

---

# MySQL NDB Cluster 9.4 Release Notes

## Abstract

This document contains release notes for the changes in each release of MySQL NDB Cluster that uses version 9.4 of the [NDB \(NDBCLUSTER\)](#) storage engine.

Each NDB Cluster 9.4 release is based on a mainline MySQL Server release and a particular version of the [NDB](#) storage engine, as shown in the version string returned by executing `SELECT VERSION()` in the `mysql` client, or by executing the `ndb_mgm` client `SHOW` or `STATUS` command; for more information, see [MySQL NDB Cluster 9.4](#).

For general information about features added in NDB Cluster 9.4, see [What is New in MySQL NDB Cluster 9.4](#). For a complete list of all bug fixes and feature changes in MySQL NDB Cluster, please refer to the changelog section for each individual NDB Cluster release.

For additional MySQL 9.4 documentation, see the [MySQL 9.4 Reference Manual](#), which includes an overview of features added in MySQL 9.4 that are not specific to NDB Cluster ([What Is New in MySQL 9.4](#)), and discussion of upgrade issues that you may encounter for upgrades from MySQL 9.3 to MySQL 9.4 ([Changes in MySQL 9.4](#)). For a complete list of all bug fixes and feature changes made in MySQL 9.4 that are not specific to [NDB](#), see [MySQL 9.4 Release Notes](#).

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

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

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

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

Document generated on: 2025-10-14 (revision: 30609)

## Table of Contents

Preface and Legal Notices .....	1
Changes in MySQL NDB Cluster 9.4.0 (2025-07-23, Innovation Release) .....	3
Index .....	6

## Preface and Legal Notices

This document contains release notes for the changes in each release of MySQL NDB Cluster that uses version 9.4 of the [NDB](#) storage engine.

### Legal Notices

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

#### License Restrictions

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate,

broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

### **Warranty Disclaimer**

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

### **Restricted Rights Notice**

If this is software, software documentation, data (as defined in the Federal Acquisition Regulation), or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

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

### **Hazardous Applications Notice**

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

### **Trademark Notice**

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

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

### **Third-Party Content, Products, and Services Disclaimer**

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

### **Use of This Documentation**

This documentation is NOT distributed under a GPL license. Use of this documentation is subject to the following terms:

You may create a printed copy of this documentation solely for your own personal use. Conversion to other formats is allowed as long as the actual content is not altered or edited in any way. You shall not publish or distribute this documentation in any form or on any media, except if you distribute the documentation in a manner similar to how Oracle disseminates it (that is, electronically for download on a Web site with the software) or on a CD-ROM or similar medium, provided however that the documentation is disseminated together with the software on the same medium. Any other use, such as any dissemination of printed copies or use of this documentation, in whole or in part, in another publication, requires the prior written consent from an authorized representative of Oracle. Oracle and/or its affiliates reserve any and all rights to this documentation not expressly granted above.

## Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at

<http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc>.

## Access to Oracle Support for Accessibility

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit

<http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> or visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> if you are hearing impaired.

## Changes in MySQL NDB Cluster 9.4.0 (2025-07-23, Innovation Release)

MySQL NDB Cluster 9.4.0 is a new Innovation release of NDB Cluster, based on MySQL Server 9.4 and including features in version 9.4 of the [NDB](#) storage engine, as well as fixing recently discovered bugs in previous NDB Cluster releases.

**Obtaining MySQL NDB Cluster 9.4.** NDB Cluster 9.4 source code and binaries can be obtained from <https://dev.mysql.com/downloads/cluster/>.

For an overview of major changes made in NDB Cluster 9.4, see [What is New in MySQL NDB Cluster 9.4](#).

This release also incorporates all bug fixes and changes made in previous NDB Cluster releases, as well as all bug fixes and feature changes which were added in mainline MySQL 9 through MySQL 9.4.0 (see [Changes in MySQL NDB Cluster 9.4.0 \(2025-07-23, Innovation Release\)](#)).



### Note

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

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

## Functionality Added or Changed

- **MySQL NDB ClusterJ:** The handling of "table not found" conditions in MySQL NDB Cluster has been improved for consistency. Previously, ClusterJ sometimes threw a [ClusterJUserException](#) and at other times a [ClusterJDatastoreException](#) when a table was not found. Now, all such

conditions are reported using a new exception class, `ClusterJTableException`, after checking the NDB error code. Additionally, a new property, `com.mysql.clusterj.table.wait.msec`, has been introduced to allow configuration of a wait-and-retry loop for "table not found" conditions.

The new feature also allows a better handling of schema changes that take place on a NDB Cluster. (Bug #37884530)

References: See also: Bug #37861635, Bug #118032.

- **MySQL NDB ClusterJ:** A single `SessionFactory` can now handle multiple databases. This is achieved through a new variant of `SessionFactory.getSession()` that takes a database name as its argument. In this new design, an umbrella `SessionFactory` manages a set of `SessionFactory` objects, each of which connects to a single database through a shared pool of connections, so that connections to databases in the same NDB cluster share the same connection.

This new feature is not enabled by default, but requires setting the property `com.mysql.clusterj.multi.database` to true, and `com.mysql.clusterj.connection.pool.size` must not be 0 (default and recommended value is 1). (Bug #37792409)

References: See also: Bug #36974092, Bug #115884.

- **MySQL NDB ClusterJ:** When a user created two `ClusterJ SessionFactory` objects with the same connection string but different database names, two different `Ndb_cluster_connection` objects were created for the two databases, which consumed a lot of extra resources.

Starting from 9.4.0, different `SessionFactory` objects will share the same underlying connection from a global connection pool when connecting to different databases in the same NDB cluster, reducing resource utilization and increasing performance. In this new implementation, each NDB Cluster now has a system name, and the connection to a cluster is by a `Connection.java` interface, which is returned by a new `Session.getConnection()` method. See the [ClusterJ API Reference](#) for details. (Bug #36974092)

References: See also: Bug #37792409, Bug #117893.

- **MySQL NDB ClusterJ:** ClusterJ has also been improved by having the implementations of the `finalize()` method replaced by the Cleaner API. (Bug #30719755)
- **MySQL NDB ClusterJ:** Enhancements have been made to improve the scalability of ClusterJ applications when using the `session.newInstance` and `session.release` interfaces. A new session cache, which is a cache of Ndb objects maintained in `DbFactoruImpl`, has been introduced, reducing an application's overhead and improving its performance. The cache size can be controlled using the `com.mysql.clusterj.max.cached.instances` configuration parameter, which enables the session cache by default with a maximum size of 100. (Bug #102510, Bug #32474777)
- Timestamps in NDB node logs are now printed by default using microsecond resolution. Data nodes can enable this feature explicitly with the data node `--ndb-log-timestamps=UTC` option; management nodes can also do so using the `ndb_mgmd` option `--ndb-log-timestamps=UTC`. For backwards compatible behavior, set this option to `LEGACY`, which uses the system time zone and resolution in seconds, as in releases prior to NDB 8.4.6; in NDB 9.4, `UTC` is the default.

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

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

- MySQL Cluster now supports explicit default expressions for columns of NDB tables used with the MySQL server. For example, you can now create a an NDB table like the one shown here:

```
CREATE TABLE t1 (  
  i INT          DEFAULT 0,
```

```
c VARCHAR(10) DEFAULT '',
f FLOAT          DEFAULT (RAND() * RAND()),
b BINARY(16)     DEFAULT (UUID_TO_BIN(UUID())),
d DATE           DEFAULT (CURRENT_DATE + INTERVAL 1 YEAR)
) ENGINE=NDBCLUSTER;
```

For information about supported column data types and default behavior, see [Data Type Default Values](#).



#### Note

Default column value expressions are supported only when NDB tables are accessed using the MySQL server (`mysqld`), and are not visible to NDB API applications.

(WL #16678)

## Bugs Fixed

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

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

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

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

- **MySQL NDB ClusterJ:** Handling of schema changes failed when multiple threads were involved. To improve the handling, `session.unloadSchema()` has been changed, and three new public methods have been defined: `boolean isStaleMetadata()`, `boolean isSchemaChangePending()`, `void awaitSchemaChange()`. When ClusterJ detects that a schema change has caused a domain type handler to become invalid, it throws an `ClusterJDatastoreException`, and the application can now do the following:
  - Call `isStaleMetadata()`.
  - If `isStaleMetadata()` returns true, the exception was caused by stale metadata due to schema changes. The thread can then call `session.unloadSchema()`. When multiple threads are involved, the first thread to enter `unloadSchema()` will acquire a lock and refresh the metadata. Any other threads calling `getDomainTypeHandler()` for the domain object while the refresh is in progress receive an exception. The threads can then retry the data operation when the refresh is finished.

After the table metadata has been refreshed, ClusterJ writes a log message containing the former and current version numbers of the table metadata.

- If `isStaleMetadata()` returns false, then user can call `isSchemaChangePending()` to see if another thread is trying to refresh the metadata. If that is the case, the user can use `awaitSchemaChange()` to wait for the lock to be released before it retries the data operation.

Our thanks to Salman Niazi of Hopsworks for contributing the changes and tests for `Session.unloadSchema()`, which make this feature possible. (Bug #37861635, Bug #37856493)

- **MySQL NDB ClusterJ:** Added support for using `LONGVARCHAR` and `DATE` types as primary keys in Cluster/J.

Our thanks to Mikael Ronström of Hopsworks for the contribution. (Bug #37782638)

- **MySQL NDB ClusterJ:** Logging for ClusterJ has been improved by adopting the pattern in `java.util.logging` of Java 1.8, in which a log message is constructed only when the corresponding log level has been enabled, thus improving the efficiency of logging. (Bug #37722667)
- **MySQL NDB ClusterJ:** When users attempted to set a `NOT NULL` string column to `NULL`, an empty string was inserted.

Our thanks to Salman Niazi of Hopsworks for contributing the fix for the issue. (Bug #117205, Bug #37476251)

- Improved password handling for file system encryption.

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

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

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

## Index

### C

`CREATE TABLESPACE`, 3

### D

`DEFAULT`, 3  
deprecation, 3

### E

encryption, 3  
errors, 3

### L

logging, 3

### M

MySQL NDB ClusterJ, 3

### N

NDB Client Programs, 3  
NDB Cluster, 3  
null, 3