must call `nextEvent()` even when `pollEvents()` returns 0. (Bug #18716991)

- **NDB Cluster APIs:** The `pollEvents()` method returned 1, even when called with a wait time equal to 0, and there were no events waiting in the queue. Now in such cases it returns 0 as expected. (Bug #18703871)
- **MySQL NDB ClusterJ:** Writing a value failed when read from a fixed-width `char` column using `utf8` to another column of the same type and length but using `latin1`. The data was returned with extra spaces after being padded during its insertion. The value is now trimmed before returning it.

This fix also corrects `Data length too long` errors during the insertion of valid `utf8` characters of 2 or more bytes. This was due to padding of the data before encoding it, rather than after. (Bug #71435, Bug #18283369)

- Processing a `NODE_FAILREP` signal that contained an invalid node ID could cause a data node to fail. (Bug #18993037, Bug #73015)

References: This issue is a regression of: Bug #16007980.

- When building out of source, some files were written to the source directory instead of the build dir. These included the `manifest.mf` files used for creating ClusterJ jars and the `pom.xml` file used by `mvn_install_ndbjtie.sh`. In addition, `ndbinfo.sql` was written to the build directory, but marked as output to the source directory in `CMakeLists.txt`. (Bug #18889568, Bug #72843)
- When the binary log injector thread commits an epoch to the binary log and this causes the log file to reach maximum size, it may need to rotate the binary log. The rotation is not performed until either all the committed transactions from all client threads are flushed to the binary log, or a maximum of 30 seconds has elapsed. In the case where all transactions were committed prior to the 30-second wait, it was possible for committed transactions from multiple client threads to belong to newer epochs than the latest epoch committed by the injector thread, causing the thread to deadlock with itself, and causing an unnecessary 30-second delay before breaking the deadlock. (Bug #18845822)
- Adding a foreign key failed with NDB Error 208 if the parent index was parent table's primary key, the primary key was not on the table's initial attributes, and the child table was not empty. (Bug #18825966)
- When an NDB table served as both the parent table and a child table for 2 different foreign keys having the same name, dropping the foreign key on the child table could cause the foreign key on the parent table to be dropped instead, leading to a situation in which it was impossible to drop the remaining foreign key. This situation can be modelled using the following `CREATE TABLE` statements:

```
CREATE TABLE parent (
  id INT NOT NULL,
  PRIMARY KEY (id)
) ENGINE=NDB;

CREATE TABLE child (
  id INT NOT NULL,
  parent_id INT,
  PRIMARY KEY (id),
  INDEX par_ind (parent_id),

  FOREIGN KEY (parent_id)
  REFERENCES parent(id)
) ENGINE=NDB;

CREATE TABLE grandchild (
  id INT,
  parent_id INT,
  INDEX par_ind (parent_id),

  FOREIGN KEY (parent_id)
  REFERENCES child(id)
```

```
) ENGINE=NDB;
```

With the tables created as just shown, the issue occurred when executing the statement `ALTER TABLE child DROP FOREIGN KEY parent_id`, because it was possible in some cases for NDB to drop the foreign key from the `grandchild` table instead. When this happened, any subsequent attempt to drop the foreign key from either the `child` or from the `grandchild` table failed. (Bug #18662582)

- It was possible for a data node restart to become stuck indefinitely in start phase 101 (see [Summary of NDB Cluster Start Phases](#)) when there were connection problems between the node being restarted and one or more subscribing API nodes.

To help prevent this from happening, a new data node configuration parameter `RestartSubscriberConnectTimeout` has been introduced, which can be used to control how long a data node restart can stall in start phase 101 before giving up and attempting to restart again. The default is 12000 ms. (Bug #18599198)

- Executing `ALTER TABLE ... REORGANIZE PARTITION` after increasing the number of data nodes in the cluster from 4 to 16 led to a crash of the data nodes. This issue was shown to be a regression caused by previous fix which added a new dump handler using a dump code that was already in use (7019), which caused the command to execute two different handlers with different semantics. The new handler was assigned a new `DUMP` code (7024). (Bug #18550318)

References: This issue is a regression of: Bug #14220269.

- Following a long series of inserts, when running with a relatively small redo log and an insufficient large value for `MaxNoOfConcurrentTransactions`, there remained transactions that were blocked by the lack of redo log and were thus not aborted in the correct state (waiting for prepare log to be sent to disk, or `LOG_QUEUED` state). This caused the redo log to remain blocked until unblocked by a completion of a local checkpoint. This could lead to a deadlock, when the blocked aborts in turned blocked global checkpoints, and blocked GCPs block LCPs. To prevent this situation from arising, we now abort immediately when we reach the `LOG_QUEUED` state in the abort state handler. (Bug #18533982)
- `ndbmt` supports multiple parallel receiver threads, each of which performs signal reception for a subset of the remote node connections (transporters) with the mapping of `remote_nodes` to receiver threads decided at node startup. Connection control is managed by the multi-instance `TRPMAN` block, which is organized as a proxy and workers, and each receiver thread has a `TRPMAN` worker running locally.

The `QMGR` block sends signals to `TRPMAN` to enable and disable communications with remote nodes. These signals are sent to the `TRPMAN` proxy, which forwards them to the workers. The workers themselves decide whether to act on signals, based on the set of remote nodes they manage.

The current issue arises because the mechanism used by the `TRPMAN` workers for determining which connections they are responsible for was implemented in such a way that each worker thought it was responsible for all connections. This resulted in the `TRPMAN` actions for `OPEN_COMORD`, `ENABLE_COMREQ`, and `CLOSE_COMREQ` being processed multiple times.

The fix keeps `TRPMAN` instances (receiver threads) executing `OPEN_COMORD`, `ENABLE_COMREQ` and `CLOSE_COMREQ` requests. In addition, the correct `TRPMAN` instance is now chosen when routing from this instance for a specific remote connection. (Bug #18518037)

- During data node failure handling, the transaction coordinator performing takeover gathers all known state information for any failed TC instance transactions, determines whether each transaction has been committed or aborted, and informs any involved API nodes so that they can report this accurately to their clients. The TC instance provides this information by sending `TCKEY_FAILREF` or `TCKEY_FAILCONF` signals to the API nodes as appropriate to each affected transaction.

In the event that this TC instance does not have a direct connection to the API node, it attempts to deliver the signal by routing it through another data node in the same node group as the failing TC, and sends a `GSN_TCKEY_FAILREFCONF_R` signal to TC block instance 0 in that data node. A problem arose in the case of multiple transaction coordinators, when this TC instance did not have a signal handler for such signals, which led it to fail.

This issue has been corrected by adding a handler to the TC proxy block which in such cases forwards the signal to one of the local TC worker instances, which in turn attempts to forward the signal on to the API node. (Bug #18455971)

- When running with a very slow main thread, and one or more transaction coordinator threads, on different CPUs, it was possible to encounter a timeout when sending a `DIH_SCAN_GET_NODESREQ` signal, which could lead to a crash of the data node. Now in such cases the timeout is avoided. (Bug #18449222)
- Failure of multiple nodes while using `ndbmt` with multiple TC threads was not handled gracefully under a moderate amount of traffic, which could in some cases lead to an unplanned shutdown of the cluster. (Bug #18069334)
- A local checkpoint (LCP) is tracked using a global LCP state (`c_lcpState`), and each NDB table has a status indicator which indicates the LCP status of that table (`tabLcpStatus`). If the global LCP state is `LCP_STATUS_IDLE`, then all the tables should have an LCP status of `TLS_COMPLETED`.

When an LCP starts, the global LCP status is `LCP_INIT_TABLES` and the thread starts setting all the NDB tables to `TLS_ACTIVE`. If any tables are not ready for LCP, the LCP initialization procedure continues with `CONTINUEB` signals until all tables have become available and been marked `TLS_ACTIVE`. When this initialization is complete, the global LCP status is set to `LCP_STATUS_ACTIVE`.

This bug occurred when the following conditions were met:

- An LCP was in the `LCP_INIT_TABLES` state, and some but not all tables had been set to `TLS_ACTIVE`.
- The master node failed before the global LCP state changed to `LCP_STATUS_ACTIVE`; that is, before the LCP could finish processing all tables.
- The `NODE_FAILREP` signal resulting from the node failure was processed before the final `CONTINUEB` signal from the LCP initialization process, so that the node failure was processed while the LCP remained in the `LCP_INIT_TABLES` state.

Following master node failure and selection of a new one, the new master queries the remaining nodes with a `MASTER_LCPREQ` signal to determine the state of the LCP. At this point, since the LCP status was `LCP_INIT_TABLES`, the LCP status was reset to `LCP_STATUS_IDLE`. However, the LCP status of the tables was not modified, so there remained tables with `TLS_ACTIVE`. Afterwards, the failed node is removed from the LCP. If the LCP status of a given table is `TLS_ACTIVE`, there is a check that the global LCP status is not `LCP_STATUS_IDLE`; this check failed and caused the data node to fail.

Now the `MASTER_LCPREQ` handler ensures that the `tabLcpStatus` for all tables is updated to `TLS_COMPLETED` when the global LCP status is changed to `LCP_STATUS_IDLE`. (Bug #18044717)

- When performing a copying `ALTER TABLE` operation, `mysqld` creates a new copy of the table to be altered. This intermediate table, which is given a name bearing the prefix `#sql-`, has an updated schema but contains no data. `mysqld` then copies the data from the original table to this intermediate table, drops the original table, and finally renames the intermediate table with the name of the original table.

`mysqld` regards such a table as a temporary table and does not include it in the output from `SHOW TABLES`; `mysqldump` also ignores an intermediate table. However, `NDB` sees no difference between such an intermediate table and any other table. This difference in how intermediate tables are viewed by `mysqld` (and MySQL client programs) and by the `NDB` storage engine can give rise to problems when performing a backup and restore if an intermediate table existed in `NDB`, possibly left over from a failed `ALTER TABLE` that used copying. If a schema backup is performed using `mysqldump` and the `mysql` client, this table is not included. However, in the case where a data backup was done using the `ndb_mgm` client's `BACKUP` command, the intermediate table was included, and was also included by `ndb_restore`, which then failed due to attempting to load data into a table which was not defined in the backed up schema.

To prevent such failures from occurring, `ndb_restore` now by default ignores intermediate tables created during `ALTER TABLE` operations (that is, tables whose names begin with the prefix `#sql-`). A new option `--exclude-intermediate-sql-tables` is added that makes it possible to override the new behavior. The option's default value is `TRUE`; to cause `ndb_restore` to revert to the old behavior and to attempt to restore intermediate tables, set this option to `FALSE`. (Bug #17882305)

- The logging of insert failures has been improved. This is intended to help diagnose occasional issues seen when writing to the `mysql.ndb_binlog_index` table. (Bug #17461625)
- The `DEFINER` column in the `INFORMATION_SCHEMA.VIEWS` table contained erroneous values for views contained in the `ndbinfo` information database. This could be seen in the result of a query such as `SELECT TABLE_NAME, DEFINER FROM INFORMATION_SCHEMA.VIEWS WHERE TABLE_SCHEMA='ndbinfo'`. (Bug #17018500)
- Employing a `CHAR` column that used the `UTF8` character set as a table's primary key column led to node failure when restarting data nodes. Attempting to restore a table with such a primary key also caused `ndb_restore` to fail. (Bug #16895311, Bug #68893)
- The `--order (-o)` option for the `ndb_select_all` utility worked only when specified as the last option, and did not work with an equals sign.

As part of this fix, the program's `--help` output was also aligned with the `--order` option's correct behavior. (Bug #64426, Bug #16374870)

## Changes in MySQL NDB Cluster 7.3.5 (5.6.17-ndb-7.3.5) (2014-04-07, General Availability)

MySQL NDB Cluster 7.3.5 is a new release of NDB Cluster, based on MySQL Server 5.6 and including features from version 7.3 of the `NDB` storage engine, as well as fixing a number of recently discovered bugs in previous NDB Cluster releases.

**Obtaining MySQL NDB Cluster 7.3.** MySQL NDB Cluster 7.3 source code and binaries can be obtained from <http://dev.mysql.com/downloads/cluster/>.

For an overview of changes made in MySQL NDB Cluster 7.3, see [What is New in NDB Cluster 7.3](#).

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 5.6 through MySQL 5.6.17 (see [Changes in MySQL 5.6.17 \(2014-03-27, General Availability\)](#)).

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

## Functionality Added or Changed

- Handling of `LongMessageBuffer` shortages and statistics has been improved as follows:
  - The default value of `LongMessageBuffer` has been increased from 4 MB to 64 MB.
  - When this resource is exhausted, a suitable informative message is now printed in the data node log describing possible causes of the problem and suggesting possible solutions.
  - `LongMessageBuffer` usage information is now shown in the `ndbinfo.memoryusage` table. See the description of this table for an example and additional information.

## Bugs Fixed

- **Important Change:** The server system variables `ndb_index_cache_entries` and `ndb_index_stat_freq`, which had been deprecated in a previous MySQL NDB Cluster release series, have now been removed. (Bug #11746486, Bug #26673)
- **Replication:** Log rotation events could cause `group_relay_log_pos` to be moved forward incorrectly within a group. This meant that, when the transaction was retried, or if the SQL thread was stopped in the middle of a transaction following one or more log rotations (such that the transaction or group spanned multiple relay log files), part or all of the group was silently skipped.

This issue has been addressed by correcting a problem in the logic used to avoid touching the coordinates of the SQL thread when updating the log position as part of a relay log rotation whereby it was possible to update the SQL thread's coordinates when not using a multithreaded slave, even in the middle of a group. (Bug #18482854)

- **Solaris:** An issue found when compiling the MySQL NDB Cluster software for Solaris platforms could lead to problems when using `ThreadConfig` on such systems. (Bug #18181656)
- **NDB Replication:** A slave in MySQL NDB Cluster Replication now monitors the progression of epoch numbers received from its immediate upstream master, which can both serve as a useful check on the low-level functioning of replication, and provide a warning in the event replication is restarted accidentally at an already-applied position.

As a result of this enhancement, an epoch ID collision has the following results, depending on the state of the slave SQL thread:

- Following a `RESET SLAVE` statement, no action is taken, in order to allow the execution of this statement without spurious warnings.
- Following `START SLAVE`, a warning is produced that the slave is being positioned at an epoch that has already been applied.
- In all other cases, the slave SQL thread is stopped against the possibility that a system malfunction has resulted in the re-application of an existing epoch.

(Bug #17461576)

References: See also: Bug #17369118.

- **NDB Cluster APIs:** When an NDB API client application received a signal with an invalid block or signal number, NDB provided only a very brief error message that did not accurately convey the nature of the problem. Now in such cases, appropriate printouts are provided when a bad signal or message is detected. In addition, the message length is now checked to make certain that it matches the size of the embedded signal. (Bug #18426180)

- **NDB Cluster APIs:** Refactoring that was performed in MySQL NDB Cluster 7.3.4 inadvertently introduced a dependency in `Ndb.hpp` on a file that is not included in the distribution, which caused NDB API applications to fail to compile. The dependency has been removed. (Bug #18293112, Bug #71803)

References: This issue is a regression of: Bug #17647637.

- **NDB Cluster APIs:** An NDB API application sends a scan query to a data node; the scan is processed by the transaction coordinator (TC). The TC forwards a `LQHKEYREQ` request to the appropriate LDM, and aborts the transaction if it does not receive a `LQHKEYCONF` response within the specified time limit. After the transaction is successfully aborted, the TC sends a `TCROLLBACKREP` to the NDBAPI client, and the NDB API client processes this message by cleaning up any `Ndb` objects associated with the transaction.

The client receives the data which it has requested in the form of `TRANSID_AI` signals, buffered for sending at the data node, and may be delivered after a delay. On receiving such a signal, NDB checks the transaction state and ID: if these are as expected, it processes the signal using the `Ndb` objects associated with that transaction.

The current bug occurs when all the following conditions are fulfilled:

- The transaction coordinator aborts a transaction due to delays and sends a `TCROLLBACPREP` signal to the client, while at the same time a `TRANSID_AI` which has been buffered for delivery at an LDM is delivered to the same client.
- The NDB API client considers the transaction complete on receipt of a `TCROLLBACKREP` signal, and immediately closes the transaction.
- The client has a separate receiver thread running concurrently with the thread that is engaged in closing the transaction.
- The arrival of the late `TRANSID_AI` interleaves with the closing of the user thread's transaction such that `TRANSID_AI` processing passes normal checks before `closeTransaction()` resets the transaction state and invalidates the receiver.

When these conditions are all met, the receiver thread proceeds to continue working on the `TRANSID_AI` signal using the invalidated receiver. Since the receiver is already invalidated, its usage results in a node failure.

Now the `Ndb` object cleanup done for `TCROLLBACKREP` includes invalidation of the transaction ID, so that, for a given transaction, any signal which is received after the `TCROLLBACKREP` arrives does not pass the transaction ID check and is silently dropped. This fix is also implemented for the `TC_COMMITREF`, `TCROLLBACKREF`, `TCKEY_FAILCONF`, and `TCKEY_FAILREF` signals as well.

See also [Operations and Signals](#), for additional information about NDB messaging. (Bug #18196562)

- **NDB Cluster APIs:** The example `ndbapi-examples/ndbapi_blob_ndbrecord/main.cpp` included an internal header file (`ndb_global.h`) not found in the MySQL NDB Cluster binary distribution. The example now uses `stdlib.h` and `string.h` instead of this file. (Bug #18096866, Bug #71409)
- **NDB Cluster APIs:** When `Dictionary::dropTable()` attempted (as a normal part of its internal operations) to drop an index used by a foreign key constraint, the drop failed. Now in such cases, invoking `dropTable()` causes all foreign keys on the table to be dropped, whether this table acts as a parent table, child table, or both.

This issue did not affect dropping of indexes using SQL statements. (Bug #18069680)

References: See also: Bug #17591531.

- **NDB Cluster APIs:** `ndb_restore` could sometimes report `Error 701 System busy with other schema operation` unnecessarily when restoring in parallel. (Bug #17916243)
- When an `ALTER TABLE` statement changed table schemas without causing a change in the table's partitioning, the new table definition did not copy the hash map from the old definition, but used the current default hash map instead. However, the table data was not reorganized according to the new hashmap, which made some rows inaccessible using a primary key lookup if the two hash maps had incompatible definitions.

To keep this situation from occurring, any `ALTER TABLE` that entails a hashmap change now triggers a reorganisation of the table. In addition, when copying a table definition in such cases, the hashmap is now also copied. (Bug #18436558)

- When certain queries generated signals having more than 18 data words prior to a node failure, such signals were not written correctly in the trace file. (Bug #18419554)
- Checking of timeouts is handled by the signal `TIME_SIGNAL`. Previously, this signal was generated by the `QMRG` NDB kernel block in the main thread, and sent to the `QMRG`, `DBLQH`, and `DBTC` blocks (see [NDB Kernel Blocks](#)) as needed to check (respectively) heartbeats, disk writes, and transaction timeouts. In `ndbmt.d` (as opposed to `ndbd`), these blocks all execute in different threads. This meant that if, for example, `QMRG` was actively working and some other thread was put to sleep, the previously sleeping thread received a large number of `TIME_SIGNAL` messages simultaneously when it was woken up again, with the effect that effective times moved very quickly in `DBLQH` as well as in `DBTC`. In `DBLQH`, this had no noticeable adverse effects, but this was not the case in `DBTC`; the latter block could not work on transactions even though time was still advancing, leading to a situation in which many operations appeared to time out because the transaction coordinator (TC) thread was comparatively slow in answering requests.

In addition, when the TC thread slept for longer than 1500 milliseconds, the data node crashed due to detecting that the timeout handling loop had not yet stopped. To rectify this problem, the generation of the `TIME_SIGNAL` has been moved into the local threads instead of `QMRG`; this provides for better control over how quickly `TIME_SIGNAL` messages are allowed to arrive. (Bug #18417623)

- The `ServerPort` and `TcpBind_INADDR_ANY` configuration parameters were not included in the output of `ndb_mgmd --print-full-config`. (Bug #18366909)
- After dropping an NDB table, neither the cluster log nor the output of the `REPORT MemoryUsage` command showed that the `IndexMemory` used by that table had been freed, even though the memory had in fact been deallocated. This issue was introduced in MySQL NDB Cluster 7.3.2. (Bug #18296810)
- `ndb_show_tables` sometimes failed with the error message `Unable to connect to management server` and immediately terminated, without providing the underlying reason for the failure. To provide more useful information in such cases, this program now also prints the most recent error from the `Ndb_cluster_connection` object used to instantiate the connection. (Bug #18276327)
- `-DWITH_NDBMTD=0` did not function correctly, which could cause the build to fail on platforms such as ARM and Raspberry Pi which do not define the memory barrier functions required to compile `ndbmt.d`. (Bug #18267919)

References: See also: Bug #16620938.

- The block threads managed by the multithreading scheduler communicate by placing signals in an out queue or job buffer which is set up between all block threads. This queue has a fixed maximum size, such that when it is filled up, the worker thread must wait for the consumer to drain the queue. In a highly loaded system, multiple threads could end up in a circular wait lock due to full out buffers, such that they were preventing each other from performing any useful work. This condition eventually led to the data node being declared dead and killed by the watchdog timer.



To fix this problem, we detect situations in which a circular wait lock is about to begin, and cause buffers which are otherwise held in reserve to become available for signal processing by queues which are highly loaded. (Bug #18229003)

- The `ndb_mgm` client `START BACKUP` command (see [Commands in the NDB Cluster Management Client](#)) could experience occasional random failures when a ping was received prior to an expected `BackupCompleted` event. Now the connection established by this command is not checked until it has been properly set up. (Bug #18165088)
- When creating a table with foreign key referencing an index in another table, it sometimes appeared possible to create the foreign key even if the order of the columns in the indexes did not match, due to the fact that an appropriate error was not always returned internally. This fix improves the error used internally to work in most cases; however, it is still possible for this situation to occur in the event that the parent index is a unique index. (Bug #18094360)
- Updating parent tables of foreign keys used excessive scan resources and so required unusually high settings for `MaxNoOfLocalScans` and `MaxNoOfConcurrentScans`. (Bug #18082045)
- Dropping a nonexistent foreign key on an NDB table (using, for example, `ALTER TABLE`) appeared to succeed. Now in such cases, the statement fails with a relevant error message, as expected. (Bug #17232212)
- Data nodes running `ndbmt_d` could stall while performing an online upgrade of a MySQL NDB Cluster containing a great many tables from a version prior to NDB 7.2.5 to version 7.2.5 or later. (Bug #16693068)

## Changes in MySQL NDB Cluster 7.3.4 (5.6.15-ndb-7.3.4) (2014-02-06, General Availability)

MySQL NDB Cluster 7.3.4 is a new release of NDB Cluster, based on MySQL Server 5.6 and including features from version 7.3 of the NDB storage engine, as well as fixing a number of recently discovered bugs in previous NDB Cluster releases.

**Obtaining MySQL NDB Cluster 7.3.** MySQL NDB Cluster 7.3 source code and binaries can be obtained from <http://dev.mysql.com/downloads/cluster/>.

For an overview of changes made in MySQL NDB Cluster 7.3, see [What is New in NDB Cluster 7.3](#).

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 5.6 through MySQL 5.6.15 (see [Changes in MySQL 5.6.15 \(2013-12-03, General Availability\)](#)).

## Bugs Fixed

- **Packaging; Solaris:** Compilation of `ndbmt_d` failed on Solaris 10 and 11 for 32-bit `x86`, and the binary was not included in the binary distributions for these platforms. (Bug #16620938)
- **Microsoft Windows:** Timers used in timing scheduler events in the NDB kernel have been refactored, in part to insure that they are monotonic on all platforms. In particular, on Windows, event intervals were previously calculated using values obtained from `GetSystemTimeAsFileTime()`, which reads directly from the system time (“wall clock”), and which may arbitrarily be reset backward or forward, leading to false watchdog or heartbeat alarms, or even node shutdown. Lack of timer monotonicity could also cause slow disk writes during backups and global checkpoints. To fix this issue, the Windows implementation now uses `QueryPerformanceCounters()` instead of

`GetSystemTimeAsFileTime()`. In the event that a monotonic timer is not found on startup of the data nodes, a warning is logged.

In addition, on all platforms, a check is now performed at compile time for available system monotonic timers, and the build fails if one cannot be found; note that `CLOCK_HIGHRES` is now supported as an alternative for `CLOCK_MONOTONIC` if the latter is not available. (Bug #17647637)

- **NDB Disk Data:** When using Disk Data tables and `ndbmta` data nodes, it was possible for the undo buffer to become overloaded, leading to a crash of the data nodes. This issue was more likely to be encountered when using Disk Data columns whose size was approximately 8K or larger. (Bug #16766493)
- **NDB Cluster APIs:** `UINT_MAX64` was treated as a signed value by Visual Studio 2010. To prevent this from happening, the value is now explicitly defined as unsigned. (Bug #17947674)

References: See also: Bug #17647637.

- **NDB Cluster APIs:** It was possible for an `Ndb` object to receive signals for handling before it was initialized, leading to thread interleaving and possible data node failure when executing a call to `Ndb::init()`. To guard against this happening, a check is now made when it is starting to receive signals that the `Ndb` object is properly initialized before any signals are actually handled. (Bug #17719439)
- **NDB Cluster APIs:** Compilation of example NDB API program files failed due to missing include directives. (Bug #17672846, Bug #70759)
- **NDB Cluster APIs:** An application, having opened two distinct instances of `Ndb_cluster_connection`, attempted to use the second connection object to send signals to itself, but these signals were blocked until the destructor was explicitly called for that connection object. (Bug #17626525)

References: This issue is a regression of: Bug #16595838.

- **ndbmemcache:** A memcached server running the NDB engine could crash after being disconnected from a cluster. (Bug #14055851)
- **MySQL NDB ClusterJ:** Call of `setPartitionKey()` failed after a previous transaction failed. It was because the state of the transaction was not properly cleaned up after the transaction failure. This fix adds the clean-up that is needed in the situation. (Bug #17885485)
- Interrupting a drop of a foreign key could cause the underlying table to become corrupt. (Bug #18041636)
- Monotonic timers on several platforms can experience issues which might result in the monotonic clock doing small jumps back in time. This is due to imperfect synchronization of clocks between multiple CPU cores and does not normally have an adverse effect on the scheduler and watchdog mechanisms; so we handle some of these cases by making backtick protection less strict, although we continue to ensure that the backtick is less than 10 milliseconds. This fix also removes several checks for backticks which are thereby made redundant. (Bug #17973819)
- Under certain specific circumstances, in a cluster having two SQL nodes, one of these could hang, and could not be accessed again even after killing the `mysqld` process and restarting it. (Bug #17875885, Bug #18080104)

References: See also: Bug #17934985.

- Poor support or lack of support on some platforms for monotonic timers caused issues with delayed signal handling by the job scheduler for the multithreaded data node. Variances (timer leaps) on such

platforms are now handled in the same way the multithreaded data node process that they are by the singlethreaded version. (Bug #17857442)

References: See also: Bug #17475425, Bug #17647637.

- In some cases, with `ndb_join_pushdown` enabled, it was possible to obtain from a valid query the error `Got error 290 'Corrupt key in TC, unable to xfrm' from NDBCLUSTER` even though the data was not actually corrupted.

It was determined that a `NULL` in a `VARCHAR` column could be used to construct a lookup key, but since `NULL` is never equal to any other value, such a lookup could simple have been eliminated instead. This `NULL` lookup in turn led to the spurious error message.

This fix takes advantage of the fact that a key lookup with `NULL` never finds any matching rows, and so `NDB` does not try to perform the lookup that would have led to the error. (Bug #17845161)

- It was theoretically possible in certain cases for a number of output functions internal to the `NDB` code to supply an uninitialized buffer as output. Now in such cases, a newline character is printed instead. (Bug #17775602, Bug #17775772)
- Use of the `localtime()` function in `NDB` multithreading code led to otherwise nondeterministic failures in `ndbmttd`. This fix replaces this function, which on many platforms uses a buffer shared among multiple threads, with `localtime_r()`, which can have allocated to it a buffer of its own. (Bug #17750252)
- When using single-threaded (`ndbd`) data nodes with `RealTimeScheduler` enabled, the CPU did not, as intended, temporarily lower its scheduling priority to normal every 10 milliseconds to give other, non-realtime threads a chance to run. (Bug #17739131)
- During arbitrator selection, `QMGR` (see [The QMGR Block](#)) runs through a series of states, the first few of which are (in order) `NULL`, `INIT`, `FIND`, `PREP1`, `PREP2`, and `START`. A check for an arbitration selection timeout occurred in the `FIND` state, even though the corresponding timer was not set until `QMGR` reached the `PREP1` and `PREP2` states. Attempting to read the resulting uninitialized timestamp value could lead to false `Could not find an arbitrator, cluster is not partition-safe` warnings.

This fix moves the setting of the timer for arbitration timeout to the `INIT` state, so that the value later read during `FIND` is always initialized. (Bug #17738720)

- The global checkpoint lag watchdog tracking the number of times a check for GCP lag was performed using the system scheduler and used this count to check for a timeout condition, but this caused a number of issues. To overcome these limitations, the GCP watchdog has been refactored to keep track of its own start times, and to calculate elapsed time by reading the (real) clock every time it is called.

In addition, any backticks (rare in any case) are now handled by taking the backward time as the new current time and calculating the elapsed time for this round as 0. Finally, any ill effects of a forward leap, which possibly could expire the watchdog timer immediately, are reduced by never calculating an elapsed time longer than the requested delay time for the watchdog timer. (Bug #17647469)

References: See also: Bug #17842035.

- The length of the interval (intended to be 10 seconds) between warnings for `GCP_COMMIT` when the GCP progress watchdog did not detect progress in a global checkpoint was not always calculated correctly. (Bug #17647213)
- Trying to drop an index used by a foreign key constraint caused data node failure. Now in such cases, the statement used to perform the drop fails. (Bug #17591531)
- In certain rare cases on commit of a transaction, an `Ndb` object was released before the transaction coordinator (`DBTC` kernel block) sent the expected `COMMIT_CONF` signal; `NDB` failed to send a

`COMMIT_ACK` signal in response, which caused a memory leak in the NDB kernel could later lead to node failure.

Now an `Ndb` object is not released until the `COMMIT_CONF` signal has actually been received. (Bug #16944817)

- Losing its connections to the management node or data nodes while a query against the `ndbinfo.memoryusage` table was in progress caused the SQL node where the query was issued to fail. (Bug #14483440, Bug #16810415)
- The `ndbd_redo_log_reader` utility now supports a `--help` option. Using this options causes the program to print basic usage information, and then to exit. (Bug #11749591, Bug #36805)

## Changes in MySQL NDB Cluster 7.3.3 (5.6.14-ndb-7.3.3) (2013-11-18, General Availability)

MySQL NDB Cluster 7.3.3 is a new release of NDB Cluster, based on MySQL Server 5.6 and including features from version 7.3 of the NDB storage engine, as well as fixing a number of recently discovered bugs in previous NDB Cluster releases.

**Obtaining MySQL NDB Cluster 7.3.** MySQL NDB Cluster 7.3 source code and binaries can be obtained from <http://dev.mysql.com/downloads/cluster/>.

For an overview of changes made in MySQL NDB Cluster 7.3, see [What is New in NDB Cluster 7.3](#).

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 5.6 through MySQL 5.6.14 (see [Changes in MySQL 5.6.14 \(2013-09-20, General Availability\)](#)).

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

### Functionality Added or Changed

- The length of time a management node waits for a heartbeat message from another management node is now configurable using the `HeartbeatIntervalMgmdMgmd` management node configuration parameter added in this release. The connection is considered dead after 3 missed heartbeats. The default value is 1500 milliseconds, or a timeout of approximately 6000 ms. (Bug #17807768, Bug #16426805)
- The MySQL NDB Cluster Auto-Installer now generates a `my.cnf` file for each `mysqld` in the cluster before starting it. For more information, see [Using the NDB Cluster Auto-Installer](#). (Bug #16994782)
- `BLOB` and `TEXT` columns are now reorganized by the `ALTER ONLINE TABLE ... REORGANIZE PARTITION` statement. (Bug #13714148)

### Bugs Fixed

- **Performance:** In a number of cases found in various locations in the MySQL NDB Cluster codebase, unnecessary iterations were performed; this was caused by failing to break out of a repeating control structure after a test condition had been met. This community-contributed fix removes the unneeded repetitions by supplying the missing breaks. (Bug #16904243, Bug #69392, Bug #16904338, Bug #69394, Bug #16778417, Bug #69171, Bug #16778494, Bug #69172, Bug #16798410, Bug #69207, Bug #16801489, Bug #69215, Bug #16904266, Bug #69393)

- **Packaging:** Portions of the documentation specific to MySQL NDB Cluster and the [NDB](#) storage engine were not included when installing from RPMs. (Bug #16303451)
- **Microsoft Windows:** The Windows error `ERROR_FILE_EXISTS` was not recognized by [NDB](#), which treated it as an unknown error. (Bug #16970960)
- **NDB Disk Data:** [NDB](#) error 899 `RowId already allocated` was raised due to a RowId “leak” which occurred under either of the following sets of circumstances:
  1. Insertion of a row into an in-memory table was rejected after an ordered index update failed due to insufficient [DataMemory](#).
  2. Insertion of a row into a Disk Data table was rejected due to lack of sufficient table space.
 (Bug #13927679)
 

References: See also: Bug #22494024, Bug #13990924.
- **NDB Replication:** Trying to use a stale `.frm` file and encountering a mismatch between table definitions could cause `mysqld` to make errors when writing to the binary log. (Bug #17250994)
- **NDB Replication:** Replaying a binary log that had been written by a `mysqld` from a MySQL Server distribution (and from not a MySQL NDB Cluster distribution), and that contained DML statements, on a MySQL NDB Cluster SQL node could lead to failure of the SQL node. (Bug #16742250)
- **NDB Cluster APIs:** The `Event::setTable()` method now supports a pointer or a reference to table as its required argument. If a null table pointer is used, the method now returns -1 to make it clear that this is what has occurred. (Bug #16329082)
- **ndbmemcache:** When attempting to start memcached with a `cache_size` larger than that of the available memory and with `preallocate=true` failed, the error message provided only a numeric code, and did not indicate what the actual source of the error was. (Bug #17509293, Bug #70403)
- **ndbmemcache:** The `CMakeLists.txt` files for `ndbmemcache` wrote into the source tree, preventing out-of-source builds. (Bug #14650456)
- **ndbmemcache:** `ndbmemcache` did not handle passed-in `BINARY` values correctly.
- Trying to restore to a table having a `BLOB` column in a different position from that of the original one caused `ndb_restore --restore_data` to fail. (Bug #17395298)
- `ndb_restore` could abort during the last stages of a restore using attribute promotion or demotion into an existing table. This could happen if a converted attribute was nullable and the backup had been run on active database. (Bug #17275798)
- It was not possible to start MySQL NDB Cluster processes created by the Auto-Installer on a Windows host running freeSSHd. (Bug #17269626)
- The `DBUTIL` data node block is now less strict about the order in which it receives certain messages from other nodes. (Bug #17052422)
- `ALTER ONLINE TABLE ... REORGANIZE PARTITION` failed when run against a table having or using a reference to a foreign key. (Bug #17036744, Bug #69619)
- `TUPKEYREQ` signals are used to read data from the tuple manager block (`DBTUP`), and are used for all types of data access, especially for scans which read many rows. A `TUPKEYREQ` specifies a series of 'columns' to be read, which can be either single columns in a specific table, or pseudocolumns, two of

which—`READ_ALL` and `READ_PACKED`—are aliases to read all columns in a table, or some subset of these columns. Pseudocolumns are used by modern NDB API applications as they require less space in the `TUPLEKEYREQ` to specify columns to be read, and can return the data in a more compact (packed) format.

This fix moves the creation and initialization of on-stack Signal objects to only those pseudocolumn reads which need to `EXECUTE_DIRECT` to other block instances, rather than for every read. In addition, the size of an on-stack signal is now varied to suit the requirements of each pseudocolumn, so that only reads of the `INDEX_STAT` pseudocolumn now require initialization (and 3KB memory each time this is performed). (Bug #17009502)

- A race condition could sometimes occur when trying to lock receive threads to cores. (Bug #17009393)
- Results from joins using a `WHERE` with an `ORDER BY ... DESC` clause were not sorted properly; the `DESC` keyword in such cases was effectively ignored. (Bug #16999886, Bug #69528)
- `RealTimeScheduler` did not work correctly with data nodes running `ndbmt.d`. (Bug #16961971)
- File system errors occurring during a local checkpoint could sometimes cause an LCP to hang with no obvious cause when they were not handled correctly. Now in such cases, such errors always cause the node to fail. Note that the LQH block always shuts down the node when a local checkpoint fails; the change here is to make likely node failure occur more quickly and to make the original file system error more visible. (Bug #16961443)
- Maintenance and checking of parent batch completion in the `SPJ` block of the `NDB` kernel was reimplemented. Among other improvements, the completion state of all ancestor nodes in the tree are now preserved. (Bug #16925513)
- Dropping a column, which was not itself a foreign key, from an `NDB` table having foreign keys failed with `ER_TABLE_DEF_CHANGED`. (Bug #16912989)
- The LCP fragment scan watchdog periodically checks for lack of progress in a fragment scan performed as part of a local checkpoint, and shuts down the node if there is no progress after a given amount of time has elapsed. This interval, formerly hard-coded as 60 seconds, can now be configured using the `LcpScanProgressTimeout` data node configuration parameter added in this release.

This configuration parameter sets the maximum time the local checkpoint can be stalled before the LCP fragment scan watchdog shuts down the node. The default is 60 seconds, which provides backward compatibility with previous releases.

You can disable the LCP fragment scan watchdog by setting this parameter to 0. (Bug #16630410)

- Added the `ndb_error_reporter` options `--connection-timeout`, which makes it possible to set a timeout for connecting to nodes, `--dry-scp`, which disables scp connections to remote hosts, and `--skip-nodegroup`, which skips all nodes in a given node group. (Bug #16602002)

References: See also: Bug #11752792, Bug #44082.

- After issuing `START BACKUP id WAIT STARTED`, if `id` had already been used for a backup ID, an error caused by the duplicate ID occurred as expected, but following this, the `START BACKUP` command never completed. (Bug #16593604, Bug #68854)
- `ndb_mgm` treated backup IDs provided to `ABORT BACKUP` commands as signed values, so that backup IDs greater than  $2^{31}$  wrapped around to negative values. This issue also affected out-of-range backup IDs, which wrapped around to negative values instead of causing errors as expected in such cases. The backup ID is now treated as an unsigned value, and `ndb_mgm` now performs proper range checking for backup ID values greater than `MAX_BACKUPS` ( $2^{32}$ ). (Bug #16585497, Bug #68798)

- When trying to specify a backup ID greater than the maximum allowed, the value was silently truncated. (Bug #16585455, Bug #68796)
- The unexpected shutdown of another data node as a starting data node received its node ID caused the latter to hang in Start Phase 1. (Bug #16007980)

References: See also: Bug #18993037.

- `SELECT ... WHERE ... LIKE` from an NDB table could return incorrect results when using `engine_condition_pushdown=ON`. (Bug #15923467, Bug #67724)
- The NDB receive thread waited unnecessarily for additional job buffers to become available when receiving data. This caused the receive mutex to be held during this wait, which could result in a busy wait when the receive thread was running with real-time priority.

This fix also handles the case where a negative return value from the initial check of the job buffer by the receive thread prevented further execution of data reception, which could possibly lead to communication blockage or configured `ReceiveBufferMemory` underutilization. (Bug #15907515)

- When the available job buffers for a given thread fell below the critical threshold, the internal multithreading job scheduler waited for job buffers for incoming rather than outgoing signals to become available, which meant that the scheduler waited the maximum timeout (1 millisecond) before resuming execution. (Bug #15907122)
- Setting `lower_case_table_names` to 1 or 2 on Windows systems caused `ALTER TABLE ... ADD FOREIGN KEY` statements against tables with names containing uppercase letters to fail with Error 155, `No such table: '(null)'`. (Bug #14826778, Bug #67354)
- Under some circumstances, a race occurred where the wrong watchdog state could be reported. A new state name `Packing Send Buffers` is added for watchdog state number 11, previously reported as `Unknown place`. As part of this fix, the state numbers for states without names are always now reported in such cases. (Bug #14824490)
- When a node fails, the Distribution Handler (`DBDIH` kernel block) takes steps together with the Transaction Coordinator (`DBTC`) to make sure that all ongoing transactions involving the failed node are taken over by a surviving node and either committed or aborted. Transactions taken over which are then committed belong in the epoch that is current at the time the node failure occurs, so the surviving nodes must keep this epoch available until the transaction takeover is complete. This is needed to maintain ordering between epochs.

A problem was encountered in the mechanism intended to keep the current epoch open which led to a race condition between this mechanism and that normally used to declare the end of an epoch. This could cause the current epoch to be closed prematurely, leading to failure of one or more surviving data nodes. (Bug #14623333, Bug #16990394)

- Exhaustion of `LongMessageBuffer` memory under heavy load could cause data nodes running `ndbmt` to fail. (Bug #14488185)
- When using dynamic listening ports for accepting connections from API nodes, the port numbers were reported to the management server serially. This required a round trip for each API node, causing the time required for data nodes to connect to the management server to grow linearly with the number of API nodes. To correct this problem, each data node now reports all dynamic ports at once. (Bug #12593774)
- `ndb_error-reporter` did not support the `--help` option. (Bug #11756666, Bug #48606)

References: See also: Bug #11752792, Bug #44082.

- Formerly, the node used as the coordinator or leader for distributed decision making between nodes (also known as the `DICT` manager—see [The DBDICTIONARY Block](#)) was indicated in the output of the `ndb_mgm` client `SHOW` command as the “master” node, although this node has no relationship to a master server in MySQL Replication. (It should also be noted that it is not necessary to know which node is the leader except when debugging `NDBCLUSTER` source code.) To avoid possible confusion, this label has been removed, and the leader node is now indicated in `SHOW` command output using an asterisk (\*) character. (Bug #11746263, Bug #24880)
- The matrix of values used for thread configuration when applying the setting of the `MaxNoOfExecutionThreads` configuration parameter has been improved to align with support for greater numbers of LDM threads. See [Multi-Threading Configuration Parameters \(ndbmtd\)](#), for more information about the changes. (Bug #75220, Bug #20215689)
- Program execution failed to break out of a loop after meeting a desired condition in a number of internal methods, performing unneeded work in all cases where this occurred. (Bug #69610, Bug #69611, Bug #69736, Bug #17030606, Bug #17030614, Bug #17160263)
- `ABORT BACKUP` in the `ndb_mgm` client (see [Commands in the NDB Cluster Management Client](#)) took an excessive amount of time to return (approximately as long as the backup would have required to complete, had it not been aborted), and failed to remove the files that had been generated by the aborted backup. (Bug #68853, Bug #17719439)
- Attribute promotion and demotion when restoring data to `NDB` tables using `ndb_restore --restore_data` with the `--promote-attributes` and `--lossy-conversions` options has been improved as follows:
  - Columns of types `CHAR`, and `VARCHAR` can now be promoted to `BINARY` and `VARBINARY`, and columns of the latter two types can be demoted to one of the first two.

Note that converted character data is not checked to conform to any character set.

- Any of the types `CHAR`, `VARCHAR`, `BINARY`, and `VARBINARY` can now be promoted to `TEXT` or `BLOB`.

When performing such promotions, the only other sort of type conversion that can be performed at the same time is between character types and binary types.

## Changes in MySQL NDB Cluster 7.3.2 (5.6.11-ndb-7.3.2) (2013-06-18, General Availability)

MySQL NDB Cluster 7.3.2 is the first GA release of NDB Cluster, based on MySQL Server 5.6 and including new features in version 7.3 of the `NDB` storage engine, as well as fixing recently discovered bugs in previous NDB Cluster releases.

**Obtaining MySQL NDB Cluster 7.3.** MySQL NDB Cluster 7.3 source code and binaries can be obtained from <http://dev.mysql.com/downloads/cluster/>.

For an overview of changes made in MySQL NDB Cluster 7.3, see [What is New in NDB Cluster 7.3](#).

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 5.6 through MySQL 5.6.11 (see [Changes in MySQL 5.6.11 \(2013-04-18, General Availability\)](#)).

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



## Functionality Added or Changed

- **ndbmemcache:** ndbmemcache now supports values of types `DATE`, `TIME`, and `DATETIME` using microsecond precision. (Bug #16593493)

## Bugs Fixed

- **Packaging; Microsoft Windows:** The MySQL NDB Cluster installer for Windows provided a nonfunctional option to install debug symbols (contained in `*.pdb` files). This option has been removed from the installer.



### Note

You can obtain the `*.pdb` debug files for a given MySQL NDB Cluster release from the Windows `.zip` archive for the same release, such as `mysql-cluster-gpl-7.2.14-win32.zip` or `mysql-cluster-gpl-7.3.2-winx64.zip`.

(Bug #16748308, Bug #69112)

- **NDB Disk Data:** The statements `CREATE TABLESPACE`, `ALTER LOGFILE GROUP`, and `ALTER TABLESPACE` failed with a syntax error when `INITIAL_SIZE` was specified using letter abbreviations such as `M` or `G`. In addition, `CREATE LOGFILE GROUP` failed when `INITIAL_SIZE`, `UNDO_BUFFER_SIZE`, or both options were specified using letter abbreviations. (Bug #13116514, Bug #16104705, Bug #62858)
- **NDB Cluster APIs:** For each log event retrieved using the MGM API, the log event category (`ndb_mgm_event_category`) was simply cast to an `enum` type, which resulted in invalid category values. Now an offset is added to the category following the cast to ensure that the value does not fall out of the allowed range.



### Note

This change was reverted by the fix for Bug #18354165. See the MySQL NDB Cluster API Developer documentation for `ndb_logevent_get_next()`, for more information.

(Bug #16723708)

References: See also: Bug #18354165.

- **NDB Cluster APIs:** The `Ndb::computeHash()` API method performs a `malloc()` if no buffer is provided for it to use. However, it was assumed that the memory thus returned would always be suitably aligned, which is not always the case. Now when `malloc()` provides a buffer to this method, the buffer is aligned after it is allocated, and before it is used. (Bug #16484617)
- **MySQL NDB ClusterJ:** ClusterJ now provides partial support for MySQL date and time data types with fractional seconds (see [Date and Time Type Overview](#)), as shown in the following table:

ClusterJ Class	MySQL Data Types	Precision
<code>java.util.Date</code>	<code>DATETIME</code>	
	millisecond	
<code>TIME</code>	millisecond	
<code>TIMESTAMP</code>	millisecond	

ClusterJ Class	MySQL Data Types	Precision
<code>java.sql.Date</code>	DATE	day
<code>java.sql.Timestamp</code>	DATETIME	millisecond
TIMESTAMP	millisecond	
<code>java.sql.Time</code>	TIME	second

(Bug #16593510)

- `mysql_upgrade` failed when upgrading from MySQL NDB Cluster 7.1.26 to MySQL NDB Cluster 7.2.13 when it attempted to invoke a stored procedure before the `mysql.proc` table had been upgraded. (Bug #16933405)

References: This issue is a regression of: Bug #16226274.

- The planned or unplanned shutdown of one or more data nodes while reading table data from the `ndbinfo` database caused a memory leak. (Bug #16932989)
- Executing `DROP TABLE` while `DBDIH` was updating table checkpoint information subsequent to a node failure could lead to a data node failure. (Bug #16904469)
- In certain cases, when starting a new SQL node, `mysqld` failed with Error 1427 `Api node died, when SUB_START_REQ reached node`. (Bug #16840741)
- Failure to use container classes specific `NDB` during node failure handling could cause leakage of commit-ack markers, which could later lead to resource shortages or additional node crashes. (Bug #16834416)
- Use of an uninitialized variable employed in connection with error handling in the `DBLQH` kernel block could sometimes lead to a data node crash or other stability issues for no apparent reason. (Bug #16834333)
- A race condition in the time between the reception of a `execNODE_FAILREP` signal by the `QMGR` kernel block and its reception by the `DBLQH` and `DBTC` kernel blocks could lead to data node crashes during shutdown. (Bug #16834242)
- The `CLUSTERLOG` command (see [Commands in the NDB Cluster Management Client](#)) caused `ndb_mgm` to crash on Solaris SPARC systems. (Bug #16834030)
- On Solaris SPARC platforms, batched key access execution of some joins could fail due to invalid memory access. (Bug #16818575)
- When 2 `NDB` tables had foreign key references to each other, it was necessary to drop the tables in the same order in which they were created. (Bug #16817928)
- The duplicate weedout algorithm introduced in MySQL 5.6 evaluates semi-joins such as subqueries using `IN`) by first performing a normal join between the outer and inner table which may create duplicates of rows from the outer (and inner) table and then removing any duplicate result rows from the outer table by comparing their primary key values. Problems could arise when `NDB` copied `VARCHAR` values using their maximum length, resulting in a binary key image which contained garbage past the actual lengths of the `VARCHAR` values, which meant that multiple instances of the same key were not binary-identical as assumed by the MySQL server.

To fix this problem, `NDB` now zero-pads such values to the maximum length of the column so that copies of the same key are treated as identical by the weedout process. (Bug #16744050)

- `DROP DATABASE` failed to work correctly when executed against a database containing `NDB` tables joined by foreign key constraints (and all such tables being contained within this database), leaving these tables in place while dropping the remaining tables in the database and reporting failure. (Bug #16692652, Bug #69008)
- When using `firstmatch=on` with the `optimizer_switch` system variable, pushed joins could return too many rows. (Bug #16664035)
- A variable used by the batched key access implementation was not initialized by `NDB` as expected. This could cause a “batch full” condition to be reported after only a single row had been batched, effectively disabling batching altogether and leading to an excessive number of round trips between `mysqld` and `NDB`. (Bug #16485658)
- When started with `--initial` and an invalid `--config-file (-f)` option, `ndb_mgmd` removed the old configuration cache before verifying the configuration file. Now in such cases, `ndb_mgmd` first checks for the file, and continues with removing the configuration cache only if the configuration file is found and is valid. (Bug #16299289)
- Creating more than 32 hash maps caused data nodes to fail. Usually new hashmaps are created only when performing reorganization after data nodes have been added or when explicit partitioning is used, such as when creating a table with the `MAX_ROWS` option, or using `PARTITION BY KEY() PARTITIONS n`. (Bug #14710311)
- Setting `foreign_key_checks = 0` had no effect on the handling of `NDB` tables. Now, doing so causes such checks of foreign key constraints to be suspended—that is, has the same effect on `NDB` tables as it has on `InnoDB` tables. (Bug #14095855, Bug #16286309)

References: See also: Bug #16286164.

## Changes in MySQL NDB Cluster 7.3.1 (5.6.10-ndb-7.3.1) (2013-04-17, Development Milestone)

MySQL NDB Cluster 7.3.1 is a new Developer Milestone release of NDB Cluster, based on MySQL Server 5.6 and previewing new features under development for version 7.3 of the `NDB` storage engine, as well as fixing a number of recently discovered bugs in previous releases.

**Obtaining MySQL NDB Cluster 7.3.** MySQL NDB Cluster 7.3 source code and binaries can be obtained from <http://dev.mysql.com/downloads/cluster/>.

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 5.6 through MySQL 5.6.10 (see [Changes in MySQL 5.6.10 \(2013-02-05, General Availability\)](#)).

- [Based on MySQL Server 5.6](#)
- [MySQL Cluster GUI Configuration Wizard](#)
- [Support for Foreign Key Constraints](#)
- [NoSQL Connector for JavaScript \(Node.js\)](#)
- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

## Based on MySQL Server 5.6

- **Important Change:** MySQL NDB Cluster SQL nodes are now based on MySQL Server 5.6. For information about feature additions and other changes made in MySQL 5.6, see [What Is New in MySQL 5.6](#).

The `mysqld` binary provided with MySQL NDB Cluster 7.3.1 is based on MySQL Server 5.6.10, and includes all MySQL Server 5.6 feature enhancements and bug fixes found in that release; see [Changes in MySQL 5.6.10 \(2013-02-05, General Availability\)](#), for information about these.

## MySQL Cluster GUI Configuration Wizard

- **Important Change:** The MySQL NDB Cluster distribution now includes a browser-based graphical configuration wizard that assists the user in configuring and deploying a MySQL NDB Cluster. This deployment can consist of an arbitrary number of nodes (within certain limits) on the user machine only, or include nodes distributed on a local network. The wizard can be launched from the command line (using the `ndb_setup` utility now included in the binary distribution) or a desktop file browser.

For more information about this tool, see [The NDB Cluster Auto-Installer](#).

## Support for Foreign Key Constraints

- **Important Change:** MySQL NDB Cluster now supports foreign key constraints between [NDB](#) tables, including support for [CASCADE](#), [SET NULL](#), and [RESTRICT](#) and [NO ACTION](#) reference options for [DELETE](#) and [UPDATE](#) actions. (MySQL currently does not support [SET DEFAULT](#).)

MySQL requires generally that all child and parent tables in foreign key relationships employ the same storage engine; thus, to use foreign keys with MySQL NDB Cluster tables, the child and parent table must each use the [NDB](#) storage engine. (It is not possible, for example, for a foreign key on an [NDB](#) table to reference an index of an [InnoDB](#) table.)

Note that MySQL NDB Cluster tables that are explicitly partitioned by [KEY](#) or [LINEAR KEY](#) may contain foreign key references or be referenced by foreign keys (or both). This is unlike the case with [InnoDB](#) tables that are user partitioned, which may not have any foreign key relationships.

You can create an [NDB](#) table having a foreign key reference on another [NDB](#) table using [CREATE TABLE ... \[CONSTRAINT\] FOREIGN KEY ... REFERENCES](#). A child table's foreign key definitions can be seen in the output of [SHOW CREATE TABLE](#); you can also obtain information about foreign keys by querying the [INFORMATION\\_SCHEMA.KEY\\_COLUMN\\_USAGE](#) table.

[FOREIGN KEY Constraints](#), provides general information about foreign key support in MySQL. For more information about the syntax supported by MySQL for foreign keys, see [Using FOREIGN KEY Constraints](#).

## NoSQL Connector for JavaScript (Node.js)

- **NDB Cluster APIs:** MySQL NDB Cluster 7.3 includes support for JavaScript applications written against Node.js with MySQL NDB Cluster and MySQL Server as a data store. The Connector for JavaScript provides a domain object model similar in many ways to that employed by ClusterJ (see [The ClusterJ API and Data Object Model](#)) and can be used with either of two back-end adapters: the `ndb` adapter, which uses the NDB API to provide high-performance native access to MySQL NDB Cluster; and the `mysql-js` adapter, which uses the MySQL Server and the `node-mysql` driver available from <https://github.com/felixge/node-mysql/>.

The Connector for JavaScript is included with the MySQL NDB Cluster distribution, and contains setup programs which can assist you with installation of the connector. You must have Node.js and MySQL NDB Cluster installed prior to running the setup scripts. The `node-mysql` driver is also required for the `mysql-js` Node.js adapter; you can install this using the package management tool included with Node.js. For more information, see [MySQL NoSQL Connector for JavaScript](#).

## Functionality Added or Changed

- **Important Change:** The behavior of and values used for the `TCP_RCV_BUF_SIZE` and `TCP_SND_BUF_SIZE` TCP configuration parameters have been improved. Formerly, the default values for these parameters were 70080 and 71540, respectively—which it was later found could lead to excessive timeouts in some circumstances—with the minimum for each of them being 1. Now, the default and recommended value is 0 for both `TCP_RCV_BUF_SIZE` and `TCP_SND_BUF_SIZE`, which allows the operating system or platform to choose the send or receive buffer size for TCP sockets. (Bug #14554519)

References: See also: Bug #14168828.

- **NDB Cluster APIs:** Added `DUMP` code 2514, which provides information about counts of transaction objects per API node. For more information, see [DUMP 2514](#). See also [Commands in the NDB Cluster Management Client](#). (Bug #15878085)
- When `ndb_restore` fails to find a table, it now includes in the error output an NDB API error code giving the reason for the failure. (Bug #16329067)
- Data node logs now provide tracking information about arbitrations, including which nodes have assumed the arbitrator role and at what times. (Bug #11761263, Bug #53736)

## Bugs Fixed

- **Important Note; NDB Replication:** Setting `binlog_row_image=minimal` caused MySQL NDB Cluster replication conflict resolution to fail.

To fix this problem, support for this variable is disabled in the NDB storage engine such that setting it or the corresponding `--binlog-row-image` server option has no effect for NDB tables. (Bug #16316828)

- **ndbmemcache:** When using a large `values` table, `memcached` failed to store flags correctly or not at all in NDB, even when it had been configured to, as follows:
  - In MySQL NDB Cluster 7.2.4 and earlier, the flags value was never stored for long values, even if configured to do so.
  - In MySQL NDB Cluster 7.2.6 and later, nonzero flags values were correctly stored for long values, but a value of zero was not stored.

In addition, because some clients (such as Spymemcached for Java) use the flags field, this fix also enables storage of flags for the values table by default. (Bug #14088078)

- **API:** `mysqld` failed to respond when `mysql_shutdown()` was invoked from a C application, or `mysqladmin shutdown` was run from the command line. (Bug #14849574)
- When an update of an NDB table changes the primary key (or part of the primary key), the operation is executed as a delete plus an insert. In some cases, the initial read operation did not retrieve all column values required by the insert, so that another read was required. This fix ensures that all required column values are included in the first read in such cases, which saves the overhead of an additional read operation. (Bug #16614114)

- Pushed joins executed when `optimizer_switch='batched_key_access=on'` was also in use returned incorrect results. (Bug #16437431)
- Selecting from the `INFORMATION_SCHEMA.KEY_COLUMN_USAGE` table while using tables with foreign keys caused `mysqld` to crash. (Bug #16246874, Bug #68224)
- Including a table as a part of a pushed join should be rejected if there are outer joined tables in between the table to be included and the tables with which it is joined with; however the check as performed for any such outer joined tables did so by checking the join type against the root of the pushed query, rather than the common ancestor of the tables being joined. (Bug #16199028)

References: See also: Bug #16198866.

- Some queries were handled differently with `ndb_join_pushdown` enabled, due to the fact that outer join conditions were not always pruned correctly from joins before they were pushed down. (Bug #16198866)

References: See also: Bug #16199028.

- Attempting to perform additional operations such as `ADD COLUMN` as part of an `ALTER [ONLINE | OFFLINE] TABLE ... RENAME ...` statement is not supported, and now fails with an `ER_NOT_SUPPORTED_YET` error. (Bug #16021021)
- Purging the binary logs could sometimes cause `mysqld` to crash. (Bug #15854719)

## Release Series Changelogs: MySQL NDB Cluster 7.3

This section contains unified changelog information for the MySQL NDB Cluster 7.3 release series.

For changelogs covering individual MySQL NDB Cluster 7.3 releases, see [NDB Cluster Release Notes](#).

For general information about features added in MySQL NDB Cluster 7.3, see [What is New in NDB Cluster 7.3](#).

For an overview of features added in MySQL 5.6 that are not specific to NDB Cluster, see [What Is New in MySQL 5.6](#). For a complete list of all bug fixes and feature changes made in MySQL 5.6 that are not specific to NDB Cluster, see the MySQL 5.6 [Release Notes](#).

## Changes in MySQL NDB Cluster 7.3.19 (5.6.38-ndb-7.3.19) (2017-10-18, General Availability)

### Bugs Fixed

- Added `DUMP` code 7027 to facilitate testing of issues relating to local checkpoints. For more information, see [DUMP 7027](#). (Bug #26661468)
- A previous fix intended to improve logging of node failure handling in the transaction coordinator included logging of transactions that could occur in normal operation, which made the resulting logs needlessly verbose. Such normal transactions are no longer written to the log in such cases. (Bug #26568782)

References: This issue is a regression of: Bug #26364729.

- Some `DUMP` codes used for the `LGMAN` kernel block were incorrectly assigned numbers in the range used for codes belonging to `DBTUX`. These have now been assigned symbolic constants and numbers in the proper range (10001, 10002, and 10003). (Bug #26365433)

- Node failure handling in the `DBTC` kernel block consists of a number of tasks which execute concurrently, and all of which must complete before TC node failure handling is complete. This fix extends logging coverage to record when each task completes, and which tasks remain, includes the following improvements:
  - Handling interactions between GCP and node failure handling interactions, in which TC takeover causes GCP participant stall at the master TC to allow it to extend the current GCI with any transactions that were taken over; the stall can begin and end in different GCP protocol states. Logging coverage is extended to cover all scenarios. Debug logging is now more consistent and understandable to users.
  - Logging done by the `QMGR` block as it monitors duration of node failure handling duration is done more frequently. A warning log is now generated every 30 seconds (instead of 1 minute), and this now includes `DBDIH` block debug information (formerly this was written separately, and less often).
  - To reduce space used, `DBTC instance number:` is shortened to `DBTC number:`.
  - A new error code is added to assist testing.

(Bug #26364729)

- A potential hundredfold signal fan-out when sending a `START_FRAG_REQ` signal could lead to a node failure due to a `job buffer full` error in start phase 5 while trying to perform a local checkpoint during a restart. (Bug #86675, Bug #26263397)

References: See also: Bug #26288247, Bug #26279522.

## Changes in MySQL NDB Cluster 7.3.18 (5.6.37-ndb-7.3.18) (2017-07-18, General Availability)

### Bugs Fixed

- Backup `.log` files contained log entries for one or more extra fragments, due to an issue with filtering out changes logged by other nodes in the same node group. This resulted in a larger `.log` file and thus use of more resources than necessary; it could also cause problems when restoring, since backups from different nodes could interfere with one another while the log was being applied. (Bug #25891014)
- Error 240 is raised when there is a mismatch between foreign key trigger columns and the values supplied to them during trigger execution, but had no error message indicating the source of the problem. (Bug #23141739)

References: See also: Bug #23068914, Bug #85857.

- `ALTER TABLE .. MAX_ROWS=0` can now be performed only by using a copying `ALTER TABLE` statement. Resetting `MAX_ROWS` to 0 can no longer be performed using `ALGORITHM=INPLACE` or the `ONLINE` keyword. (Bug #21960004)
- When compiling the NDB kernel with `gcc` version 6.0.0 or later, it is now built using `-flifetime-dse=1`. (Bug #85381, Bug #25690926)

## Changes in MySQL NDB Cluster 7.3.17 (5.6.36-ndb-7.3.17) (2017-04-10, General Availability)

## Bugs Fixed

- **NDB Disk Data:** Stale data from NDB Disk Data tables that had been dropped could potentially be included in backups due to the fact that disk scans were enabled for these. To prevent this possibility, disk scans are now disabled—as are other types of scans—when taking a backup. (Bug #84422, Bug #25353234)
- **NDB Disk Data:** In some cases, setting dynamic in-memory columns of an NDB Disk Data table to `NULL` was not handled correctly. (Bug #79253, Bug #22195588)

## Changes in MySQL NDB Cluster 7.3.16 (5.6.35-ndb-7.3.16) (2017-01-17, General Availability)

### Bugs Fixed

- `ndb_restore` did not restore tables having more than 341 columns correctly. This was due to the fact that the buffer used to hold table metadata read from `.ctl` files was of insufficient size, so that only part of the table descriptor could be read from it in such cases. This issue is fixed by increasing the size of the buffer used by `ndb_restore` for file reads. (Bug #25182956)

References: See also: Bug #25302901.

- The `rand()` function was used to produce a unique table ID and table version needed to identify a schema operation distributed between multiple SQL nodes, relying on the assumption that `rand()` would never produce the same numbers on two different instances of `mysqld`. It was later determined that this is not the case, and that in fact it is very likely for the same random numbers to be produced on all SQL nodes.

This fix removes the usage of `rand()` for producing a unique table ID or version, and instead uses a sequence in combination with the node ID of the coordinator. This guarantees uniqueness until the counter for the sequence wraps, which should be sufficient for this purpose.

The effects of this duplication could be observed as timeouts in the log (for example `NDB create db: waiting max 119 sec for distributing`) when restarting multiple `mysqld` processes simultaneously or nearly so, or when issuing the same `CREATE DATABASE` or `DROP DATABASE` statement on multiple SQL nodes. (Bug #24926009)

- Long message buffer exhaustion when firing immediate triggers could result in row ID leaks; this could later result in persistent `RowId already allocated` errors (NDB Error 899). (Bug #23723110)

References: See also: Bug #19506859, Bug #13927679.

- when a parent NDB table in a foreign key relationship was updated, the update cascaded to a child table as expected, but the change was not cascaded to a child table of this child table (that is, to a grandchild of the original parent). This can be illustrated using the tables generated by the following `CREATE TABLE` statements:

```
CREATE TABLE parent(
  id INT PRIMARY KEY AUTO_INCREMENT,
  col1 INT UNIQUE,
  col2 INT
) ENGINE NDB;

CREATE TABLE child(
  refl INT UNIQUE,
  FOREIGN KEY fkl(refl)
  REFERENCES parent(col1) ON UPDATE CASCADE
```



```

) ENGINE NDB;

CREATE TABLE grandchild(
  ref2 INT,
  FOREIGN KEY fk2(ref2)
    REFERENCES child(ref1) ON UPDATE CASCADE
) ENGINE NDB;

```

Table `child` is a child of table `parent`; table `grandchild` is a child of table `child`, and a grandchild of `parent`. In this scenario, a change to column `coll` of `parent` cascaded to `ref1` in table `child`, but it was not always propagated in turn to `ref2` in table `grandchild`. (Bug #83743, Bug #25063506)

## Changes in MySQL NDB Cluster 7.3.15 (5.6.34-ndb-7.3.15) (2016-10-18, General Availability)

### Bugs Fixed

- **NDB Cluster APIs:** Reuse of transaction IDs could occur when `Ndb` objects were created and deleted concurrently. As part of this fix, the NDB API methods `lock_ndb_objects()` and `unlock_ndb_objects` are now declared as `const`. (Bug #23709232)
- Removed an invalid assertion to the effect that all cascading child scans are closed at the time API connection records are released following an abort of the main transaction. The assertion was invalid because closing of scans in such cases is by design asynchronous with respect to the main transaction, which means that subscans may well take some time to close after the main transaction is closed. (Bug #23709284)
- A number of potential buffer overflow issues were found and fixed in the `NDB` codebase. (Bug #23152979)
- When a data node has insufficient redo buffer during a system restart, it does not participate in the restart until after the other nodes have started. After this, it performs a takeover of its fragments from the nodes in its node group that have already started; during this time, the cluster is already running and user activity is possible, including DML and DDL operations.

During a system restart, table creation is handled differently in the `DIH` kernel block than normally, as this creation actually consists of reloading table definition data from disk on the master node. Thus, `DIH` assumed that any table creation that occurred before all nodes had restarted must be related to the restart and thus always on the master node. However, during the takeover, table creation can occur on non-master nodes due to user activity; when this happened, the cluster underwent a forced shutdown.

Now an extra check is made during system restarts to detect in such cases whether the executing node is the master node, and use that information to determine whether the table creation is part of the restart proper, or is taking place during a subsequent takeover. (Bug #23028418)

- When restoring a backup taken from a database containing tables that had foreign keys, `ndb_restore` disabled the foreign keys for data, but not for the logs. (Bug #83155, Bug #24736950)
- Several object constructors and similar functions in the `NDB` codebase did not always perform sanity checks when creating new instances. These checks are now performed under such circumstances. (Bug #77408, Bug #21286722)
- An internal call to `malloc()` was not checked for `NULL`. The function call was replaced with a direct write. (Bug #77375, Bug #21271194)

## Changes in MySQL NDB Cluster 7.3.14 (5.6.31-ndb-7.3.14) (2016-07-18, General Availability)

## Bugs Fixed

- **Incompatible Change:** When the data nodes are only partially connected to the API nodes, a node used for a pushdown join may get its request from a transaction coordinator on a different node, without (yet) being connected to the API node itself. In such cases, the `NodeInfo` object for the requesting API node contained no valid info about the software version of the API node, which caused the `DBSPJ` block to assume (incorrectly) when aborting to assume that the API node used `NDB` version 7.2.4 or earlier, requiring the use of a backward compatibility mode to be used during query abort which sent a node failure error instead of the real error causing the abort.

Now, whenever this situation occurs, it is assumed that, if the `NDB` software version is not yet available, the API node version is greater than 7.2.4. (Bug #23049170)

- **NDB Cluster APIs:** Deletion of `Ndb` objects used a disproportionately high amount of CPU. (Bug #22986823)
- Reserved send buffer for the loopback transporter, introduced in MySQL NDB Cluster 7.4.8 and used by API and management nodes for administrative signals, was calculated incorrectly. (Bug #23093656, Bug #22016081)

References: This issue is a regression of: Bug #21664515.

- During a node restart, re-creation of internal triggers used for verifying the referential integrity of foreign keys was not reliable, because it was possible that not all distributed TC and LDM instances agreed on all trigger identities. To fix this problem, an extra step is added to the node restart sequence, during which the trigger identities are determined by querying the current master node. (Bug #23068914)

References: See also: Bug #23221573.

- Following the forced shutdown of one of the 2 data nodes in a cluster where `NoOfReplicas=2`, the other data node shut down as well, due to arbitration failure. (Bug #23006431)
- `ClusterMgr` is an internal component of `NDB API` and `ndb_mgmd` processes, part of `TransporterFacade`—which in turn is a wrapper around the transporter registry—and shared with data nodes. This component is responsible for a number of tasks including connection setup requests; sending and monitoring of heartbeats; provision of node state information; handling of cluster disconnects and reconnects; and forwarding of cluster state indicators. `ClusterMgr` maintains a count of live nodes which is incremented on receiving a report of a node having connected (`reportConnected()` method call), and decremented on receiving a report that a node has disconnected (`reportDisconnected()`) from `TransporterRegistry`. This count is checked within `reportDisconnected()` to verify that it is greater than zero.

The issue addressed here arose when node connections were very brief due to send buffer exhaustion (among other potential causes) and the check just described failed. This occurred because, when a node did not fully connect, it was still possible for the connection attempt to trigger a `reportDisconnected()` call in spite of the fact that the connection had not yet been reported to `ClusterMgr`; thus, the pairing of `reportConnected()` and `reportDisconnected()` calls was not guaranteed, which could cause the count of connected nodes to be set to zero even though there remained nodes that were still in fact connected, causing node crashes with debug builds of MySQL NDB Cluster, and potential errors or other adverse effects with release builds.

To fix this issue, `ClusterMgr::reportDisconnected()` now verifies that a disconnected node had actually finished connecting completely before checking and decrementing the number of connected nodes. (Bug #21683144, Bug #22016081)

References: See also: Bug #21664515, Bug #21651400.

- To reduce the possibility that a node's loopback transporter becomes disconnected from the transporter registry by `reportError()` due to send buffer exhaustion (implemented by the fix for Bug #21651400), a portion of the send buffer is now reserved for the use of this transporter. (Bug #21664515, Bug #22016081)

References: See also: Bug #21651400, Bug #21683144.

- The loopback transporter is similar to the TCP transporter, but is used by a node to send signals to itself as part of many internal operations. Like the TCP transporter, it could be disconnected due to certain conditions including send buffer exhaustion, but this could result in blocking of `TransporterFacade` and so cause multiple issues within an `ndb_mgmd` or API node process. To prevent this, a node whose loopback transporter becomes disconnected is now simply shut down, rather than allowing the node process to hang. (Bug #21651400, Bug #22016081)

References: See also: Bug #21683144, Bug #21664515.

## Changes in MySQL NDB Cluster 7.3.13 (5.6.29-ndb-7.3.13) (2016-04-18, General Availability)

### Bugs Fixed

- **NDB Cluster APIs:** Executing a transaction with an `NdbIndexOperation` based on an obsolete unique index caused the data node process to fail. Now the index is checked in such cases, and if it cannot be used the transaction fails with an appropriate error. (Bug #79494, Bug #22299443)
- During node failure handling, the request structure used to drive the cleanup operation was not maintained correctly when the request was executed. This led to inconsistencies that were harmless during normal operation, but these could lead to assertion failures during node failure handling, with subsequent failure of additional nodes. (Bug #22643129)
- The previous fix for a lack of mutex protection for the internal `TransporterFacade::deliver_signal()` function was found to be incomplete in some cases. (Bug #22615274)

References: This issue is a regression of: Bug #77225, Bug #21185585.

- When setup of the binary log as an atomic operation on one SQL node failed, this could trigger a state in other SQL nodes in which they appeared to detect the SQL node participating in schema change distribution, whereas it had not yet completed binary log setup. This could in turn cause a deadlock on the global metadata lock when the SQL node still retrying binary log setup needed this lock, while another mysqld had taken the lock for itself as part of a schema change operation. In such cases, the second SQL node waited for the first one to act on its schema distribution changes, which it was not yet able to do. (Bug #22494024)
- Duplicate key errors could occur when `ndb_restore` was run on a backup containing a unique index. This was due to the fact that, during restoration of data, the database can pass through one or more inconsistent states prior to completion, such an inconsistent state possibly having duplicate values for a column which has a unique index. (If the restoration of data is preceded by a run with `--disable-indexes` and followed by one with `--rebuild-indexes`, these errors are avoided.)

Added a check for unique indexes in the backup which is performed only when restoring data, and which does not process tables that have explicitly been excluded. For each unique index found, a warning is now printed. (Bug #22329365)

- Restoration of metadata with `ndb_restore -m` occasionally failed with the error message `Failed to create index...` when creating a unique index. While diagnosing this problem, it was found that the

internal error `PREPARE_SEIZE_ERROR` (a temporary error) was reported as an unknown error. Now in such cases, `ndb_restore` retries the creation of the unique index, and `PREPARE_SEIZE_ERROR` is reported as NDB Error 748 *Busy during read of event table*. (Bug #21178339)

References: See also: Bug #22989944.

- Optimization of signal sending by buffering and sending them periodically, or when the buffer became full, could cause `SUB_GCP_COMPLETE_ACK` signals to be excessively delayed. Such signals are sent for each node and epoch, with a minimum interval of `TimeBetweenEpochs`; if they are not received in time, the `SUMA` buffers can overflow as a result. The overflow caused API nodes to be disconnected, leading to current transactions being aborted due to node failure. This condition made it difficult for long transactions (such as altering a very large table), to be completed. Now in such cases, the `ACK` signal is sent without being delayed. (Bug #18753341)

## Changes in MySQL NDB Cluster 7.3.12 (5.6.28-ndb-7.3.12) (2016-01-19, General Availability)

### Bugs Fixed

- **Important Change:** A fix made in MySQL NDB Cluster 7.3.11 and MySQL NDB Cluster 7.4.8 caused `ndb_restore` to perform unique key checks even when operating in modes which do not restore data, such as when using the program's `--restore_epoch` or `--print_data` option.

That change in behavior caused existing valid backup routines to fail; to keep this issue from affecting this and future releases, the previous fix has been reverted. This means that the requirement added in those versions that `ndb_restore` be run `--disable-indexes` or `--rebuild-indexes` when used on tables containing unique indexes is also lifted. (Bug #22345748)

References: See also: Bug #22329365. Reverted patches: Bug #57782, Bug #11764893.

- **Important Note:** If an NDB table having a foreign key was dropped while one of the data nodes was stopped, the data node later failed when trying to restart. (Bug #18554390)
- **NDB Cluster APIs:** The binary log injector did not work correctly with `TE_INCONSISTENT` event type handling by `Ndb::nextEvent()`. (Bug #22135541)

References: See also: Bug #20646496.

- **NDB Cluster APIs:** `Ndb::pollEvents()` and `pollEvents2()` were slow to receive events, being dependent on other client threads or blocks to perform polling of transporters on their behalf. This fix allows a client thread to perform its own transporter polling when it has to wait in either of these methods.

Introduction of transporter polling also revealed a problem with missing mutex protection in the `ndbcluster_binlog` handler, which has been added as part of this fix. (Bug #79311, Bug #20957068, Bug #22224571)

- In debug builds, a `WAIT_EVENT` while polling caused excessive logging to stdout. (Bug #22203672)
- When executing a schema operation such as `CREATE TABLE` on a MySQL NDB Cluster with multiple SQL nodes, it was possible for the SQL node on which the operation was performed to time out while waiting for an acknowledgement from the others. This could occur when different SQL nodes had different settings for `--ndb-log-updated-only`, `--ndb-log-update-as-write`, or other `mysqld` options effecting binary logging by `NDB`.

This happened due to the fact that, in order to distribute schema changes between them, all SQL nodes subscribe to changes in the `ndb_schema` system table, and that all SQL nodes are made aware of

each others subscriptions by subscribing to `TE_SUBSCRIBE` and `TE_UNSUBSCRIBE` events. The names of events to subscribe to are constructed from the table names, adding `REPL$` or `REPLF$` as a prefix. `REPLF$` is used when full binary logging is specified for the table. The issue described previously arose because different values for the options mentioned could lead to different events being subscribed to by different SQL nodes, meaning that all SQL nodes were not necessarily aware of each other, so that the code that handled waiting for schema distribution to complete did not work as designed.

To fix this issue, MySQL NDB Cluster now treats the `ndb_schema` table as a special case and enforces full binary logging at all times for this table, independent of any settings for `mysqld` binary logging options. (Bug #22174287, Bug #79188)

- Attempting to create an NDB table having greater than the maximum supported combined width for all `BIT` columns (4096) caused data node failure when these columns were defined with `COLUMN_FORMAT DYNAMIC`. (Bug #21889267)
- Creating a table with the maximum supported number of columns (512) all using `COLUMN_FORMAT DYNAMIC` led to data node failures. (Bug #21863798)
- Using `ndb_mgm STOP -f` to force a node shutdown even when it triggered a complete shutdown of the cluster, it was possible to lose data when a sufficient number of nodes were shut down, triggering a cluster shutdown, and the timing was such that `SUMA` handovers had been made to nodes already in the process of shutting down. (Bug #17772138)
- The internal `NdbEventBuffer::set_total_buckets()` method calculated the number of remaining buckets incorrectly. This caused any incomplete epoch to be prematurely completed when the `SUB_START_CONF` signal arrived out of order. Any events belonging to this epoch arriving later were then ignored, and so effectively lost, which resulted in schema changes not being distributed correctly among SQL nodes. (Bug #79635, Bug #22363510)
- Compilation of MySQL NDB Cluster failed on SUSE Linux Enterprise Server 12. (Bug #79429, Bug #22292329)
- Schema events were appended to the binary log out of order relative to non-schema events. This was caused by the fact that the binary log injector did not properly handle the case where schema events and non-schema events were from different epochs.

This fix modifies the handling of events from the two schema and non-schema event streams such that events are now always handled one epoch at a time, starting with events from the oldest available epoch, without regard to the event stream in which they occur. (Bug #79077, Bug #22135584, Bug #20456664)

- NDB failed during a node restart due to the status of the current local checkpoint being set but not as active, even though it could have other states under such conditions. (Bug #78780, Bug #21973758)
- The value set for `spintime` by the `ThreadConfig` parameter was not calculated correctly, causing the spin to continue for longer than actually specified. (Bug #78525, Bug #21886476)

## Changes in MySQL NDB Cluster 7.3.11 (5.6.27-ndb-7.3.11) (2015-10-19, General Availability)

### Bugs Fixed

- **Important Change:** When `ndb_restore` was run without `--disable-indexes` or `--rebuild-indexes` on a table having a unique index, it was possible for rows to be restored in an order that resulted in duplicate values, causing it to fail with duplicate key errors. Running `ndb_restore` on such a table now requires using at least one of these options; failing to do so now results in an error. (Bug #57782, Bug #11764893)

References: See also: Bug #22329365, Bug #22345748.

- **NDB Cluster APIs:** While executing `dropEvent()`, if the coordinator `DBDICT` failed after the subscription manager (`SUMA` block) had removed all subscriptions but before the coordinator had deleted the event from the system table, the dropped event remained in the table, causing any subsequent drop or create event with the same name to fail with `NDB` error 1419 `Subscription already dropped` or error 746 `Event name already exists`. This occurred even when calling `dropEvent()` with a nonzero force argument.

Now in such cases, error 1419 is ignored, and `DBDICT` deletes the event from the table. (Bug #21554676)

- **NDB Cluster APIs:** The internal value representing the latest global checkpoint was not always updated when a completed epoch of event buffers was inserted into the event queue. This caused subsequent calls to `Ndb::pollEvents()` and `pollEvents2()` to fail when trying to obtain the correct GCI for the events available in the event buffers. This could also result in later calls to `nextEvent()` or `nextEvent2()` seeing events that had not yet been discovered. (Bug #78129, Bug #21651536)
- Backup block states were reported incorrectly during backups. (Bug #21360188)

References: See also: Bug #20204854, Bug #21372136.

- When a data node is known to have been alive by other nodes in the cluster at a given global checkpoint, but its `sysfile` reports a lower GCI, the higher GCI is used to determine which global checkpoint the data node can recreate. This caused problems when the data node being started had a clean file system (GCI = 0), or when it was more than more global checkpoint behind the other nodes.

Now in such cases a higher GCI known by other nodes is used only when it is at most one GCI ahead. (Bug #19633824)

References: See also: Bug #20334650, Bug #21899993. This issue is a regression of: Bug #29167.

- When restoring a specific database or databases with the `--include-databases` or `--exclude-databases` option, `ndb_restore` attempted to apply foreign keys on tables in databases which were not among those being restored. (Bug #18560951)
- After restoring the database schema from backup using `ndb_restore`, auto-discovery of restored tables in transactions having multiple statements did not work correctly, resulting in `Deadlock found when trying to get lock; try restarting transaction` errors.

This issue was encountered both in the `mysql` client, as well as when such transactions were executed by application programs using Connector/J and possibly other MySQL APIs.

Prior to upgrading, this issue can be worked around by executing `SELECT TABLE_NAME, TABLE_SCHEMA FROM INFORMATION_SCHEMA.TABLES WHERE ENGINE = 'NDBCLUSTER'` on all SQL nodes following the restore operation, before executing any other statements. (Bug #18075170)

- `ndb_desc` used with the `--extra-partition-info` and `--blob-info` options failed when run against a table containing one or more `TINYBLOB` columns. (Bug #14695968)
- Trying to create an `NDB` table with a composite foreign key referencing a composite primary key of the parent table failed when one of the columns in the composite foreign key was the table's primary key and in addition this column also had a unique key. (Bug #78150, Bug #21664899)
- When attempting to enable index statistics, creation of the required system tables, events and event subscriptions often fails when multiple `mysqld` processes using index statistics are started concurrently

in conjunction with starting, restarting, or stopping the cluster, or with node failure handling. This is normally recoverable, since the affected `mysqld` process or processes can (and do) retry these operations shortly thereafter. For this reason, such failures are no longer logged as warnings, but merely as informational events. (Bug #77760, Bug #21462846)

- Adding a unique key to an `NDB` table failed when the table already had a foreign key. Prior to upgrading, you can work around this issue by creating the unique key first, then adding the foreign key afterwards, using a separate `ALTER TABLE` statement. (Bug #77457, Bug #20309828)

## Changes in MySQL NDB Cluster 7.3.10 (5.6.25-ndb-7.3.10) (2015-07-13, General Availability)

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

### Functionality Added or Changed

- A number of improvements, listed here, have been made with regard to handling issues that could arise when an overload arose due to a great number of inserts being performed during a local checkpoint (LCP):
  - Failures sometimes occurred during restart processing when trying to execute the undo log, due to a problem with finding the end of the log. This happened when there remained unwritten pages at the end of the first undo file when writing to the second undo file, which caused the execution of undo logs in reverse order and so execute old or even nonexistent log records.

This is fixed by ensuring that execution of the undo log begins with the proper end of the log, and, if started earlier, that any unwritten or faulty pages are ignored.

- It was possible to fail during an LCP, or when performing a `COPY_FRAGREQ`, due to running out of operation records. We fix this by making sure that LCPs and `COPY_FRAG` use resources reserved for operation records, as was already the case with scan records. In addition, old code for ACC operations that was no longer required but that could lead to failures was removed.
- When an LCP was performed while loading a table, it was possible to hit a livelock during LCP scans, due to the fact that each record that was inserted into new pages after the LCP had started had its `LCP_SKIP` flag set. Such records were discarded as intended by the LCP scan, but when inserts occurred faster than the LCP scan could discard records, the scan appeared to hang. As part of this issue, the scan failed to report any progress to the LCP watchdog, which after 70 seconds of livelock killed the process. This issue was observed when performing on the order of 250000 inserts per second over an extended period of time (120 seconds or more), using a single LDM.

This part of the fix makes a number of changes, listed here:

- We now ensure that pages created after the LCP has started are not included in LCP scans; we also ensure that no records inserted into those pages have their `LCP_SKIP` flag set.
- Handling of the scan protocol is changed such that a certain amount of progress is made by the LCP regardless of load; we now report progress to the LCP watchdog so that we avoid failure in the event that an LCP is making progress but not writing any records.
- We now take steps to guarantee that LCP scans proceed more quickly than inserts can occur, by ensuring that scans are prioritized this scanning activity, and thus, that the LCP is in fact (eventually) completed.

- In addition, scanning is made more efficient, by prefetching tuples; this helps avoid stalls while fetching memory in the CPU.
- Row checksums for preventing data corruption now include the tuple header bits.

(Bug #76373, Bug #20727343, Bug #76741, Bug #69994, Bug #20903880, Bug #76742, Bug #20904721, Bug #76883, Bug #20980229)

## Bugs Fixed

- **Important Change; NDB Cluster APIs:** Added the method `Ndb::isExpectingHigherQueuedEpochs()` to the NDB API to detect when additional, newer event epochs were detected by `pollEvents2()`.

The behavior of `Ndb::pollEvents()` has also been modified such that it now returns `NDB_FAILURE_GCI` (equal to `~(Uint64) 0`) when a cluster failure has been detected. (Bug #18753887)

- **NDB Cluster APIs:** Added the `Column::getSizeInBytesForRecord()` method, which returns the size required for a column by an `NdbRecord`, depending on the column's type (text/blob, or other). (Bug #21067283)
- **NDB Cluster APIs:** Creation and destruction of `Ndb_cluster_connection` objects by multiple threads could make use of the same application lock, which in some cases led to failures in the global dictionary cache. To alleviate this problem, the creation and destruction of several internal NDB API objects have been serialized. (Bug #20636124)
- **NDB Cluster APIs:** A number of timeouts were not handled correctly in the NDB API. (Bug #20617891)
- **NDB Cluster APIs:** When an `Ndb` object created prior to a failure of the cluster was reused, the event queue of this object could still contain data node events originating from before the failure. These events could reference “old” epochs (from before the failure occurred), which in turn could violate the assumption made by the `nextEvent()` method that epoch numbers always increase. This issue is addressed by explicitly clearing the event queue in such cases. (Bug #18411034)

References: See also: Bug #20888668.

- After restoring the database metadata (but not any data) by running `ndb_restore --restore_meta` (or `-m`), SQL nodes would hang while trying to `SELECT` from a table in the database to which the metadata was restored. In such cases the attempt to query the table now fails as expected, since the table does not actually exist until `ndb_restore` is executed with `--restore_data` (`-r`). (Bug #21184102)

References: See also: Bug #16890703.

- When a great many threads opened and closed blocks in the NDB API in rapid succession, the internal `close_clnt()` function synchronizing the closing of the blocks waited an insufficiently long time for a self-signal indicating potential additional signals needing to be processed. This led to excessive CPU usage by `ndb_mgmd`, and prevented other threads from opening or closing other blocks. This issue is fixed by changing the function polling call to wait on a specific condition to be woken up (that is, when a signal has in fact been executed). (Bug #21141495)
- Previously, multiple send threads could be invoked for handling sends to the same node; these threads then competed for the same send lock. While the send lock blocked the additional send threads, work threads could be passed to other nodes.



This issue is fixed by ensuring that new send threads are not activated while there is already an active send thread assigned to the same node. In addition, a node already having an active send thread assigned to it is no longer visible to other, already active, send threads; that is, such a node is longer added to the node list when a send thread is currently assigned to it. (Bug #20954804, Bug #76821)

- Queueing of pending operations when the redo log was overloaded (`DefaultOperationRedoProblemAction` API node configuration parameter) could lead to timeouts when data nodes ran out of redo log space (`P_TAIL_PROBLEM` errors). Now when the redo log is full, the node aborts requests instead of queuing them. (Bug #20782580)

References: See also: Bug #20481140.

- NDB statistics queries could be delayed by the error delay set for `ndb_index_stat_option` (default 60 seconds) when the index that was queried had been marked with internal error. The same underlying issue could also cause `ANALYZE TABLE` to hang when executed against an NDB table having multiple indexes where an internal error occurred on one or more but not all indexes.

Now in such cases, any existing statistics are returned immediately, without waiting for any additional statistics to be discovered. (Bug #20553313, Bug #20707694, Bug #76325)

- The multithreaded scheduler sends to remote nodes either directly from each worker thread or from dedicated send threadsL, depending on the cluster's configuration. This send might transmit all, part, or none of the available data from the send buffers. While there remained pending send data, the worker or send threads continued trying to send in a loop. The actual size of the data sent in the most recent attempt to perform a send is now tracked, and used to detect lack of send progress by the send or worker threads. When no progress has been made, and there is no other work outstanding, the scheduler takes a 1 millisecond pause to free up the CPU for use by other threads. (Bug #18390321)

References: See also: Bug #20929176, Bug #20954804.

- In some cases, attempting to restore a table that was previously backed up failed with a `File Not Found` error due to a missing table fragment file. This occurred as a result of the NDB kernel `BACKUP` block receiving a `Busy` error while trying to obtain the table description, due to other traffic from external clients, and not retrying the operation.

The fix for this issue creates two separate queues for such requests—one for internal clients such as the `BACKUP` block or `ndb_restore`, and one for external clients such as API nodes—and prioritizing the internal queue.

Note that it has always been the case that external client applications using the NDB API (including MySQL applications running against an SQL node) are expected to handle `Busy` errors by retrying transactions at a later time; this expectation is *not* changed by the fix for this issue. (Bug #17878183)

References: See also: Bug #17916243.

- In some cases, the `DBDICTIONARY` block failed to handle repeated `GET_TABINFOREQ` signals after the first one, leading to possible node failures and restarts. This could be observed after setting a sufficiently high value for `MaxNoOfExecutionThreads` and low value for `LcpScanProgressTimeout`. (Bug #77433, Bug #21297221)
- Client lookup for delivery of API signals to the correct client by the internal `TransporterFacade::deliver_signal()` function had no mutex protection, which could cause issues such as timeouts encountered during testing, when other clients connected to the same `TransporterFacade`. (Bug #77225, Bug #21185585)

- It was possible to end up with a lock on the send buffer mutex when send buffers became a limiting resource, due either to insufficient send buffer resource configuration, problems with slow or failing communications such that all send buffers became exhausted, or slow receivers failing to consume what was sent. In this situation worker threads failed to allocate send buffer memory for signals, and attempted to force a send in order to free up space, while at the same time the send thread was busy trying to send to the same node or nodes. All of these threads competed for taking the send buffer mutex, which resulted in the lock already described, reported by the watchdog as `Stuck in Send`. This fix is made in two parts, listed here:
  1. The send thread no longer holds the global send thread mutex while getting the send buffer mutex; it now releases the global mutex prior to locking the send buffer mutex. This keeps worker threads from getting stuck in send in such cases.
  2. Locking of the send buffer mutex done by the send threads now uses a try-lock. If the try-lock fails, the node to make the send to is reinserted at the end of the list of send nodes in order to be retried later. This removes the `Stuck in Send` condition for the send threads.

(Bug #77081, Bug #21109605)

## Changes in MySQL NDB Cluster 7.3.9 (5.6.24-ndb-7.3.9) (2015-04-14, General Availability)

### Bugs Fixed

- **Important Change:** The maximum failure time calculation used to ensure that normal node failure handling mechanisms are given time to handle survivable cluster failures (before global checkpoint watchdog mechanisms start to kill nodes due to GCP delays) was excessively conservative, and neglected to consider that there can be at most  $\text{number\_of\_data\_nodes} / \text{NoOfReplicas}$  node failures before the cluster can no longer survive. Now the value of `NoOfReplicas` is properly taken into account when performing this calculation.

This fix adds the `TimeBetweenGlobalCheckpointsTimeout` data node configuration parameter, which makes the minimum timeout between global checkpoints settable by the user. This timeout was previously fixed internally at 120000 milliseconds, which is now the default value for this parameter. (Bug #20069617, Bug #20069624)

References: See also: Bug #19858151, Bug #20128256, Bug #20135976.

- **NDB Cluster APIs:** When a transaction is started from a cluster connection, `Table` and `Index` schema objects may be passed to this transaction for use. If these schema objects have been acquired from a different connection (`Ndb_cluster_connection` object), they can be deleted at any point by the deletion or disconnection of the owning connection. This can leave a connection with invalid schema objects, which causes an NDB API application to fail when these are dereferenced.

To avoid this problem, if your application uses multiple connections, you can now set a check to detect sharing of schema objects between connections when passing a schema object to a transaction, using the `NdbTransaction::setSchemaObjectOwnerChecks()` method added in this release. When this check is enabled, the schema objects having the same names are acquired from the connection and compared to the schema objects passed to the transaction. Failure to match causes the application to fail with an error. (Bug #19785977)

- **NDB Cluster APIs:** The increase in the default number of hashmap buckets (`DefaultHashMapSize` API node configuration parameter) from 240 to 3480 in MySQL NDB Cluster 7.2.11 increased the size of the internal `DictHashMapInfo::HashMap` type considerably. This type was allocated on the stack in some `getTable()` calls which could lead to stack overflow issues for NDB API users.

To avoid this problem, the hashmap is now dynamically allocated from the heap. (Bug #19306793)

- **NDB Cluster APIs:** A scan operation, whether it is a single table scan or a query scan used by a pushed join, stores the result set in a buffer. This maximum size of this buffer is calculated and preallocated before the scan operation is started. This buffer may consume a considerable amount of memory; in some cases we observed a 2 GB buffer footprint in tests that executed 100 parallel scans with 2 single-threaded ([ndbd](#)) data nodes. This memory consumption was found to scale linearly with additional fragments.

A number of root causes, listed here, were discovered that led to this problem:

- Result rows were unpacked to full [NdbRecord](#) format before they were stored in the buffer. If only some but not all columns of a table were selected, the buffer contained empty space (essentially wasted).
- Due to the buffer format being unpacked, [VARCHAR](#) and [VARBINARY](#) columns always had to be allocated for the maximum size defined for such columns.
- [BatchByteSize](#) and [MaxScanBatchSize](#) values were not taken into consideration as a limiting factor when calculating the maximum buffer size.

These issues became more evident in NDB 7.2 and later MySQL NDB Cluster release series. This was due to the fact buffer size is scaled by [BatchSize](#), and that the default value for this parameter was increased fourfold (from 64 to 256) beginning with MySQL NDB Cluster 7.2.1.

This fix causes result rows to be buffered using the packed format instead of the unpacked format; a buffered scan result row is now not unpacked until it becomes the current row. In addition, [BatchByteSize](#) and [MaxScanBatchSize](#) are now used as limiting factors when calculating the required buffer size.

Also as part of this fix, refactoring has been done to separate handling of buffered (packed) from handling of unbuffered result sets, and to remove code that had been unused since NDB 7.0 or earlier. The [NdbRecord](#) class declaration has also been cleaned up by removing a number of unused or redundant member variables. (Bug #73781, Bug #75599, Bug #19631350, Bug #20408733)

- It was found during testing that problems could arise when the node registered as the arbitrator disconnected or failed during the arbitration process.

In this situation, the node requesting arbitration could never receive a positive acknowledgement from the registered arbitrator; this node also lacked a stable set of members and could not initiate selection of a new arbitrator.

Now in such cases, when the arbitrator fails or loses contact during arbitration, the requesting node immediately fails rather than waiting to time out. (Bug #20538179)

- The values of the [Ndb\\_last\\_commit\\_epoch\\_server](#) and [Ndb\\_last\\_commit\\_epoch\\_session](#) status variables were incorrectly reported on some platforms. To correct this problem, these values are now stored internally as `long long`, rather than `long`. (Bug #20372169)
- When a data node fails or is being restarted, the remaining nodes in the same nodegroup resend to subscribers any data which they determine has not already been sent by the failed node. Normally, when a data node (actually, the [SUMA](#) kernel block) has sent all data belonging to an epoch for which it is responsible, it sends a [SUB\\_GCP\\_COMPLETE\\_REP](#) signal, together with a count, to all subscribers, each of which responds with a [SUB\\_GCP\\_COMPLETE\\_ACK](#). When [SUMA](#) receives this acknowledgment from all subscribers, it reports this to the other nodes in the same nodegroup so that they know that there is no need to resend this data in case of a subsequent node failure. If a node failed before all subscribers

sent this acknowledgement but before all the other nodes in the same nodegroup received it from the failing node, data for some epochs could be sent (and reported as complete) twice, which could lead to an unplanned shutdown.

The fix for this issue adds to the count reported by `SUB_GCP_COMPLETE_ACK` a list of identifiers which the receiver can use to keep track of which buckets are completed and to ignore any duplicate reported for an already completed bucket. (Bug #17579998)

- When performing a restart, it was sometimes possible to find a log end marker which had been written by a previous restart, and that should have been invalidated. Now when searching for the last page to invalidate, the same search algorithm is used as when searching for the last page of the log to read. (Bug #76207, Bug #20665205)
- When reading and copying transporter short signal data, it was possible for the data to be copied back to the same signal with overlapping memory. (Bug #75930, Bug #20553247)
- When a bulk delete operation was committed early to avoid an additional round trip, while also returning the number of affected rows, but failed with a timeout error, an SQL node performed no verification that the transaction was in the Committed state. (Bug #74494, Bug #20092754)

References: See also: Bug #19873609.

- An `ALTER TABLE` statement containing comments and a partitioning option against an NDB table caused the SQL node on which it was executed to fail. (Bug #74022, Bug #19667566)

## Changes in MySQL NDB Cluster 7.3.8 (5.6.22-ndb-7.3.8) (2015-01-21, General Availability)

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

### Functionality Added or Changed

- **Performance:** Recent improvements made to the multithreaded scheduler were intended to optimize the cache behavior of its internal data structures, with members of these structures placed such that those local to a given thread do not overflow into a cache line which can be accessed by another thread. Where required, extra padding bytes are inserted to isolate cache lines owned (or shared) by other threads, thus avoiding invalidation of the entire cache line if another thread writes into a cache line not entirely owned by itself. This optimization improved MT Scheduler performance by several percent.

It has since been found that the optimization just described depends on the global instance of struct `thr_repository` starting at a cache line aligned base address as well as the compiler not rearranging or adding extra padding to the scheduler struct; it was also found that these prerequisites were not guaranteed (or even checked). Thus this cache line optimization has previously worked only when `g_thr_repository` (that is, the global instance) ended up being cache line aligned only by accident. In addition, on 64-bit platforms, the compiler added extra padding words in struct `thr_safe_pool` such that attempts to pad it to a cache line aligned size failed.

The current fix ensures that `g_thr_repository` is constructed on a cache line aligned address, and the constructors modified so as to verify cacheline aligned addresses where these are assumed by design.

Results from internal testing show improvements in MT Scheduler read performance of up to 10% in some cases, following these changes. (Bug #18352514)

- **NDB Cluster APIs:** Two new example programs, demonstrating reads and writes of `CHAR`, `VARCHAR`, and `VARBINARY` column values, have been added to `storage/ndb/ndbapi-examples` in the MySQL NDB Cluster source tree. For more information about these programs, including source code listings, see [NDB API Simple Array Example](#), and [NDB API Simple Array Example Using Adapter](#).

## Bugs Fixed

- **NDB Disk Data:** An update on many rows of a large Disk Data table could in some rare cases lead to node failure. In the event that such problems are observed with very large transactions on Disk Data tables you can now increase the number of page entries allocated for disk page buffer memory by raising the value of the `DiskPageBufferEntries` data node configuration parameter added in this release. (Bug #19958804)
- **NDB Disk Data:** In some cases, during `DICT` master takeover, the new master could crash while attempting to roll forward an ongoing schema transaction. (Bug #19875663, Bug #74510)
- **NDB Disk Data:** When a node acting as a `DICT` master fails, the arbitrator selects another node to take over in place of the failed node. During the takeover procedure, which includes cleaning up any schema transactions which are still open when the master failed, the disposition of the uncommitted schema transaction is decided. Normally this transaction be rolled back, but if it has completed a sufficient portion of a commit request, the new master finishes processing the commit. Until the fate of the transaction has been decided, no new `TRANS_END_REQ` messages from clients can be processed. In addition, since multiple concurrent schema transactions are not supported, takeover cleanup must be completed before any new transactions can be started.

A similar restriction applies to any schema operations which are performed in the scope of an open schema transaction. The counter used to coordinate schema operation across all nodes is employed both during takeover processing and when executing any non-local schema operations. This means that starting a schema operation while its schema transaction is in the takeover phase causes this counter to be overwritten by concurrent uses, with unpredictable results.

The scenarios just described were handled previously using a pseudo-random delay when recovering from a node failure. Now we check before the new master has rolled forward or backwards any schema transactions remaining after the failure of the previous master and avoid starting new schema transactions or performing operations using old transactions until takeover processing has cleaned up after the abandoned transaction. (Bug #19874809, Bug #74503)

- **NDB Disk Data:** When a node acting as `DICT` master fails, it is still possible to request that any open schema transaction be either committed or aborted by sending this request to the new `DICT` master. In this event, the new master takes over the schema transaction and reports back on whether the commit or abort request succeeded. In certain cases, it was possible for the new master to be misidentified—that is, the request was sent to the wrong node, which responded with an error that was interpreted by the client application as an aborted schema transaction, even in cases where the transaction could have been successfully committed, had the correct node been contacted. (Bug #74521, Bug #19880747)
- **NDB Cluster APIs:** It was possible to delete an `Ndb_cluster_connection` object while there remained instances of `Ndb` using references to it. Now the `Ndb_cluster_connection` destructor waits for all related `Ndb` objects to be released before completing. (Bug #19999242)

References: See also: Bug #19846392.

- **NDB Cluster APIs:** The buffer allocated by an `NdbScanOperation` for receiving scanned rows was not released until the `NdbTransaction` owning the scan operation was closed. This could lead to excessive memory usage in an application where multiple scans were created within the same transaction, even if these scans were closed at the end of their lifecycle, unless `NdbScanOperation::close()` was invoked with the `releaseOp` argument equal to

`true`. Now the buffer is released whenever the cursor navigating the result set is closed with `NdbScanOperation::close()`, regardless of the value of this argument. (Bug #75128, Bug #20166585)

- The global checkpoint commit and save protocols can be delayed by various causes, including slow disk I/O. The `DIH` master node monitors the progress of both of these protocols, and can enforce a maximum lag time during which the protocols are stalled by killing the node responsible for the lag when it reaches this maximum. This `DIH` master GCP monitor mechanism did not perform its task more than once per master node; that is, it failed to continue monitoring after detecting and handling a GCP stop. (Bug #20128256)

References: See also: Bug #19858151, Bug #20069617, Bug #20062754.

- When running `mysql_upgrade` on a MySQL NDB Cluster SQL node, the expected drop of the `performance_schema` database on this node was instead performed on all SQL nodes connected to the cluster. (Bug #20032861)
- A number of problems relating to the fired triggers pool have been fixed, including the following issues:
  - When the fired triggers pool was exhausted, `NDB` returned Error 218 (`Out of LongMessageBuffer`). A new error code 221 is added to cover this case.
  - An additional, separate case in which Error 218 was wrongly reported now returns the correct error.
  - Setting low values for `MaxNoOfFiredTriggers` led to an error when no memory was allocated if there was only one hash bucket.
  - An aborted transaction now releases any fired trigger records it held. Previously, these records were held until its `ApiConnectRecord` was reused by another transaction.
  - In addition, for the `Fired Triggers` pool in the internal `ndbinfo.ndb$poools` table, the high value always equalled the total, due to the fact that all records were momentarily seized when initializing them. Now the high value shows the maximum following completion of initialization.

(Bug #19976428)

- Online reorganization when using `ndbmtd` data nodes and with binary logging by `mysqld` enabled could sometimes lead to failures in the `TRIX` and `DBLQH` kernel blocks, or in silent data corruption. (Bug #19903481)

References: See also: Bug #19912988.

- The local checkpoint scan fragment watchdog and the global checkpoint monitor can each exclude a node when it is too slow when participating in their respective protocols. This exclusion was implemented by simply asking the failing node to shut down, which in case this was delayed (for whatever reason) could prolong the duration of the GCP or LCP stall for other, unaffected nodes.

To minimize this time, an isolation mechanism has been added to both protocols whereby any other live nodes forcibly disconnect the failing node after a predetermined amount of time. This allows the failing node the opportunity to shut down gracefully (after logging debugging and other information) if possible, but limits the time that other nodes must wait for this to occur. Now, once the remaining live nodes have processed the disconnection of any failing nodes, they can commence failure handling and restart the related protocol or protocol, even if the failed node takes an excessively long time to shut down. (Bug #19858151)

References: See also: Bug #20128256, Bug #20069617, Bug #20062754.

- A watchdog failure resulted from a hang while freeing a disk page in `TUP_COMMITREQ`, due to use of an uninitialized block variable. (Bug #19815044, Bug #74380)
- Multiple threads crashing led to multiple sets of trace files being printed and possibly to deadlocks. (Bug #19724313)
- When a client retried against a new master a schema transaction that failed previously against the previous master while the latter was restarting, the lock obtained by this transaction on the new master prevented the previous master from progressing past start phase 3 until the client was terminated, and resources held by it were cleaned up. (Bug #19712569, Bug #74154)
- When using the `NDB` storage engine, the maximum possible length of a database or table name is 63 characters, but this limit was not always strictly enforced. This meant that a statement using a name having 64 characters such `CREATE DATABASE`, `DROP DATABASE`, or `ALTER TABLE RENAME` could cause the SQL node on which it was executed to fail. Now such statements fail with an appropriate error message. (Bug #19550973)
- When a new data node started, API nodes were allowed to attempt to register themselves with the data node for executing transactions before the data node was ready. This forced the API node to wait an extra heartbeat interval before trying again.

To address this issue, a number of `HA_ERR_NO_CONNECTION` errors (Error 4009) that could be issued during this time have been changed to `Cluster temporarily unavailable` errors (Error 4035), which should allow API nodes to use new data nodes more quickly than before. As part of this fix, some errors which were incorrectly categorised have been moved into the correct categories, and some errors which are no longer used have been removed. (Bug #19524096, Bug #73758)

- When executing very large pushdown joins involving one or more indexes each defined over several columns, it was possible in some cases for the `DBSPJ` block (see [The DBSPJ Block](#)) in the `NDB` kernel to generate `SCAN_FRAGREQ` signals that were excessively large. This caused data nodes to fail when these could not be handled correctly, due to a hard limit in the kernel on the size of such signals (32K). This fix bypasses that limitation by breaking up `SCAN_FRAGREQ` data that is too large for one such signal, and sending the `SCAN_FRAGREQ` as a chunked or fragmented signal instead. (Bug #19390895)
- `ndb_index_stat` sometimes failed when used against a table containing unique indexes. (Bug #18715165)
- Queries against tables containing a `CHAR(0)` columns failed with `ERROR 1296 (HY000): Got error 4547 'RecordSpecification has overlapping offsets' from NDBCLUSTER`. (Bug #14798022)
- In the `NDB` kernel, it was possible for a `TransporterFacade` object to reset a buffer while the data contained by the buffer was being sent, which could lead to a race condition. (Bug #75041, Bug #20112981)
- `mysql_upgrade` failed to drop and recreate the `ndbinfo` database and its tables as expected. (Bug #74863, Bug #20031425)
- Due to a lack of memory barriers, MySQL NDB Cluster programs such as `ndbmt.d` did not compile on `POWER` platforms. (Bug #74782, Bug #20007248)
- In some cases, when run against a table having an `AFTER DELETE` trigger, a `DELETE` statement that matched no rows still caused the trigger to execute. (Bug #74751, Bug #19992856)
- A basic requirement of the `NDB` storage engine's design is that the transporter registry not attempt to receive data (`TransporterRegistry::performReceive()`) from and update the connection status (`TransporterRegistry::update_connections()`) of the same set of transporters concurrently, due to the fact that the updates perform final cleanup and reinitialization of buffers used when receiving

data. Changing the contents of these buffers while reading or writing to them could lead to "garbage" or inconsistent signals being read or written.

During the course of work done previously to improve the implementation of the transporter facade, a mutex intended to protect against the concurrent use of the `performReceive()` and `update_connections()` methods on the same transporter was inadvertently removed. This fix adds a watchdog check for concurrent usage. In addition, `update_connections()` and `performReceive()` calls are now serialized together while polling the transporters. (Bug #74011, Bug #19661543)

- `ndb_restore` failed while restoring a table which contained both a built-in conversion on the primary key and a staging conversion on a `TEXT` column.

During staging, a `BLOB` table is created with a primary key column of the target type. However, a conversion function was not provided to convert the primary key values before loading them into the staging blob table, which resulted in corrupted primary key values in the staging `BLOB` table. While moving data from the staging table to the target table, the `BLOB` read failed because it could not find the primary key in the `BLOB` table.

Now all `BLOB` tables are checked to see whether there are conversions on primary keys of their main tables. This check is done after all the main tables are processed, so that conversion functions and parameters have already been set for the main tables. Any conversion functions and parameters used for the primary key in the main table are now duplicated in the `BLOB` table. (Bug #73966, Bug #19642978)

- Corrupted messages to data nodes sometimes went undetected, causing a bad signal to be delivered to a block which aborted the data node. This failure in combination with disconnecting nodes could in turn cause the entire cluster to shut down.

To keep this from happening, additional checks are now made when unpacking signals received over TCP, including checks for byte order, compression flag (which must not be used), and the length of the next message in the receive buffer (if there is one).

Whenever two consecutive unpacked messages fail the checks just described, the current message is assumed to be corrupted. In this case, the transporter is marked as having bad data and no more unpacking of messages occurs until the transporter is reconnected. In addition, an entry is written to the cluster log containing the error as well as a hex dump of the corrupted message. (Bug #73843, Bug #19582925)

- `ndb_restore --print_data` truncated `TEXT` and `BLOB` column values to 240 bytes rather than 256 bytes. (Bug #65467, Bug #14571512)
- Transporter send buffers were not updated properly following a failed send. (Bug #45043, Bug #20113145)

## Changes in MySQL NDB Cluster 7.3.7 (5.6.21-ndb-7.3.7) (2014-10-17, General Availability)

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

### Functionality Added or Changed

- After adding new data nodes to the configuration file of a MySQL NDB Cluster having many API nodes, but prior to starting any of the data node processes, API nodes tried to connect to these



“missing” data nodes several times per second, placing extra loads on management nodes and the network. To reduce unnecessary traffic caused in this way, it is now possible to control the amount of time that an API node waits between attempts to connect to data nodes which fail to respond; this is implemented in two new API node configuration parameters `StartConnectBackoffMaxTime` and `ConnectBackoffMaxTime`.

Time elapsed during node connection attempts is not taken into account when applying these parameters, both of which are given in milliseconds with approximately 100 ms resolution. As long as the API node is not connected to any data nodes as described previously, the value of the `StartConnectBackoffMaxTime` parameter is applied; otherwise, `ConnectBackoffMaxTime` is used.

In a MySQL NDB Cluster with many unstarted data nodes, the values of these parameters can be raised to circumvent connection attempts to data nodes which have not yet begun to function in the cluster, as well as moderate high traffic to management nodes.

For more information about the behavior of these parameters, see [Defining SQL and Other API Nodes in an NDB Cluster](#). (Bug #17257842)

- Added the `--exclude-missing-tables` option for `ndb_restore`. When enabled, the option causes tables present in the backup but not in the target database to be ignored. (Bug #57566, Bug #11764704)

## Bugs Fixed

- **NDB Cluster APIs:** The fix for Bug #16723708 stopped the `ndb_logevent_get_next()` function from casting a log event's `ndb_mgm_event_category` to an `enum` type, but this change interfered with existing applications, and so the function's original behavior is now reinstated. A new MGM API function exhibiting the corrected behavior `ndb_logevent_get_next2()` has been added in this release to take the place of the reverted function, for use in applications that do not require backward compatibility. In all other respects apart from this, the new function is identical with its predecessor. (Bug #18354165)

References: Reverted patches: Bug #16723708.

- **NDB Cluster APIs:** NDB API scans leaked `Ndb_cluster_connection` objects after `nextResult()` was called when an operation resulted in an error. This leak locked up the corresponding connection objects in the `DBTC` kernel block until the connection was closed. (Bug #17730825, Bug #20170731)
- When assembling error messages of the form `Incorrect state for node n state: node_state`, written when the transporter failed to connect, the node state was used in place of the node ID in a number of instances, which resulted in errors of this type for which the node state was reported incorrectly. (Bug #19559313, Bug #73801)
- In some cases, transporter receive buffers were reset by one thread while being read by another. This happened when a race condition occurred between a thread receiving data and another thread initiating disconnect of the transporter (disconnection clears this buffer). Concurrency logic has now been implemented to keep this race from taking place. (Bug #19552283, Bug #73790)
- The failure of a data node could in some situations cause a set of API nodes to fail as well due to the sending of a `CLOSE_COMREQ` signal that was sometimes not completely initialized. (Bug #19513967)
- A more detailed error report is printed in the event of a critical failure in one of the `NDB` internal `sendSignal*()` methods, prior to crashing the process, as was already implemented for `sendSignal()`, but was missing from the more specialized `sendSignalNoRelease()` method. Having a crash of this type correctly reported can help with identifying configuration hardware issues in some cases. (Bug #19414511)

References: See also: Bug #19390895.

- `ndb_restore` failed to restore the cluster's metadata when there were more than approximately 17 K data objects. (Bug #19202654)
- The fix for a previous issue with the handling of multiple node failures required determining the number of TC instances the failed node was running, then taking them over. The mechanism to determine this number sometimes provided an invalid result which caused the number of TC instances in the failed node to be set to an excessively high value. This in turn caused redundant takeover attempts, which wasted time and had a negative impact on the processing of other node failures and of global checkpoints. (Bug #19193927)

References: This issue is a regression of: Bug #18069334.

- Parallel transactions performing reads immediately preceding a delete on the same tuple could cause the NDB kernel to crash. This was more likely to occur when separate TC threads were specified using the `ThreadConfig` configuration parameter. (Bug #19031389)
- Attribute promotion between different `TEXT` types (any of `TINYTEXT`, `TEXT`, `MEDIUMTEXT`, and `LONGTEXT`) by `ndb_restore` was not handled properly in some cases. In addition, `TEXT` values are now truncated according to the limits set by `mysqld` (for example, values converted to `TINYTEXT` from another type are truncated to 256 bytes). In the case of columns using a multibyte character set, the value is truncated to the end of the last well-formed character.

Also as a result of this fix, conversion to a `TEXT` column of any size that uses a different character set from the original is now disallowed. (Bug #18875137)

- To assist with diagnostic issues where many watchdog warnings are raised, it is now possible to activate (or deactivate) a killer watchdog using `DUMP 2610` in the `ndb_mgm` client. When set, this shuts down the data node on which the next watchdog warning occurs, providing a trace log. See `DUMP 2610`, for more information. (Bug #18703922)
- The NDB optimized node recovery mechanism attempts to transfer only relevant page changes to a starting node in order to speed the recovery process; this is done by having the starting node indicate the index of the last global checkpoint (GCI) in which it participated, so that the node that was already running copies only data for rows which have changed since that GCI. Every row has a GCI metacolumn which facilitates this; for a deleted row, the slot formerly storing this row's data contains a GCI value, and for deleted pages, every row on the missing page is considered changed and thus needs to be sent.

When these changes are received by the starting node, this node performs a lookup for the page and index to determine what they contain. This lookup could cause a real underlying page to be mapped against the logical page ID, even when this page contained no data.

One way in which this issue could manifest itself occurred after cluster `DataMemory` usage approached maximum, and deletion of many rows followed by a rolling restart of the data nodes was performed with the expectation that this would free memory, but in fact it was possible in this scenario for memory not to be freed and in some cases for memory usage actually to increase to its maximum.

This fix solves these issues by ensuring that a real physical page is mapped to a logical ID during node recovery only when this page contains actual data which needs to be stored. (Bug #18683398, Bug #18731008)

- When a data node sent a `MISSING_DATA` signal due to a buffer overflow and no event data had yet been sent for the current epoch, the dummy event list created to handle this inconsistency was not deleted after the information in the dummy event list was transferred to the completed list. (Bug #18410939)
- Incorrect calculation of the next autoincrement value following a manual insertion towards the end of a cached range could result in duplicate values sometimes being used. This issue could

manifest itself when using certain combinations of values for `auto_increment_increment`, `auto_increment_offset`, and `ndb_autoincrement_prefetch_sz`.

This issue has been fixed by modifying the calculation to make sure that the next value from the cache as computed by NDB is of the form `auto_increment_offset + (N * auto_increment_increment)`. This avoids any rounding up by the MySQL Server of the returned value, which could result in duplicate entries when the rounded-up value fell outside the range of values cached by NDB. (Bug #17893872)

- `ndb_show_tables --help` output contained misleading information about the `--database (-d)` option. In addition, the long form of the option (`--database`) did not work properly. (Bug #17703874)
- Using the `--help` option with `ndb_print_file` caused the program to segfault. (Bug #17069285)
- For multithreaded data nodes, some threads do communicate often, with the result that very old signals can remain at the top of the signal buffers. When performing a thread trace, the signal dumper calculated the latest signal ID from what it found in the signal buffers, which meant that these old signals could be erroneously counted as the newest ones. Now the signal ID counter is kept as part of the thread state, and it is this value that is used when dumping signals for trace files. (Bug #73842, Bug #19582807)

## Changes in MySQL NDB Cluster 7.3.6 (5.6.19-ndb-7.3.6) (2014-07-11, General Availability)

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

### Functionality Added or Changed

- **NDB Cluster APIs:** Added as an aid to debugging the ability to specify a human-readable name for a given `Ndb` object and later to retrieve it. These operations are implemented, respectively, as the `setNdbObjectName()` and `getNdbObjectName()` methods.

To make tracing of event handling between a user application and NDB easier, you can use the reference (from `getReference()` followed by the name (if provided) in printouts; the reference ties together the application `Ndb` object, the event buffer, and the NDB storage engine's `SUMA` block. (Bug #18419907)

### Bugs Fixed

- **NDB Disk Data:** Setting the undo buffer size used by `InitialLogFileGroup` to a value greater than that set by `SharedGlobalMemory` prevented data nodes from starting; the data nodes failed with Error 1504 `Out of logbuffer memory`. While the failure itself is expected behavior, the error message did not provide sufficient information to diagnose the actual source of the problem; now in such cases, a more specific error message `Out of logbuffer memory (specify smaller undo_buffer_size or increase SharedGlobalMemory)` is supplied. (Bug #11762867, Bug #55515)
- **NDB Cluster APIs:** When two tables had different foreign keys with the same name, `ndb_restore` considered this a name conflict and failed to restore the schema. As a result of this fix, a slash character (`/`) is now expressly disallowed in foreign key names, and the naming format `parent_id/child_id/fk_name` is now enforced by the NDB API. (Bug #18824753)
- **NDB Cluster APIs:** When an NDB data node indicates a buffer overflow via an empty epoch, the event buffer places an inconsistent data event in the event queue. When this was consumed, it was not removed from the event queue as expected, causing subsequent `nextEvent()` calls to return 0. This

caused event consumption to stall because the inconsistency remained flagged forever, while event data accumulated in the queue.

Event data belonging to an empty inconsistent epoch can be found either at the beginning or somewhere in the middle. `pollEvents()` returns 0 for the first case. This fix handles the second case: calling `nextEvent()` call dequeues the inconsistent event before it returns. In order to benefit from this fix, user applications must call `nextEvent()` even when `pollEvents()` returns 0. (Bug #18716991)

- **NDB Cluster APIs:** The `pollEvents()` method returned 1, even when called with a wait time equal to 0, and there were no events waiting in the queue. Now in such cases it returns 0 as expected. (Bug #18703871)
- Processing a `NODE_FAILREP` signal that contained an invalid node ID could cause a data node to fail. (Bug #18993037, Bug #73015)

References: This issue is a regression of: Bug #16007980.

- When building out of source, some files were written to the source directory instead of the build dir. These included the `manifest.mf` files used for creating ClusterJ jars and the `pom.xml` file used by `mvn_install_ndbjtie.sh`. In addition, `ndbinfo.sql` was written to the build directory, but marked as output to the source directory in `CMakeLists.txt`. (Bug #18889568, Bug #72843)
- When the binary log injector thread commits an epoch to the binary log and this causes the log file to reach maximum size, it may need to rotate the binary log. The rotation is not performed until either all the committed transactions from all client threads are flushed to the binary log, or a maximum of 30 seconds has elapsed. In the case where all transactions were committed prior to the 30-second wait, it was possible for committed transactions from multiple client threads to belong to newer epochs than the latest epoch committed by the injector thread, causing the thread to deadlock with itself, and causing an unnecessary 30-second delay before breaking the deadlock. (Bug #18845822)
- Adding a foreign key failed with NDB Error 208 if the parent index was parent table's primary key, the primary key was not on the table's initial attributes, and the child table was not empty. (Bug #18825966)
- When an NDB table served as both the parent table and a child table for 2 different foreign keys having the same name, dropping the foreign key on the child table could cause the foreign key on the parent table to be dropped instead, leading to a situation in which it was impossible to drop the remaining foreign key. This situation can be modelled using the following `CREATE TABLE` statements:

```
CREATE TABLE parent (  
  id INT NOT NULL,  
  PRIMARY KEY (id)  
) ENGINE=NDB;  
  
CREATE TABLE child (  
  id INT NOT NULL,  
  parent_id INT,  
  PRIMARY KEY (id),  
  INDEX par_ind (parent_id),  
  
  FOREIGN KEY (parent_id)  
  REFERENCES parent(id)  
) ENGINE=NDB;  
  
CREATE TABLE grandchild (  
  id INT,  
  parent_id INT,  
  INDEX par_ind (parent_id),  
  
  FOREIGN KEY (parent_id)  
  REFERENCES child(id)
```

```
) ENGINE=NDB;
```

With the tables created as just shown, the issue occurred when executing the statement `ALTER TABLE child DROP FOREIGN KEY parent_id`, because it was possible in some cases for NDB to drop the foreign key from the `grandchild` table instead. When this happened, any subsequent attempt to drop the foreign key from either the `child` or from the `grandchild` table failed. (Bug #18662582)

- It was possible for a data node restart to become stuck indefinitely in start phase 101 (see [Summary of NDB Cluster Start Phases](#)) when there were connection problems between the node being restarted and one or more subscribing API nodes.

To help prevent this from happening, a new data node configuration parameter `RestartSubscriberConnectTimeout` has been introduced, which can be used to control how long a data node restart can stall in start phase 101 before giving up and attempting to restart again. The default is 12000 ms. (Bug #18599198)

- Executing `ALTER TABLE ... REORGANIZE PARTITION` after increasing the number of data nodes in the cluster from 4 to 16 led to a crash of the data nodes. This issue was shown to be a regression caused by previous fix which added a new dump handler using a dump code that was already in use (7019), which caused the command to execute two different handlers with different semantics. The new handler was assigned a new `DUMP` code (7024). (Bug #18550318)

References: This issue is a regression of: Bug #14220269.

- Following a long series of inserts, when running with a relatively small redo log and an insufficient large value for `MaxNoOfConcurrentTransactions`, there remained transactions that were blocked by the lack of redo log and were thus not aborted in the correct state (waiting for prepare log to be sent to disk, or `LOG_QUEUED` state). This caused the redo log to remain blocked until unblocked by a completion of a local checkpoint. This could lead to a deadlock, when the blocked aborts in turned blocked global checkpoints, and blocked GCPs block LCPs. To prevent this situation from arising, we now abort immediately when we reach the `LOG_QUEUED` state in the abort state handler. (Bug #18533982)
- `ndbmt` supports multiple parallel receiver threads, each of which performs signal reception for a subset of the remote node connections (transporters) with the mapping of `remote_nodes` to receiver threads decided at node startup. Connection control is managed by the multi-instance `TRPMAN` block, which is organized as a proxy and workers, and each receiver thread has a `TRPMAN` worker running locally.

The `QMGR` block sends signals to `TRPMAN` to enable and disable communications with remote nodes. These signals are sent to the `TRPMAN` proxy, which forwards them to the workers. The workers themselves decide whether to act on signals, based on the set of remote nodes they manage.

The current issue arises because the mechanism used by the `TRPMAN` workers for determining which connections they are responsible for was implemented in such a way that each worker thought it was responsible for all connections. This resulted in the `TRPMAN` actions for `OPEN_COMORD`, `ENABLE_COMREQ`, and `CLOSE_COMREQ` being processed multiple times.

The fix keeps `TRPMAN` instances (receiver threads) executing `OPEN_COMORD`, `ENABLE_COMREQ` and `CLOSE_COMREQ` requests. In addition, the correct `TRPMAN` instance is now chosen when routing from this instance for a specific remote connection. (Bug #18518037)

- During data node failure handling, the transaction coordinator performing takeover gathers all known state information for any failed TC instance transactions, determines whether each transaction has been committed or aborted, and informs any involved API nodes so that they can report this accurately to their clients. The TC instance provides this information by sending `TCKEY_FAILREF` or `TCKEY_FAILCONF` signals to the API nodes as appropriate top each affected transaction.

In the event that this TC instance does not have a direct connection to the API node, it attempts to deliver the signal by routing it through another data node in the same node group as the failing TC, and sends a `GSN_TCKEY_FAILREFCONF_R` signal to TC block instance 0 in that data node. A problem arose in the case of multiple transaction coordinators, when this TC instance did not have a signal handler for such signals, which led it to fail.

This issue has been corrected by adding a handler to the TC proxy block which in such cases forwards the signal to one of the local TC worker instances, which in turn attempts to forward the signal on to the API node. (Bug #18455971)

- When running with a very slow main thread, and one or more transaction coordinator threads, on different CPUs, it was possible to encounter a timeout when sending a `DIH_SCAN_GET_NODESREQ` signal, which could lead to a crash of the data node. Now in such cases the timeout is avoided. (Bug #18449222)
- Failure of multiple nodes while using `ndbmt` with multiple TC threads was not handled gracefully under a moderate amount of traffic, which could in some cases lead to an unplanned shutdown of the cluster. (Bug #18069334)
- A local checkpoint (LCP) is tracked using a global LCP state (`c_lcpState`), and each NDB table has a status indicator which indicates the LCP status of that table (`tabLcpStatus`). If the global LCP state is `LCP_STATUS_IDLE`, then all the tables should have an LCP status of `TLS_COMPLETED`.

When an LCP starts, the global LCP status is `LCP_INIT_TABLES` and the thread starts setting all the NDB tables to `TLS_ACTIVE`. If any tables are not ready for LCP, the LCP initialization procedure continues with `CONTINUEB` signals until all tables have become available and been marked `TLS_ACTIVE`. When this initialization is complete, the global LCP status is set to `LCP_STATUS_ACTIVE`.

This bug occurred when the following conditions were met:

- An LCP was in the `LCP_INIT_TABLES` state, and some but not all tables had been set to `TLS_ACTIVE`.
- The master node failed before the global LCP state changed to `LCP_STATUS_ACTIVE`; that is, before the LCP could finish processing all tables.
- The `NODE_FAILREP` signal resulting from the node failure was processed before the final `CONTINUEB` signal from the LCP initialization process, so that the node failure was processed while the LCP remained in the `LCP_INIT_TABLES` state.

Following master node failure and selection of a new one, the new master queries the remaining nodes with a `MASTER_LCPREQ` signal to determine the state of the LCP. At this point, since the LCP status was `LCP_INIT_TABLES`, the LCP status was reset to `LCP_STATUS_IDLE`. However, the LCP status of the tables was not modified, so there remained tables with `TLS_ACTIVE`. Afterwards, the failed node is removed from the LCP. If the LCP status of a given table is `TLS_ACTIVE`, there is a check that the global LCP status is not `LCP_STATUS_IDLE`; this check failed and caused the data node to fail.

Now the `MASTER_LCPREQ` handler ensures that the `tabLcpStatus` for all tables is updated to `TLS_COMPLETED` when the global LCP status is changed to `LCP_STATUS_IDLE`. (Bug #18044717)

- When performing a copying `ALTER TABLE` operation, `mysqld` creates a new copy of the table to be altered. This intermediate table, which is given a name bearing the prefix `#sql-`, has an updated schema but contains no data. `mysqld` then copies the data from the original table to this intermediate table, drops the original table, and finally renames the intermediate table with the name of the original table.

`mysqld` regards such a table as a temporary table and does not include it in the output from `SHOW TABLES`; `mysqldump` also ignores an intermediate table. However, `NDB` sees no difference between such an intermediate table and any other table. This difference in how intermediate tables are viewed by `mysqld` (and MySQL client programs) and by the `NDB` storage engine can give rise to problems when performing a backup and restore if an intermediate table existed in `NDB`, possibly left over from a failed `ALTER TABLE` that used copying. If a schema backup is performed using `mysqldump` and the `mysql` client, this table is not included. However, in the case where a data backup was done using the `ndb_mgm` client's `BACKUP` command, the intermediate table was included, and was also included by `ndb_restore`, which then failed due to attempting to load data into a table which was not defined in the backed up schema.

To prevent such failures from occurring, `ndb_restore` now by default ignores intermediate tables created during `ALTER TABLE` operations (that is, tables whose names begin with the prefix `#sql-`). A new option `--exclude-intermediate-sql-tables` is added that makes it possible to override the new behavior. The option's default value is `TRUE`; to cause `ndb_restore` to revert to the old behavior and to attempt to restore intermediate tables, set this option to `FALSE`. (Bug #17882305)

- The logging of insert failures has been improved. This is intended to help diagnose occasional issues seen when writing to the `mysql.ndb_binlog_index` table. (Bug #17461625)
- The `DEFINER` column in the `INFORMATION_SCHEMA.VIEWS` table contained erroneous values for views contained in the `ndbinfo` information database. This could be seen in the result of a query such as `SELECT TABLE_NAME, DEFINER FROM INFORMATION_SCHEMA.VIEWS WHERE TABLE_SCHEMA='ndbinfo'`. (Bug #17018500)
- Employing a `CHAR` column that used the `UTF8` character set as a table's primary key column led to node failure when restarting data nodes. Attempting to restore a table with such a primary key also caused `ndb_restore` to fail. (Bug #16895311, Bug #68893)
- The `--order (-o)` option for the `ndb_select_all` utility worked only when specified as the last option, and did not work with an equals sign.

As part of this fix, the program's `--help` output was also aligned with the `--order` option's correct behavior. (Bug #64426, Bug #16374870)

## Changes in MySQL NDB Cluster 7.3.5 (5.6.17-ndb-7.3.5) (2014-04-07, General Availability)

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

### Functionality Added or Changed

- Handling of `LongMessageBuffer` shortages and statistics has been improved as follows:
  - The default value of `LongMessageBuffer` has been increased from 4 MB to 64 MB.
  - When this resource is exhausted, a suitable informative message is now printed in the data node log describing possible causes of the problem and suggesting possible solutions.
  - `LongMessageBuffer` usage information is now shown in the `ndbinfo.memoryusage` table. See the description of this table for an example and additional information.

## Bugs Fixed

- **Important Change:** The server system variables `ndb_index_cache_entries` and `ndb_index_stat_freq`, which had been deprecated in a previous MySQL NDB Cluster release series, have now been removed. (Bug #11746486, Bug #26673)
- **Solaris:** An issue found when compiling the MySQL NDB Cluster software for Solaris platforms could lead to problems when using `ThreadConfig` on such systems. (Bug #18181656)
- **NDB Cluster APIs:** When an NDB API client application received a signal with an invalid block or signal number, `NDB` provided only a very brief error message that did not accurately convey the nature of the problem. Now in such cases, appropriate printouts are provided when a bad signal or message is detected. In addition, the message length is now checked to make certain that it matches the size of the embedded signal. (Bug #18426180)
- **NDB Cluster APIs:** Refactoring that was performed in MySQL NDB Cluster 7.3.4 inadvertently introduced a dependency in `Ndb.hpp` on a file that is not included in the distribution, which caused NDB API applications to fail to compile. The dependency has been removed. (Bug #18293112, Bug #71803)

References: This issue is a regression of: Bug #17647637.

- **NDB Cluster APIs:** An NDB API application sends a scan query to a data node; the scan is processed by the transaction coordinator (TC). The TC forwards a `LQHKEYREQ` request to the appropriate LDM, and aborts the transaction if it does not receive a `LQHKEYCONF` response within the specified time limit. After the transaction is successfully aborted, the TC sends a `TCROLLBACKREP` to the NDBAPI client, and the NDB API client processes this message by cleaning up any `Ndb` objects associated with the transaction.

The client receives the data which it has requested in the form of `TRANSID_AI` signals, buffered for sending at the data node, and may be delivered after a delay. On receiving such a signal, `NDB` checks the transaction state and ID: if these are as expected, it processes the signal using the `Ndb` objects associated with that transaction.

The current bug occurs when all the following conditions are fulfilled:

- The transaction coordinator aborts a transaction due to delays and sends a `TCROLLBACPREP` signal to the client, while at the same time a `TRANSID_AI` which has been buffered for delivery at an LDM is delivered to the same client.
- The NDB API client considers the transaction complete on receipt of a `TCROLLBACKREP` signal, and immediately closes the transaction.
- The client has a separate receiver thread running concurrently with the thread that is engaged in closing the transaction.
- The arrival of the late `TRANSID_AI` interleaves with the closing of the user thread's transaction such that `TRANSID_AI` processing passes normal checks before `closeTransaction()` resets the transaction state and invalidates the receiver.

When these conditions are all met, the receiver thread proceeds to continue working on the `TRANSID_AI` signal using the invalidated receiver. Since the receiver is already invalidated, its usage results in a node failure.

Now the `Ndb` object cleanup done for `TCROLLBACKREP` includes invalidation of the transaction ID, so that, for a given transaction, any signal which is received after the `TCROLLBACKREP` arrives does not pass the transaction ID check and is silently dropped. This fix is also implemented for the `TC_COMMITREF`, `TCROLLBACKREF`, `TCKEY_FAILCONF`, and `TCKEY_FAILREF` signals as well.



See also [Operations and Signals](#), for additional information about NDB messaging. (Bug #18196562)

- **NDB Cluster APIs:** The example `ndbapi-examples/ndbapi_blob_ndbrecord/main.cpp` included an internal header file (`ndb_global.h`) not found in the MySQL NDB Cluster binary distribution. The example now uses `stdlib.h` and `string.h` instead of this file. (Bug #18096866, Bug #71409)
- **NDB Cluster APIs:** When `Dictionary::dropTable()` attempted (as a normal part of its internal operations) to drop an index used by a foreign key constraint, the drop failed. Now in such cases, invoking `dropTable()` causes all foreign keys on the table to be dropped, whether this table acts as a parent table, child table, or both.

This issue did not affect dropping of indexes using SQL statements. (Bug #18069680)

References: See also: Bug #17591531.

- **NDB Cluster APIs:** `ndb_restore` could sometimes report `Error 701 System busy with other schema operation` unnecessarily when restoring in parallel. (Bug #17916243)
- When an `ALTER TABLE` statement changed table schemas without causing a change in the table's partitioning, the new table definition did not copy the hash map from the old definition, but used the current default hash map instead. However, the table data was not reorganized according to the new hashmap, which made some rows inaccessible using a primary key lookup if the two hash maps had incompatible definitions.

To keep this situation from occurring, any `ALTER TABLE` that entails a hashmap change now triggers a reorganisation of the table. In addition, when copying a table definition in such cases, the hashmap is now also copied. (Bug #18436558)

- When certain queries generated signals having more than 18 data words prior to a node failure, such signals were not written correctly in the trace file. (Bug #18419554)
- Checking of timeouts is handled by the signal `TIME_SIGNAL`. Previously, this signal was generated by the `QMGR` NDB kernel block in the main thread, and sent to the `QMRG`, `DBLQH`, and `DBTC` blocks (see [NDB Kernel Blocks](#)) as needed to check (respectively) heartbeats, disk writes, and transaction timeouts. In `ndbmtd` (as opposed to `ndbd`), these blocks all execute in different threads. This meant that if, for example, `QMGR` was actively working and some other thread was put to sleep, the previously sleeping thread received a large number of `TIME_SIGNAL` messages simultaneously when it was woken up again, with the effect that effective times moved very quickly in `DBLQH` as well as in `DBTC`. In `DBLQH`, this had no noticeable adverse effects, but this was not the case in `DBTC`; the latter block could not work on transactions even though time was still advancing, leading to a situation in which many operations appeared to time out because the transaction coordinator (TC) thread was comparatively slow in answering requests.

In addition, when the TC thread slept for longer than 1500 milliseconds, the data node crashed due to detecting that the timeout handling loop had not yet stopped. To rectify this problem, the generation of the `TIME_SIGNAL` has been moved into the local threads instead of `QMGR`; this provides for better control over how quickly `TIME_SIGNAL` messages are allowed to arrive. (Bug #18417623)

- The `ServerPort` and `TcpBind_INADDR_ANY` configuration parameters were not included in the output of `ndb_mgmd --print-full-config`. (Bug #18366909)
- After dropping an NDB table, neither the cluster log nor the output of the `REPORT MemoryUsage` command showed that the `IndexMemory` used by that table had been freed, even though the memory had in fact been deallocated. This issue was introduced in MySQL NDB Cluster 7.3.2. (Bug #18296810)

- `ndb_show_tables` sometimes failed with the error message `Unable to connect to management server` and immediately terminated, without providing the underlying reason for the failure. To provide more useful information in such cases, this program now also prints the most recent error from the `Ndb_cluster_connection` object used to instantiate the connection. (Bug #18276327)
- `-DWITH_NDBMTD=0` did not function correctly, which could cause the build to fail on platforms such as ARM and Raspberry Pi which do not define the memory barrier functions required to compile `ndbmt.d`. (Bug #18267919)

References: See also: Bug #16620938.

- The block threads managed by the multithreading scheduler communicate by placing signals in an out queue or job buffer which is set up between all block threads. This queue has a fixed maximum size, such that when it is filled up, the worker thread must wait for the consumer to drain the queue. In a highly loaded system, multiple threads could end up in a circular wait lock due to full out buffers, such that they were preventing each other from performing any useful work. This condition eventually led to the data node being declared dead and killed by the watchdog timer.

To fix this problem, we detect situations in which a circular wait lock is about to begin, and cause buffers which are otherwise held in reserve to become available for signal processing by queues which are highly loaded. (Bug #18229003)

- The `ndb_mgm` client `START BACKUP` command (see [Commands in the NDB Cluster Management Client](#)) could experience occasional random failures when a ping was received prior to an expected `BackupCompleted` event. Now the connection established by this command is not checked until it has been properly set up. (Bug #18165088)
- When creating a table with foreign key referencing an index in another table, it sometimes appeared possible to create the foreign key even if the order of the columns in the indexes did not match, due to the fact that an appropriate error was not always returned internally. This fix improves the error used internally to work in most cases; however, it is still possible for this situation to occur in the event that the parent index is a unique index. (Bug #18094360)
- Updating parent tables of foreign keys used excessive scan resources and so required unusually high settings for `MaxNoOfLocalScans` and `MaxNoOfConcurrentScans`. (Bug #18082045)
- Dropping a nonexistent foreign key on an `NDB` table (using, for example, `ALTER TABLE`) appeared to succeed. Now in such cases, the statement fails with a relevant error message, as expected. (Bug #17232212)
- Data nodes running `ndbmt.d` could stall while performing an online upgrade of a MySQL NDB Cluster containing a great many tables from a version prior to NDB 7.2.5 to version 7.2.5 or later. (Bug #16693068)

## Changes in MySQL NDB Cluster 7.3.4 (5.6.15-ndb-7.3.4) (2014-02-06, General Availability)

### Bugs Fixed

- **Packaging; Solaris:** Compilation of `ndbmt.d` failed on Solaris 10 and 11 for 32-bit `x86`, and the binary was not included in the binary distributions for these platforms. (Bug #16620938)
- **Microsoft Windows:** Timers used in timing scheduler events in the `NDB` kernel have been refactored, in part to insure that they are monotonic on all platforms. In particular, on Windows, event intervals were previously calculated using values obtained from `GetSystemTimeAsFileTime()`, which

reads directly from the system time (“wall clock”), and which may arbitrarily be reset backward or forward, leading to false watchdog or heartbeat alarms, or even node shutdown. Lack of timer monotonicity could also cause slow disk writes during backups and global checkpoints. To fix this issue, the Windows implementation now uses `QueryPerformanceCounters()` instead of `GetSystemTimeAsFileTime()`. In the event that a monotonic timer is not found on startup of the data nodes, a warning is logged.

In addition, on all platforms, a check is now performed at compile time for available system monotonic timers, and the build fails if one cannot be found; note that `CLOCK_HIGHRES` is now supported as an alternative for `CLOCK_MONOTONIC` if the latter is not available. (Bug #17647637)

- **NDB Disk Data:** When using Disk Data tables and `ndbmt` data nodes, it was possible for the undo buffer to become overloaded, leading to a crash of the data nodes. This issue was more likely to be encountered when using Disk Data columns whose size was approximately 8K or larger. (Bug #16766493)
- **NDB Cluster APIs:** `UINT_MAX64` was treated as a signed value by Visual Studio 2010. To prevent this from happening, the value is now explicitly defined as unsigned. (Bug #17947674)

References: See also: Bug #17647637.

- **NDB Cluster APIs:** It was possible for an `Ndb` object to receive signals for handling before it was initialized, leading to thread interleaving and possible data node failure when executing a call to `Ndb::init()`. To guard against this happening, a check is now made when it is starting to receive signals that the `Ndb` object is properly initialized before any signals are actually handled. (Bug #17719439)
- **NDB Cluster APIs:** Compilation of example NDB API program files failed due to missing include directives. (Bug #17672846, Bug #70759)
- **NDB Cluster APIs:** An application, having opened two distinct instances of `Ndb_cluster_connection`, attempted to use the second connection object to send signals to itself, but these signals were blocked until the destructor was explicitly called for that connection object. (Bug #17626525)

References: This issue is a regression of: Bug #16595838.

- Interrupting a drop of a foreign key could cause the underlying table to become corrupt. (Bug #18041636)
- Monotonic timers on several platforms can experience issues which might result in the monotonic clock doing small jumps back in time. This is due to imperfect synchronization of clocks between multiple CPU cores and does not normally have an adverse effect on the scheduler and watchdog mechanisms; so we handle some of these cases by making backtick protection less strict, although we continue to ensure that the backtick is less than 10 milliseconds. This fix also removes several checks for backticks which are thereby made redundant. (Bug #17973819)
- Under certain specific circumstances, in a cluster having two SQL nodes, one of these could hang, and could not be accessed again even after killing the `mysqld` process and restarting it. (Bug #17875885, Bug #18080104)

References: See also: Bug #17934985.

- Poor support or lack of support on some platforms for monotonic timers caused issues with delayed signal handling by the job scheduler for the multithreaded data node. Variances (timer leaps) on such platforms are now handled in the same way the multithreaded data node process that they are by the singlethreaded version. (Bug #17857442)

References: See also: Bug #17475425, Bug #17647637.

- In some cases, with `ndb_join_pushdown` enabled, it was possible to obtain from a valid query the error `Got error 290 'Corrupt key in TC, unable to xfrm' from NDBCLUSTER` even though the data was not actually corrupted.

It was determined that a `NULL` in a `VARCHAR` column could be used to construct a lookup key, but since `NULL` is never equal to any other value, such a lookup could simple have been eliminated instead. This `NULL` lookup in turn led to the spurious error message.

This fix takes advantage of the fact that a key lookup with `NULL` never finds any matching rows, and so `NDB` does not try to perform the lookup that would have led to the error. (Bug #17845161)

- It was theoretically possible in certain cases for a number of output functions internal to the `NDB` code to supply an uninitialized buffer as output. Now in such cases, a newline character is printed instead. (Bug #17775602, Bug #17775772)
- Use of the `localtime()` function in `NDB` multithreading code led to otherwise nondeterministic failures in `ndbmttd`. This fix replaces this function, which on many platforms uses a buffer shared among multiple threads, with `localtime_r()`, which can have allocated to it a buffer of its own. (Bug #17750252)
- When using single-threaded (`ndbd`) data nodes with `RealTimeScheduler` enabled, the CPU did not, as intended, temporarily lower its scheduling priority to normal every 10 milliseconds to give other, non-realtime threads a chance to run. (Bug #17739131)
- During arbitrator selection, `QMGR` (see [The QMGR Block](#)) runs through a series of states, the first few of which are (in order) `NULL`, `INIT`, `FIND`, `PREP1`, `PREP2`, and `START`. A check for an arbitration selection timeout occurred in the `FIND` state, even though the corresponding timer was not set until `QMGR` reached the `PREP1` and `PREP2` states. Attempting to read the resulting uninitialized timestamp value could lead to false `Could not find an arbitrator, cluster is not partition-safe` warnings.

This fix moves the setting of the timer for arbitration timeout to the `INIT` state, so that the value later read during `FIND` is always initialized. (Bug #17738720)

- The global checkpoint lag watchdog tracking the number of times a check for GCP lag was performed using the system scheduler and used this count to check for a timeout condition, but this caused a number of issues. To overcome these limitations, the GCP watchdog has been refactored to keep track of its own start times, and to calculate elapsed time by reading the (real) clock every time it is called.

In addition, any backticks (rare in any case) are now handled by taking the backward time as the new current time and calculating the elapsed time for this round as 0. Finally, any ill effects of a forward leap, which possibly could expire the watchdog timer immediately, are reduced by never calculating an elapsed time longer than the requested delay time for the watchdog timer. (Bug #17647469)

References: See also: Bug #17842035.

- The length of the interval (intended to be 10 seconds) between warnings for `GCP_COMMIT` when the GCP progress watchdog did not detect progress in a global checkpoint was not always calculated correctly. (Bug #17647213)
- Trying to drop an index used by a foreign key constraint caused data node failure. Now in such cases, the statement used to perform the drop fails. (Bug #17591531)
- In certain rare cases on commit of a transaction, an `Ndb` object was released before the transaction coordinator (`DBTC` kernel block) sent the expected `COMMIT_CONF` signal; `NDB` failed to send a `COMMIT_ACK` signal in response, which caused a memory leak in the `NDB` kernel could later lead to node failure.

Now an `Ndb` object is not released until the `COMMIT_CONF` signal has actually been received. (Bug #16944817)

- Losing its connections to the management node or data nodes while a query against the `ndbinfo.memoryusage` table was in progress caused the SQL node where the query was issued to fail. (Bug #14483440, Bug #16810415)
- The `ndbd_redo_log_reader` utility now supports a `--help` option. Using this options causes the program to print basic usage information, and then to exit. (Bug #11749591, Bug #36805)

## Changes in MySQL NDB Cluster 7.3.3 (5.6.14-ndb-7.3.3) (2013-11-18, General Availability)

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

### Functionality Added or Changed

- The length of time a management node waits for a heartbeat message from another management node is now configurable using the `HeartbeatIntervalMgmdMgmd` management node configuration parameter added in this release. The connection is considered dead after 3 missed heartbeats. The default value is 1500 milliseconds, or a timeout of approximately 6000 ms. (Bug #17807768, Bug #16426805)
- The MySQL NDB Cluster Auto-Installer now generates a `my.cnf` file for each `mysqld` in the cluster before starting it. For more information, see [Using the NDB Cluster Auto-Installer](#). (Bug #16994782)
- `BLOB` and `TEXT` columns are now reorganized by the `ALTER ONLINE TABLE ... REORGANIZE PARTITION` statement. (Bug #13714148)

### Bugs Fixed

- **Performance:** In a number of cases found in various locations in the MySQL NDB Cluster codebase, unnecessary iterations were performed; this was caused by failing to break out of a repeating control structure after a test condition had been met. This community-contributed fix removes the unneeded repetitions by supplying the missing breaks. (Bug #16904243, Bug #69392, Bug #16904338, Bug #69394, Bug #16778417, Bug #69171, Bug #16778494, Bug #69172, Bug #16798410, Bug #69207, Bug #16801489, Bug #69215, Bug #16904266, Bug #69393)
- **Packaging:** Portions of the documentation specific to MySQL NDB Cluster and the `NDB` storage engine were not included when installing from RPMs. (Bug #16303451)
- **Microsoft Windows:** The Windows error `ERROR_FILE_EXISTS` was not recognized by `NDB`, which treated it as an unknown error. (Bug #16970960)
- **NDB Disk Data:** `NDB` error 899 `RowId already allocated` was raised due to a `RowId` “leak” which occurred under either of the following sets of circumstances:
  1. Insertion of a row into an in-memory table was rejected after an ordered index update failed due to insufficient `DataMemory`.
  2. Insertion of a row into a Disk Data table was rejected due to lack of sufficient table space.(Bug #13927679)

References: See also: Bug #22494024, Bug #13990924.

- **NDB Cluster APIs:** The `Event::setTable()` method now supports a pointer or a reference to table as its required argument. If a null table pointer is used, the method now returns -1 to make it clear that this is what has occurred. (Bug #16329082)
- Trying to restore to a table having a `BLOB` column in a different position from that of the original one caused `ndb_restore --restore_data` to fail. (Bug #17395298)
- `ndb_restore` could abort during the last stages of a restore using attribute promotion or demotion into an existing table. This could happen if a converted attribute was nullable and the backup had been run on active database. (Bug #17275798)
- It was not possible to start MySQL NDB Cluster processes created by the Auto-Installer on a Windows host running freeSSHd. (Bug #17269626)
- The `DBUTIL` data node block is now less strict about the order in which it receives certain messages from other nodes. (Bug #17052422)
- `ALTER ONLINE TABLE ... REORGANIZE PARTITION` failed when run against a table having or using a reference to a foreign key. (Bug #17036744, Bug #69619)
- `TUPKEYREQ` signals are used to read data from the tuple manager block (`DBTUP`), and are used for all types of data access, especially for scans which read many rows. A `TUPKEYREQ` specifies a series of 'columns' to be read, which can be either single columns in a specific table, or pseudocolumns, two of which—`READ_ALL` and `READ_PACKED`—are aliases to read all columns in a table, or some subset of these columns. Pseudocolumns are used by modern NDB API applications as they require less space in the `TUPKEYREQ` to specify columns to be read, and can return the data in a more compact (packed) format.

This fix moves the creation and initialization of on-stack Signal objects to only those pseudocolumn reads which need to `EXECUTE_DIRECT` to other block instances, rather than for every read. In addition, the size of an on-stack signal is now varied to suit the requirements of each pseudocolumn, so that only reads of the `INDEX_STAT` pseudocolumn now require initialization (and 3KB memory each time this is performed). (Bug #17009502)

- A race condition could sometimes occur when trying to lock receive threads to cores. (Bug #17009393)
- Results from joins using a `WHERE` with an `ORDER BY ... DESC` clause were not sorted properly; the `DESC` keyword in such cases was effectively ignored. (Bug #16999886, Bug #69528)
- `RealTimeScheduler` did not work correctly with data nodes running `ndbmt.d`. (Bug #16961971)
- File system errors occurring during a local checkpoint could sometimes cause an LCP to hang with no obvious cause when they were not handled correctly. Now in such cases, such errors always cause the node to fail. Note that the LQH block always shuts down the node when a local checkpoint fails; the change here is to make likely node failure occur more quickly and to make the original file system error more visible. (Bug #16961443)
- Maintenance and checking of parent batch completion in the `SPJ` block of the `NDB` kernel was reimplemented. Among other improvements, the completion state of all ancestor nodes in the tree are now preserved. (Bug #16925513)
- Dropping a column, which was not itself a foreign key, from an `NDB` table having foreign keys failed with `ER_TABLE_DEF_CHANGED`. (Bug #16912989)
- The LCP fragment scan watchdog periodically checks for lack of progress in a fragment scan performed as part of a local checkpoint, and shuts down the node if there is no progress after a given amount of

time has elapsed. This interval, formerly hard-coded as 60 seconds, can now be configured using the `LcpScanProgressTimeout` data node configuration parameter added in this release.

This configuration parameter sets the maximum time the local checkpoint can be stalled before the LCP fragment scan watchdog shuts down the node. The default is 60 seconds, which provides backward compatibility with previous releases.

You can disable the LCP fragment scan watchdog by setting this parameter to 0. (Bug #16630410)

- Added the `ndb_error_reporter` options `--connection-timeout`, which makes it possible to set a timeout for connecting to nodes, `--dry-scp`, which disables scp connections to remote hosts, and `--skip-nodegroup`, which skips all nodes in a given node group. (Bug #16602002)

References: See also: Bug #11752792, Bug #44082.

- After issuing `START BACKUP id WAIT STARTED`, if `id` had already been used for a backup ID, an error caused by the duplicate ID occurred as expected, but following this, the `START BACKUP` command never completed. (Bug #16593604, Bug #68854)
- `ndb_mgm` treated backup IDs provided to `ABORT BACKUP` commands as signed values, so that backup IDs greater than  $2^{31}$  wrapped around to negative values. This issue also affected out-of-range backup IDs, which wrapped around to negative values instead of causing errors as expected in such cases. The backup ID is now treated as an unsigned value, and `ndb_mgm` now performs proper range checking for backup ID values greater than `MAX_BACKUPS` ( $2^{32}$ ). (Bug #16585497, Bug #68798)
- When trying to specify a backup ID greater than the maximum allowed, the value was silently truncated. (Bug #16585455, Bug #68796)
- The unexpected shutdown of another data node as a starting data node received its node ID caused the latter to hang in Start Phase 1. (Bug #16007980)

References: See also: Bug #18993037.

- `SELECT ... WHERE ... LIKE` from an NDB table could return incorrect results when using `engine_condition_pushdown=ON`. (Bug #15923467, Bug #67724)
- The NDB receive thread waited unnecessarily for additional job buffers to become available when receiving data. This caused the receive mutex to be held during this wait, which could result in a busy wait when the receive thread was running with real-time priority.

This fix also handles the case where a negative return value from the initial check of the job buffer by the receive thread prevented further execution of data reception, which could possibly lead to communication blockage or configured `ReceiveBufferMemory` underutilization. (Bug #15907515)

- When the available job buffers for a given thread fell below the critical threshold, the internal multithreading job scheduler waited for job buffers for incoming rather than outgoing signals to become available, which meant that the scheduler waited the maximum timeout (1 millisecond) before resuming execution. (Bug #15907122)
- Setting `lower_case_table_names` to 1 or 2 on Windows systems caused `ALTER TABLE ... ADD FOREIGN KEY` statements against tables with names containing uppercase letters to fail with Error 155, `No such table: '(null)'`. (Bug #14826778, Bug #67354)
- Under some circumstances, a race occurred where the wrong watchdog state could be reported. A new state name `Packing Send Buffers` is added for watchdog state number 11, previously reported as `Unknown place`. As part of this fix, the state numbers for states without names are always now reported in such cases. (Bug #14824490)

- When a node fails, the Distribution Handler ([DBDIH](#) kernel block) takes steps together with the Transaction Coordinator ([DBTC](#)) to make sure that all ongoing transactions involving the failed node are taken over by a surviving node and either committed or aborted. Transactions taken over which are then committed belong in the epoch that is current at the time the node failure occurs, so the surviving nodes must keep this epoch available until the transaction takeover is complete. This is needed to maintain ordering between epochs.

A problem was encountered in the mechanism intended to keep the current epoch open which led to a race condition between this mechanism and that normally used to declare the end of an epoch. This could cause the current epoch to be closed prematurely, leading to failure of one or more surviving data nodes. (Bug #14623333, Bug #16990394)

- Exhaustion of [LongMessageBuffer](#) memory under heavy load could cause data nodes running [ndbmysqld](#) to fail. (Bug #14488185)
- When using dynamic listening ports for accepting connections from API nodes, the port numbers were reported to the management server serially. This required a round trip for each API node, causing the time required for data nodes to connect to the management server to grow linearly with the number of API nodes. To correct this problem, each data node now reports all dynamic ports at once. (Bug #12593774)
- [ndb\\_error-reporter](#) did not support the `--help` option. (Bug #11756666, Bug #48606)

References: See also: Bug #11752792, Bug #44082.

- Formerly, the node used as the coordinator or leader for distributed decision making between nodes (also known as the [DICT](#) manager—see [The DBDICT Block](#)) was indicated in the output of the [ndb\\_mgm](#) client `SHOW` command as the “master” node, although this node has no relationship to a master server in MySQL Replication. (It should also be noted that it is not necessary to know which node is the leader except when debugging [NDBCLUSTER](#) source code.) To avoid possible confusion, this label has been removed, and the leader node is now indicated in `SHOW` command output using an asterisk (\*) character. (Bug #11746263, Bug #24880)
- The matrix of values used for thread configuration when applying the setting of the [MaxNoOfExecutionThreads](#) configuration parameter has been improved to align with support for greater numbers of LDM threads. See [Multi-Threading Configuration Parameters \(ndbmysqld\)](#), for more information about the changes. (Bug #75220, Bug #20215689)
- Program execution failed to break out of a loop after meeting a desired condition in a number of internal methods, performing unneeded work in all cases where this occurred. (Bug #69610, Bug #69611, Bug #69736, Bug #17030606, Bug #17030614, Bug #17160263)
- `ABORT BACKUP` in the [ndb\\_mgm](#) client (see [Commands in the NDB Cluster Management Client](#)) took an excessive amount of time to return (approximately as long as the backup would have required to complete, had it not been aborted), and failed to remove the files that had been generated by the aborted backup. (Bug #68853, Bug #17719439)
- Attribute promotion and demotion when restoring data to [NDB](#) tables using `ndb_restore --restore_data` with the `--promote-attributes` and `--lossy-conversions` options has been improved as follows:
  - Columns of types [CHAR](#), and [VARCHAR](#) can now be promoted to [BINARY](#) and [VARBINARY](#), and columns of the latter two types can be demoted to one of the first two.

Note that converted character data is not checked to conform to any character set.

- Any of the types [CHAR](#), [VARCHAR](#), [BINARY](#), and [VARBINARY](#) can now be promoted to [TEXT](#) or [BLOB](#).



When performing such promotions, the only other sort of type conversion that can be performed at the same time is between character types and binary types.

## Changes in MySQL NDB Cluster 7.3.2 (5.6.11-ndb-7.3.2) (2013-06-18, General Availability)

### Bugs Fixed

- **Packaging; Microsoft Windows:** The MySQL NDB Cluster installer for Windows provided a nonfunctional option to install debug symbols (contained in \*.pdb files). This option has been removed from the installer.



#### Note

You can obtain the \*.pdb debug files for a given MySQL NDB Cluster release from the Windows .zip archive for the same release, such as `mysql-cluster-gpl-7.2.14-win32.zip` or `mysql-cluster-gpl-7.3.2-winx64.zip`.

(Bug #16748308, Bug #69112)

- **NDB Disk Data:** The statements `CREATE TABLESPACE`, `ALTER LOGFILE GROUP`, and `ALTER TABLESPACE` failed with a syntax error when `INITIAL_SIZE` was specified using letter abbreviations such as `M` or `G`. In addition, `CREATE LOGFILE GROUP` failed when `INITIAL_SIZE`, `UNDO_BUFFER_SIZE`, or both options were specified using letter abbreviations. (Bug #13116514, Bug #16104705, Bug #62858)
- **NDB Cluster APIs:** For each log event retrieved using the MGM API, the log event category (`ndb_mgm_event_category`) was simply cast to an `enum` type, which resulted in invalid category values. Now an offset is added to the category following the cast to ensure that the value does not fall out of the allowed range.



#### Note

This change was reverted by the fix for Bug #18354165. See the MySQL NDB Cluster API Developer documentation for `ndb_logevent_get_next()`, for more information.

(Bug #16723708)

References: See also: Bug #18354165.

- **NDB Cluster APIs:** The `Ndb::computeHash()` API method performs a `malloc()` if no buffer is provided for it to use. However, it was assumed that the memory thus returned would always be suitably aligned, which is not always the case. Now when `malloc()` provides a buffer to this method, the buffer is aligned after it is allocated, and before it is used. (Bug #16484617)
- `mysql_upgrade` failed when upgrading from MySQL NDB Cluster 7.1.26 to MySQL NDB Cluster 7.2.13 when it attempted to invoke a stored procedure before the `mysql.proc` table had been upgraded. (Bug #16933405)

References: This issue is a regression of: Bug #16226274.

- The planned or unplanned shutdown of one or more data nodes while reading table data from the `ndbinfo` database caused a memory leak. (Bug #16932989)

- Executing `DROP TABLE` while `DBDIH` was updating table checkpoint information subsequent to a node failure could lead to a data node failure. (Bug #16904469)
- In certain cases, when starting a new SQL node, `mysqld` failed with Error 1427 `Api node died, when SUB_START_REQ reached node`. (Bug #16840741)
- Failure to use container classes specific `NDB` during node failure handling could cause leakage of commit-ack markers, which could later lead to resource shortages or additional node crashes. (Bug #16834416)
- Use of an uninitialized variable employed in connection with error handling in the `DBLQH` kernel block could sometimes lead to a data node crash or other stability issues for no apparent reason. (Bug #16834333)
- A race condition in the time between the reception of a `execNODE_FAILREP` signal by the `QMGR` kernel block and its reception by the `DBLQH` and `DBTC` kernel blocks could lead to data node crashes during shutdown. (Bug #16834242)
- The `CLUSTERLOG` command (see [Commands in the NDB Cluster Management Client](#)) caused `ndb_mgm` to crash on Solaris SPARC systems. (Bug #16834030)
- On Solaris SPARC platforms, batched key access execution of some joins could fail due to invalid memory access. (Bug #16818575)
- When 2 `NDB` tables had foreign key references to each other, it was necessary to drop the tables in the same order in which they were created. (Bug #16817928)
- The duplicate weedout algorithm introduced in MySQL 5.6 evaluates semi-joins such as subqueries using `IN`) by first performing a normal join between the outer and inner table which may create duplicates of rows from the outer (and inner) table and then removing any duplicate result rows from the outer table by comparing their primary key values. Problems could arise when `NDB` copied `VARCHAR` values using their maximum length, resulting in a binary key image which contained garbage past the actual lengths of the `VARCHAR` values, which meant that multiple instances of the same key were not binary-identical as assumed by the MySQL server.

To fix this problem, `NDB` now zero-pads such values to the maximum length of the column so that copies of the same key are treated as identical by the weedout process. (Bug #16744050)

- `DROP DATABASE` failed to work correctly when executed against a database containing `NDB` tables joined by foreign key constraints (and all such tables being contained within this database), leaving these tables in place while dropping the remaining tables in the database and reporting failure. (Bug #16692652, Bug #69008)
- When using `firstmatch=on` with the `optimizer_switch` system variable, pushed joins could return too many rows. (Bug #16664035)
- A variable used by the batched key access implementation was not initialized by `NDB` as expected. This could cause a “batch full” condition to be reported after only a single row had been batched, effectively disabling batching altogether and leading to an excessive number of round trips between `mysqld` and `NDB`. (Bug #16485658)
- When started with `--initial` and an invalid `--config-file (-f)` option, `ndb_mgmd` removed the old configuration cache before verifying the configuration file. Now in such cases, `ndb_mgmd` first checks for the file, and continues with removing the configuration cache only if the configuration file is found and is valid. (Bug #16299289)
- Creating more than 32 hash maps caused data nodes to fail. Usually new hashmaps are created only when performing reorganization after data nodes have been added or when explicit partitioning

is used, such as when creating a table with the `MAX_ROWS` option, or using `PARTITION BY KEY() PARTITIONS n`. (Bug #14710311)

- Setting `foreign_key_checks = 0` had no effect on the handling of `NDB` tables. Now, doing so causes such checks of foreign key constraints to be suspended—that is, has the same effect on `NDB` tables as it has on `InnoDB` tables. (Bug #14095855, Bug #16286309)

References: See also: Bug #16286164.

## Changes in MySQL NDB Cluster 7.3.1 (5.6.10-ndb-7.3.1) (2013-04-17, Development Milestone)

- [Based on MySQL Server 5.6](#)
- [MySQL Cluster GUI Configuration Wizard](#)
- [Support for Foreign Key Constraints](#)
- [NoSQL Connector for JavaScript \(Node.js\)](#)
- [Functionality Added or Changed](#)
- [Bugs Fixed](#)

### Based on MySQL Server 5.6

- **Important Change:** MySQL NDB Cluster SQL nodes are now based on MySQL Server 5.6. For information about feature additions and other changes made in MySQL 5.6, see [What Is New in MySQL 5.6](#).

The `mysqld` binary provided with MySQL NDB Cluster 7.3.1 is based on MySQL Server 5.6.10, and includes all MySQL Server 5.6 feature enhancements and bug fixes found in that release; see [Changes in MySQL 5.6.10 \(2013-02-05, General Availability\)](#), for information about these.

### MySQL Cluster GUI Configuration Wizard

- **Important Change:** The MySQL NDB Cluster distribution now includes a browser-based graphical configuration wizard that assists the user in configuring and deploying a MySQL NDB Cluster. This deployment can consist of an arbitrary number of nodes (within certain limits) on the user machine only, or include nodes distributed on a local network. The wizard can be launched from the command line (using the `ndb_setup` utility now included in the binary distribution) or a desktop file browser.

For more information about this tool, see [The NDB Cluster Auto-Installer](#).

### Support for Foreign Key Constraints

- **Important Change:** MySQL NDB Cluster now supports foreign key constraints between `NDB` tables, including support for `CASCADE`, `SET NULL`, and `RESTRICT` and `NO ACTION` reference options for `DELETE` and `UPDATE` actions. (MySQL currently does not support `SET DEFAULT`.)

MySQL requires generally that all child and parent tables in foreign key relationships employ the same storage engine; thus, to use foreign keys with MySQL NDB Cluster tables, the child and parent table must each use the `NDB` storage engine. (It is not possible, for example, for a foreign key on an `NDB` table to reference an index of an `InnoDB` table.)

Note that MySQL NDB Cluster tables that are explicitly partitioned by `KEY` or `LINEAR KEY` may contain foreign key references or be referenced by foreign keys (or both). This is unlike the case with `InnoDB` tables that are user partitioned, which may not have any foreign key relationships.

You can create an [NDB table](#) having a foreign key reference on another [NDB table](#) using `CREATE TABLE ... [CONSTRAINT] FOREIGN KEY ... REFERENCES`. A child table's foreign key definitions can be seen in the output of `SHOW CREATE TABLE`; you can also obtain information about foreign keys by querying the `INFORMATION_SCHEMA.KEY_COLUMN_USAGE` table.

[FOREIGN KEY Constraints](#), provides general information about foreign key support in MySQL. For more information about the syntax supported by MySQL for foreign keys, see [Using FOREIGN KEY Constraints](#).

## NoSQL Connector for JavaScript (Node.js)

- **NDB Cluster APIs:** MySQL NDB Cluster 7.3 includes support for JavaScript applications written against Node.js with MySQL NDB Cluster and MySQL Server as a data store. The Connector for JavaScript provides a domain object model similar in many ways to that employed by ClusterJ (see [The ClusterJ API and Data Object Model](#)) and can be used with either of two back-end adapters: the `ndb` adapter, which uses the NDB API to provide high-performance native access to MySQL NDB Cluster; and the `mysql-js` adapter, which uses the MySQL Server and the `node-mysql` driver available from <https://github.com/felixge/node-mysql/>.

The Connector for JavaScript is included with the MySQL NDB Cluster distribution, and contains setup programs which can assist you with installation of the connector. You must have Node.js and MySQL NDB Cluster installed prior to running the setup scripts. The `node-mysql` driver is also required for the `mysql-js` Node.js adapter; you can install this using the package management tool included with Node.js. For more information, see [MySQL NoSQL Connector for JavaScript](#).

## Functionality Added or Changed

- **Important Change:** The behavior of and values used for the `TCP_RCV_BUF_SIZE` and `TCP_SND_BUF_SIZE` TCP configuration parameters have been improved. Formerly, the default values for these parameters were 70080 and 71540, respectively—which it was later found could lead to excessive timeouts in some circumstances—with the minimum for each of them being 1. Now, the default and recommended value is 0 for both `TCP_RCV_BUF_SIZE` and `TCP_SND_BUF_SIZE`, which allows the operating system or platform to choose the send or receive buffer size for TCP sockets. (Bug #14554519)

References: See also: Bug #14168828.

- **NDB Cluster APIs:** Added `DUMP` code 2514, which provides information about counts of transaction objects per API node. For more information, see [DUMP 2514](#). See also [Commands in the NDB Cluster Management Client](#). (Bug #15878085)
- When `ndb_restore` fails to find a table, it now includes in the error output an NDB API error code giving the reason for the failure. (Bug #16329067)
- Data node logs now provide tracking information about arbitrations, including which nodes have assumed the arbitrator role and at what times. (Bug #11761263, Bug #53736)

## Bugs Fixed

- **API:** `mysqld` failed to respond when `mysql_shutdown()` was invoked from a C application, or `mysqladmin shutdown` was run from the command line. (Bug #14849574)
- When an update of an [NDB table](#) changes the primary key (or part of the primary key), the operation is executed as a delete plus an insert. In some cases, the initial read operation did not retrieve all column values required by the insert, so that another read was required. This fix ensures that all required column

values are included in the first read in such cases, which saves the overhead of an additional read operation. (Bug #16614114)

- Pushed joins executed when `optimizer_switch='batched_key_access=on'` was also in use returned incorrect results. (Bug #16437431)
- Selecting from the `INFORMATION_SCHEMA.KEY_COLUMN_USAGE` table while using tables with foreign keys caused `mysqld` to crash. (Bug #16246874, Bug #68224)
- Including a table as a part of a pushed join should be rejected if there are outer joined tables in between the table to be included and the tables with which it is joined with; however the check as performed for any such outer joined tables did so by checking the join type against the root of the pushed query, rather than the common ancestor of the tables being joined. (Bug #16199028)

References: See also: Bug #16198866.

- Some queries were handled differently with `ndb_join_pushdown` enabled, due to the fact that outer join conditions were not always pruned correctly from joins before they were pushed down. (Bug #16198866)

References: See also: Bug #16199028.

- Attempting to perform additional operations such as `ADD COLUMN` as part of an `ALTER [ONLINE | OFFLINE] TABLE ... RENAME ...` statement is not supported, and now fails with an `ER_NOT_SUPPORTED_YET` error. (Bug #16021021)
- Purging the binary logs could sometimes cause `mysqld` to crash. (Bug #15854719)

## Index

### Symbols

[#include](#), 41, 82  
[#sql-\\*](#), 32, 75  
[\\*](#), 44, 85  
[--binlog-row-image](#), 51  
[--config-file](#), 48, 89  
[--database](#), 29, 72  
[--disable-indexes](#), 16, 61  
[--exclude-databases](#), 16, 61  
[--exclude-intermediate-sql-tables](#), 32, 75  
[--exclude-missing-tables](#), 29, 72  
[--help](#), 29, 41, 72, 82  
[--include-databases](#), 16, 61  
[--initial](#), 48, 89  
[--lossy-conversions](#), 44, 85  
[--ndb-log-update-as-write](#), 14, 60  
[--ndb-log-updated-only](#), 14, 60  
[--order](#), 32, 75  
[--print-full-config](#), 37, 79  
[--print\\_data](#), 24, 68  
[--promote-attributes](#), 44, 85  
[--rebuild-indexes](#), 16, 61  
[--restore\\_data](#), 18, 44, 63, 85  
[--restore\\_meta](#), 18, 63  
[-d](#), 29, 72

- DWITH\_NDBMTD, 37, 79
- flifetime-dse, 6, 55
- .ctl files, 8, 56
- .FRM, 44
- / (slash character), 32, 75
- ~Ndb, 10, 57
- ~Ndb\_cluster\_connection(), 41, 82

## A

- ABORT BACKUP, 44, 85
- add node, 24, 68
- add nodes, 29, 72
- AFTER DELETE, 24, 68
- alignment, 48, 89
- ALTER LOGFILE GROUP, 48, 89
- ALTER ONLINE TABLE, 44, 85
- ALTER TABLE, 16, 32, 37, 51, 61, 75, 79, 91
- ALTER TABLE RENAME, 24, 68
- ALTER TABLESPACE, 48, 89
- ANALYZE TABLE, 18, 63
- AnyValue, 5
- API, 51, 91
- API nodes, 12, 29, 32, 48, 59, 72, 75, 89
- API\_REGREQ, 24, 68
- arbitration, 10, 57
- arbitrator, 51, 91
- arbitrator selection, 41, 82
- attribute demotion, 29, 72
- attribute promotion, 29, 44, 72, 85
- auto\_increment\_increment, 29, 72
- auto\_increment\_offset, 29, 72

## B

- backup, 6, 16, 18, 41, 44, 55, 61, 63, 82, 85
- backup and restore, 16
- backup files, 44, 85
- BackupCompleted, 37, 79
- backups, 7, 55
- BatchByteSize, 21, 66
- BatchSize, 21, 66
- BINARY, 44
- binary log, 44
- binary log injector, 14, 60
- binlog, 44
- binlog injector, 14, 32, 60, 75
- binlog rotation, 32, 75
- BKA, 48, 51, 89, 91
- BLOB, 18, 24, 29, 41, 44, 68, 82, 85
- BLOB tables, 24, 68
- blocks, 18, 63
- buffer allocation, 24, 68
- buffer overload, 41, 82
- buffers, 48, 89

build, 32, 75  
bulk deletes, 21, 66  
Busy error, 18, 63

## C

cache line, 24, 68  
cache\_size, 44  
cascading\_scans\_count, 9, 57  
changes  
    NDB Cluster, 54  
CHAR, 24, 32, 68, 75  
CHAR(0), 24, 68  
circular wait lock, 37, 79  
close\_clnt(), 18, 63  
CLOSE\_COMREQ, 29, 32, 72, 75  
cluster failure and recovery, 18, 63  
Cluster temporary unavailable, 24, 68  
ClusterJDatastoreException, 9  
ClusterJPA, 6  
CLUSTERLOG, 48, 89  
ClusterMgr, 10, 57  
CMakeLists.txt, 32, 44, 75  
column accessors, 29  
Column::getSizeInBytesForRecord(), 18, 63  
COLUMN\_FORMAT DYNAMIC, 14, 60  
COMMENT, 21, 66  
commitAckMarker, 48, 89  
COMMIT\_ACK, 41, 82  
COMMIT\_CONF, 41, 82  
COMMIT\_REF, 41, 82  
community, 44, 85  
community contribution, 44, 85  
compiling, 6, 14, 24, 37, 44, 55, 60, 68, 79  
composite keys, 16, 61  
configuration, 6, 51, 91  
conflict detection, 32  
conflict resolution, 24, 32, 51  
ConnectBackoffMaxTime, 29, 72  
connection-timeout, 44, 85  
constraints, 51, 91  
COPY\_FRAGREQ, 18, 63  
CREATE DATABASE, 24, 68  
CREATE LOGFILE GROUP, 48, 89  
CREATE TABLE, 9, 51, 57, 91  
CREATE TABLESPACE, 48, 89  
createEvent(), 16, 61  
create\_old\_temporals, 16  
cross-version replication, 16  
c\_lcpState, 32, 75

## D

Data length too long, 32  
data node halts, 24

data node restarts, 10, 57  
data nodes, 32, 44, 48, 51, 75, 85, 89, 91  
data redistribution, 37, 79  
DataMemory, 29, 44, 72, 85  
date and time types, 48  
DBDICT, 16, 18, 24, 61, 63, 68  
DBDIH, 9, 32, 44, 57, 75, 85  
Dbdih::CheckGcpStopLab(), 41, 82  
DBLQH, 24, 37, 41, 44, 48, 68, 79, 82, 85, 89  
Dblqh::TcConnectRec, 48, 89  
DBSPJ, 10, 24, 44, 57, 68, 85  
DBTC, 5, 10, 29, 32, 37, 41, 44, 48, 54, 57, 72, 75, 79, 82, 85, 89  
DBTRIX, 24, 68  
DBTUP, 44, 85  
DbtupExecQuery.cpp, 29, 72  
DBTUX, 5, 54  
DBUTIL, 44, 85  
deadlocks, 16, 18, 61, 63  
debug, 14, 48, 60, 89  
DefaultHashMapSize, 21, 66  
DefaultOperationRedoProblemAction, 18, 63  
DEFINER, 32, 75  
DELETE, 32  
DESC, 44, 85  
DICT manager, 44, 85  
DICT master, 24, 68  
DictHashMapInfo::HashMap, 21, 66  
Dictionary, 16, 61  
Dictionary::dropTable(), 37, 79  
Dictionary::getTable(), 21, 66  
DIH master, 24, 68  
DIH\_SCAN\_GET\_NODESREQ, 32, 75  
disconnection, 41  
DISCONNECT\_REP, 21, 66  
DiskBufferPageEntries, 24, 68  
distributions, 44  
DML, 44  
documentation, 44, 85  
DROP COLUMN, 44, 85  
DROP DATABASE, 24, 48, 68, 89  
drop events, 14  
DROP FOREIGN KEY, 32, 75  
DROP INDEX, 41, 82  
drop operations, 37, 79  
DROP TABLE, 14, 37, 48, 60, 79, 89  
dropEvent(), 16, 61  
dry-scp, 44, 85  
DUMP, 51, 91  
DUMP 2610, 29, 72  
DUMP 7019, 32, 75  
DUMP 7024, 32, 75  
DUMP 7027, 5, 54  
DUMP 9991, 21, 66  
DUMP codes, 5, 54



duplicate IDs, 44, 85  
duplicate key, 16, 61  
duplicate keys, 12, 59  
DYNAMIC, 7, 55

## E

ENABLE\_COMREQ, 32, 75  
epochs, 37, 44, 85  
error 1419, 16, 61  
Error 218, 24, 68  
Error 221, 24, 68  
error 240, 6, 55  
error 4547, 24, 68  
error 746, 16, 61  
Error 899, 8, 56  
error 899, 44, 85  
error handling, 14, 37, 44, 51, 60, 79, 85, 91  
error logging, 24  
errors, 12, 29, 32, 37, 44, 59, 72, 75, 79, 85  
ERROR\_FILE\_EXISTS, 44, 85  
ER\_DUP\_ENTRY, 29, 72  
event buffer, 32, 75  
event queue, 18, 63  
Event::setTable(), 44, 85  
examples, 24, 37, 41, 68, 79, 82  
exceptions tables, 24  
EXECUTE\_DIRECT, 44, 85

## F

failure handling, 14  
FAIL\_REP, 21, 66  
File not found error, 18, 63  
FILES, 5  
fired triggers, 24, 68  
firstmatch optimization, 48, 89  
FKs, 37, 41, 44, 51, 79, 82, 85, 91  
flags, 51  
FLUSH BINARY LOG, 24  
FLUSH LOGS, 32  
forced shutdown, 10, 57  
foreign keys, 8, 9, 10, 14, 16, 32, 37, 41, 44, 48, 56, 57, 57, 60, 61, 75, 79, 82, 85, 89  
foreign\_key\_checks, 48, 89  
fractional seconds, 48  
fragment scans, 44, 85  
freeSSHd, 44, 85

## G

garbage collection, 18  
gcc, 6, 55  
GCI, 16, 29, 61, 72  
GCP, 21, 41, 66, 82  
GCP lag, 41, 82

GCP monitor, 24, 68  
 GCP stop, 24, 68  
 GCP\_COMMIT, 41, 82  
 getConnectionPoolSessionCounts(), 10  
 getNdbObjectName(), 32, 75  
 GetSystemTimeAsFileTime, 41, 82  
 getTupleIdFromNdb(), 29, 72  
 GET\_TABINFOREQ, 18, 63  
 GET\_TABLINFOREF, 18, 63  
 GET\_TABINFOREQ, 18, 63  
 grandchild, 8, 56  
 group\_relay\_log\_pos, 37  
 GSIRReader, 21, 66  
 GSN\_TCKEY\_FAILREFCONF\_R, 32, 75  
 g\_thr\_repository, 24, 68

## H

hash maps, 37, 48, 79, 89  
 HA\_ERR\_NO\_CONNECTION, 24, 68  
 HeartbeatIntervalMgmdMgmd, 44, 85  
 help, 32, 44, 75, 85

## I

IBM POWER, 24, 68  
 identifiers, 24, 68  
 Important Change, 6, 14, 16, 18, 21, 37, 51, 60, 61, 63, 66, 79, 91  
 Important Note, 14, 51, 60  
 IN, 4  
 IN(), 48, 89  
 Incompatible Change, 10, 57  
 index statistics, 16, 18, 61, 63  
 IndexMemory, 37, 79  
 INDEX\_STAT, 44, 85  
 INFORMATION\_SCHEMA, 5, 32, 51, 75, 91  
 InitialLogFileGroup, 32, 75  
 INITIAL\_SIZE, 48, 89  
 INSERT, 18, 32, 63, 75  
 insert, 44, 85  
 installer, 44, 85  
 installing, 48, 89  
 internals, 29, 44, 72, 85  
 invalidateLcplInfoAfterSr(), 14, 60  
 isConsistent(), 14, 60

## J

JavaScript, 51, 91  
 job buffer, 5, 54  
 job buffers, 44, 85  
 job scheduler, 41, 44, 82, 85  
 JOIN, 44, 85

## K

KEY\_COLUMN\_USAGE, 51, 91

**L**

last page of log, 21, 66  
latin1, 32  
LCP, 5, 14, 44, 54, 60, 85  
LCP scans, 18, 63  
LCP states, 32, 75  
LcpScanProgressTimeout, 18, 44, 63, 85  
LCP\_SKIP, 18, 63  
LDM, 10, 57  
LDM threads, 44, 85  
leak, 29, 72  
LGMAN, 5, 41, 54, 82  
LIKE, 44, 85  
livelocks, 18, 63  
localtime(), 41, 82  
localtime\_r(), 41, 82  
locking, 24, 68  
lock\_ndb\_objects(), 9, 57  
log files, 6, 55  
log messages, 37, 79  
log rotation, 32  
logging, 5, 14, 16, 51, 54, 60, 61, 91  
long long, 21, 66  
LongMessageBuffer, 8, 24, 37, 44, 56, 68, 79, 85  
LONGVARBINARY, 4  
Loopback transporter, 10, 57  
lost connection, 10  
lower\_case\_table\_names, 44, 85  
LQH > 4, 37, 79  
LQHKEYREQ, 37, 79

**M**

malloc(), 9, 48, 57, 89  
manifest.mf, 32, 75  
Master, 44, 85  
MASTER\_LCPREQ, 32, 75  
MaxBufferedEpochs, 12, 59  
MaxNoOfConcurrentScans, 37, 79  
MaxNoOfExecutionThreads, 18, 44, 63, 85  
MaxNoOfFiredTriggers, 24, 68  
MaxNoOfLocalScans, 37, 79  
MaxScanBatchSize, 21, 66  
max\_binlog\_size, 32  
max\_failure\_time, 21, 66  
MAX\_NULL\_BITS, 14, 60  
memory allocation, 29, 72  
memory barrier macros, 37, 79  
message corruption, 24, 68  
metadata, 12, 29, 59, 72  
MGM API, 29, 48, 72, 89  
mgmd, 44, 85  
Microsoft Windows, 41, 44, 48, 82, 85, 89  
MISSING\_DATA, 29, 72

MT scheduler, 18, 63  
MT Scheduler, 24, 68  
mt-scheduler, 18, 37, 63, 79  
multiple connections, 21, 66  
multiple mysqlds, 16, 61  
multiple-statement transactions, 16, 61  
multithreaded, 18, 63  
mvn\_install\_ndbjtie.sh, 32, 75  
my.cnf, 44, 85  
MySQL NDB ClusterJ, 6, 9, 10, 18, 24, 29, 32, 41, 48  
mysql-js, 51, 91  
mysql.proc, 48, 89  
mysqladmin, 51, 91  
mysqld, 14, 16, 21, 24, 41, 44, 48, 51, 60, 61, 66, 68, 82, 85, 89, 91  
mysql\_shutdown, 51, 91  
mysql\_upgrade, 24, 48, 68, 89

## N

Ndb, 24, 68  
NDB Cluster, 4, 5, 5, 6, 8, 9, 10, 12, 14, 16, 18, 21, 24, 29, 32, 37, 41, 44, 48, 51, 54, 55, 56, 57, 57, 59, 60, 61, 63, 66, 68, 72, 75, 79, 82, 85, 89, 91  
NDB Cluster APIs, 4, 9, 10, 12, 14, 16, 18, 21, 24, 29, 32, 37, 41, 44, 48, 51, 57, 57, 59, 60, 61, 63, 66, 68, 72, 75, 79, 82, 85, 89, 91  
NDB Disk Data, 7, 24, 32, 41, 44, 48, 55, 68, 75, 82, 85, 89  
NDB distributed triggers, 10, 57  
Ndb object, 32, 41, 75, 82  
NDB Replication, 5, 6, 14, 16, 24, 29, 32, 37, 44, 51  
NDB\$EPOCH, 32  
NDB\$EPOCH\_TRANS, 32  
Ndb.hpp, 37, 79  
Ndb::closeTransaction(), 37, 79  
Ndb::computeHash(), 48, 89  
Ndb::getNextEventOpInEpoch3(), 5  
Ndb::init(), 41, 82  
Ndb::isExpectingHigherQueuedEpochs(), 18, 63  
Ndb::nextEvent(), 18, 32, 63, 75  
Ndb::pollEvents(), 18, 32, 63, 75  
ndbd, 41, 82  
NdbDictInterface::dictSignal(), 24, 68  
NdbDictionary, 44, 85  
ndbd\_redo\_log\_reader, 41, 82  
NdbEventBuffer::m\_latestGCI, 16, 61  
NDBFS, 5  
NdbFS, 44, 85  
NdbIndexOperation, 12, 59  
NdbIndexScanOperation::setBound(), 4  
ndbinfo, 24, 32, 48, 68, 75, 89  
ndbinfo.memoryusage, 37, 41, 79, 82  
ndbinfo.ndb\$pools, 24, 68  
ndbinfo.sql, 32, 75  
NDBJTie, 6  
ndbmemcache, 41, 44, 48, 51

ndbmtd, 24, 29, 32, 37, 41, 44, 68, 72, 75, 79, 82, 85  
NdbReceiver, 21, 66  
NdbRecord, 21, 37, 66, 79  
NdbScanOperation, 24, 68  
NdbScanOperation::close(), 24, 68  
NdbScanOperation::nextresult(), 29, 72  
NdbTick(), 41, 82  
NdbTick\_Elapsed(), 41, 82  
NdbTransaction, 24, 68  
NdbTransaction::setSchemObjectOwnerChecks(), 21, 66  
ndb\_autoincrement\_prefetch\_sz, 29, 72  
ndb\_binlog\_index, 32, 75  
ndb\_binlog\_setup(), 12, 59  
Ndb\_cluster\_connection, 9, 18, 24, 29, 37, 41, 57, 63, 68, 72, 79, 82  
ndb\_desc, 16, 61  
ndb\_error\_reporter, 44, 85  
ndb\_global.h, 37, 79  
ndb\_index\_cache\_entries, 37, 79  
ndb\_index\_stat, 24, 68  
ndb\_index\_stat\_freq, 37, 79  
ndb\_index\_stat\_option, 18, 63  
ndb\_join\_pushdown, 41, 51, 82, 91  
Ndb\_last\_commit\_epoch\_server, 21, 24, 66  
Ndb\_last\_commit\_epoch\_session, 21, 24, 66  
ndb\_logevent\_get\_next(), 29, 72  
ndb\_logevent\_get\_next2(), 29, 72  
NDB\_MAX\_TUPLE\_SIZE, 14, 60  
ndb\_mgm, 37, 44, 48, 79, 85, 89  
ndb\_mgmd, 10, 37, 44, 48, 57, 79, 85, 89  
ndb\_mgm\_event\_category, 29, 48, 72, 89  
ndb\_print\_file, 29, 72  
ndb\_restore, 4, 8, 9, 12, 14, 16, 18, 24, 29, 32, 37, 44, 51, 56, 57,  
59, 60, 61, 63, 68, 72, 75, 79, 85, 91  
ndb\_schema, 14, 60  
ndb\_select\_all, 32, 75  
ndb\_show\_tables, 29, 37, 72, 79  
Ndb\_slave\_max\_replicated\_epoch, 24  
nextEvent(), 14, 16, 60, 61  
nextEvent2(), 16, 61  
node connections, 41, 82  
node failover, 24, 68  
node failure, 16, 32, 44, 61, 75, 85  
node failure handling, 5, 12, 24, 29, 32, 37, 48, 54, 59, 68, 72, 75, 79, 89  
node recovery, 29, 72  
node restart, 5, 14, 54, 60  
node restarts, 24, 68  
node shutdown, 48, 89  
node starts, 16, 61  
node state, 29, 72  
node takeover, 24, 29, 32, 44, 68, 72, 75, 85  
Node.js, 51, 91  
nodeFailure error, 10, 57  
NodeInfo, 10, 57  
NODE\_FAILREP, 24, 32, 68, 75

noOfConnectedNodes, 10, 57  
NoOfReplicas, 21, 66  
NotMaster, 24, 68  
NULL, 4, 7, 41, 44, 55, 82, 85  
null, 9, 57

## O

object creation, 18, 63  
object destruction, 18, 63  
OFFLINE, 51, 91  
ON UPDATE CASCADE, 8, 56  
ONLINE, 51, 91  
online add nodes, 44, 85  
online reorganization, 24, 68  
OPEN\_COMORD, 32, 75  
optimizer, 51, 91  
optimizer\_switch, 51, 91  
ORDER BY, 5, 44, 85  
output buffers, 41, 82  
O\_SYNC, 5

## P

Packaging, 41, 44, 48, 82, 85, 89  
parallel schema operations, 8, 56  
parent batch completion, 44, 85  
PARTITION, 21, 66  
partition info, 16, 61  
Performance, 24, 44, 68, 85  
Performance Schema, 24, 68  
PK, 24, 32, 68, 75  
pollEvent(), 14, 60  
pollEvents(), 14, 16, 60, 61  
pollEvents2(), 14, 16, 60, 61  
preallocate, 44  
PREPARE\_SEIZE\_ERROR, 12, 59  
primary key, 51, 91  
primary keys, 16, 61  
PROPERTY\_CLUSTER\_CONNECT\_TIMEOUT\_MGM, 24  
purging logs, 51, 91  
pushdown conditions, 51, 91  
pushdown joins, 24, 68  
pushed joins, 51, 91  
P\_TAIL\_PROBLEM, 18, 63

## Q

QMGR, 5, 29, 32, 37, 41, 48, 54, 72, 75, 79, 82, 89  
queries, 29  
QueryPerformanceCounters, 41, 82

## R

rand(), 8, 56  
range checking, 44, 85  
read-only, 14

READ\_ALL, 44, 85  
READ\_PACKED, 44, 85  
RealTimeScheduler, 41, 44, 82, 85  
receive buffers, 29, 72  
receive thread, 44, 85  
receive threads, 18, 44, 63, 85  
ReceiveBufferMemory, 44, 85  
receiver thread, 32, 75  
RecordSpecification, 24, 68  
redo, 18, 63  
relay log rotation, 37  
releaseOp, 24, 68  
RENAME, 51, 91  
REORGANIZE PARTITION, 32, 44, 75, 85  
Replication, 37  
reportConnected(), 10, 57  
reportDisconnected(), 10, 57  
resource recovery, 37, 79  
restarts, 18, 21, 32, 48, 63, 66, 75, 89  
RestartSubscriberConnectTimeout, 32, 75  
restore, 16, 61  
result buffers, 21, 66  
rolling restart, 29, 72  
RPM, 44, 85

## S

ScanFrag watchdog, 24, 68  
scans, 7, 55  
SCAN\_FRAGREQ, 24, 68  
schema events, 14, 60  
schema operations, 14, 60  
schema transactions, 24, 68  
schemaTrans, 24, 68  
SELECT, 18, 63  
send buffer, 10, 18, 57, 63  
send buffers, 24, 68  
send threads, 18, 63  
SendBuffer, 10, 57  
sendFragmentedSignal(), 24, 68  
sendSignal(), 29, 72  
sendSignalNoRelease(), 29, 72  
server system variables, 37, 79  
ServerPort, 37, 79  
setNdbObjectName(), 32, 75  
setPartitionKey(), 41  
setsockopt(), 51, 91  
SharedGlobalMemory, 32, 75  
SHOW, 44, 85  
SHOW BINLOG EVENTS, 24  
signal corruption, 24, 68  
signal dump, 29, 72  
signal handling, 21, 32, 41, 66, 75, 82  
signal ID, 29, 72

- slave, 37
- slave\_parallel\_workers, 6
- Solaris, 37, 41, 79, 82
- SPARC, 48, 89
- spintimer, 14, 60
- SPJ, 48, 89
- SQL node, 12, 51, 59, 91
- SQL nodes, 12, 59
- stack overflow, 21, 66
- START BACKUP, 37, 44, 79, 85
- start phase 101, 32, 75
- start, stop, restart, 16, 61
- StartConnectBackoffMaxTime, 29, 72
- startup, 44, 85
- STOP -f, 14, 60
- Stuck in Send, 18, 63
- subqueries, 48, 89
- subscription events, 14, 60
- SUB\_GCP\_COMPLETE\_ACK, 12, 21, 59, 66
- SUB\_GCP\_COMPLETE\_REP, 14, 21, 60, 66
- SUB\_START\_CONF, 14, 60
- SUB\_START\_REQ, 48, 89
- SUMA, 12, 14, 16, 59, 60, 61
- sysfile, 16, 61
- system restart, 9, 57

## T

- tabLcpStatus, 32, 75
- table corruption, 41, 82
- tablespace, 44, 85
- Table\_Map, 5
- TC, 29, 32, 72, 75
- TcDumpApiConnectRecSummary, 51, 91
- TCKEY\_FAILCONF, 32, 75
- TCKEY\_FAILREF, 32, 75
- TCP transporter, 10, 24, 57, 68
- TcpBind\_INADDR\_ANY, 37, 79
- TCP\_RCV\_BUF\_SIZE, 51, 91
- TCP\_SND\_BUF\_SIZE, 51, 91
- TCROLLBACKREP, 37, 79
- TEXT, 24, 29, 44, 68, 72, 85
- TE\_SUBSCRIBE, 14, 60
- thread contention, 18, 63
- thread synchronization, 18, 63
- ThreadConfig, 14, 29, 37, 60, 72, 79
- throttling, 41, 82
- thr\_safe\_pool, 24, 68
- TimeBetweenEpochs, 12, 59
- TimeBetweenGlobalCheckpointsTimeout, 21, 66
- timeouts, 14, 18, 24, 32, 37, 41, 44, 60, 63, 75, 79, 82, 85
- timers, 41, 82
- TIME\_SIGNAL, 37, 79
- TINYBLOB, 16, 61



- TINYTEXT, 29, 72
- TOTAL\_BUCKETS\_INIT, 14, 60
- trace, 29, 72
- trace files, 24, 37, 68, 79
- transaction ID, 9, 57
- transactions, 24, 68
- TRANSID\_AI, 37, 79
- Transporter::doDisconnect(), 29, 72
- Transporter::doSend(), 24, 68
- TransporterFacade, 10, 57
- TransporterFacade::deliver\_signal(), 18, 63
- TransporterFacade::reset\_send\_buffer(), 24, 68
- TransporterFacade::theMutexPtr, 24, 68
- TransporterRegistry, 10, 57
- TransporterRegistry::connect\_server(), 29, 72
- TransporterRegistry::performReceive(), 29, 72
- TransporterRegistry::performRecive(), 24, 68
- TransporterRegistry::reportError(), 10, 57
- TransporterRegistry::update\_connections(), 24, 68
- transporters, 29, 72
- TRANS\_END\_REQ, 24, 68
- triggers, 24, 68
- TRMAN, 29, 72
- TRPMAN, 32, 75
- TUPKEYREQ, 44, 85
- TUP\_COMMITREQ, 24, 68
- type conversions, 24, 68
- types, 29, 72

## U

- UINT\_MAX64, 41, 82
- undo buffer, 32, 75
- undo log, 18, 63
- UNDO\_BUFFER\_SIZE, 48, 89
- UNIQUE, 24, 68
- unique index, 12, 59
- unique indexes, 12, 59
- unique key, 16, 61
- unique key checks, 14, 60
- unique keys, 12, 16, 59, 61
- unlock\_ndb\_objects(), 9, 57
- updates, 37, 51, 79, 91
- upgrades, 16, 37, 79
- UTF8, 32, 75
- utf8, 32

## V

- values table, 51
- VARBINARY, 21, 24, 66, 68
- VARCHAR, 21, 24, 41, 48, 66, 68, 82, 89
- Visual Studio, 41, 82

## **W**

WAIT STARTED, 44, 85  
WAIT\_EVENT, 14, 60  
watchdog checks, 41, 82  
watchdog state, 44, 85  
watchdog warnings, 29, 72  
WHERE, 44, 85  
wildcards, 24  
work threads, 18, 63

## **X**

x86, 41, 82

## **Z**

ZFAIL\_CLOSING, 21, 66