MySQL Cluster 4.1 Release Notes

Abstract

Beginning with MySQL 4.1.14, MySQL Cluster changes for MySQL 4.1 Server releases can be found in the MySQL 4.1 Server Release Notes.

This document contains release notes for older releases of MySQL Cluster (before 4.1.14) that use version 4.1 of the NDBCLUSTER storage engine.

For additional MySQL 4.1 documentation, see the MySQL 3.23, 4.0, 4.1 Reference Manual. For a complete list of all bugfixes and feature changes made in MySQL 4.1 that are not specific to MySQL Cluster, see MySQL 4.1 Server Release Notes.

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.

For additional documentation on MySQL products, including translations of the documentation into other languages, and downloadable versions in variety of formats, including HTML and PDF formats, see the MySQL Documentation Library.

Document generated on: 2016-05-10 (revision: 8859)

Table of Contents

Preface and Legal Notices

Beginning with MySQL 4.1.14, MySQL Cluster changes for MySQL 4.1 Server releases can be found in the MySQL 4.1 Server Release Notes.

This document contains release notes for older releases of MySQL Cluster (before 4.1.14) that use version 4.1 of the NDBCLUSTER storage engine.

Legal Notices

Changes in MySQL Cluster-4.1.13 (15 July 2005)

Note

Starting with MySQL 4.1.14, changes for MySQL Cluster can be found in the MySQL 4.1 Server Release Notes.

Functionality added or changed:

Bugs fixed:

- (Bug #11132) Connections between data nodes and management nodes were not being closed following shutdown of `ndb_mgmd`.
- (Bug #11050) `ndb_mgm> show` printed incorrectly after master data node failure.
- (Bug #10956) More than 7 node restarts with `--initial` caused cluster to fail.
- (Bug #9826) (Bug #10948) Schema change (`DROP TABLE`, `ALTER TABLE`) crashed HPUX and PPC32.
- (Bug #9025) Data nodes failed to restart on 64-bit Solaris.
- (Bug #11166) Insert records were incorrectly applied by `ndb_restore`, thus making restoration from backup inconsistent if the binary log contained inserts.
- (Bug #8918) (Bug #9363) (Bug #10711) (Bug #10058) (Bug #9025) Cluster would time out and crash after first query; setting DataMemory to more than 2GB prevented cluster from starting; calling `ndb_select_count()` crashed the cluster. (64-bit Unix OSes)
- (Bug #10190) When making a backup of a cluster where `NumberOfReplicas` was equal to 1, the backup's metadata was corrupted. (Linux)
- (Bug #9945) `ALTER TABLE` caused server crash. (Linux/390)
- (Bug #11133) A delete operation performed as part of a transaction caused an erroneous result.
- (Bug #10294) Not allowing sufficient parallelism in cluster configuration (for example, `NoOfTransactions` too small) caused `ndb_restore` to fail without generating any error messages.
- (Bug #11290) Setting TransactionInactiveTimeout= 0 did not result in an infinite timeout.

Changes in MySQL Cluster-4.1.12 (13 May 2005)

Functionality added or changed:

Bugs fixed:

- (Bug #10471) Backup can become inconsistent with certain combinations of multiple-row updates
- (Bug #10287) ndb_select_all "delimiter" option non functional
- (Bug #10142) Unhandled resource shortage in UNIQUE index code

- (Bug #10029) crash in ordered index scan after db full
- (Bug #10001) 2 NDB nodes get signal 6 (abort) in DBTC
- (Bug #9969) 4012 - has misleading error message
- (Bug #9960) START BACKUP reports failure albeit succeeding
- (Bug #9924) ABORT BACKUP 1 crashes 4 node cluster
- (Bug #9892) Index activation file during node recovery
- (Bug #9891) Crash in DBACC (line 7004) during commit
- (Bug #9865) SELECT does not function properly
- (Bug #9839) Column with AUTOINC contains -1 Value on node stop
- (Bug #9757) Uncompleted node failure after gracefully stopping node
- (Bug #9749) Transactions causes deadlock in ACC
- (Bug #9724) Node fails to start: Message: File has already been opened
- (Bug #9691) UPDATE fails on attempt to update primary key
- (Bug #9675) Auto-increment not working with INSERT..SELECT and NDB storage
- (Bug #9318) drop database does not drop ndb tables
- (Bug #9280) Memory leak in cluster when dependent sub-queries are used
- (Bug #8928) create table with keys will shutdown the cluster
- Creating a table did not work for a cluster with 6 nodes. (Bug #8928) Databases with 1, 2, 4, 8, ... (2^n nodes) did not have the problem. After a rolling upgrade, restart each node manually by restarting it with the `--initial` option. Otherwise, use dump and restore after an upgrade.

Changes in MySQL Cluster-4.1.11 (01 April 2005)

Functionality added or changed:

Bugs fixed:

- (Bug #9916) DbaccMain.cpp / DBACC (Line: 4876) / Pointer too large
- (Bug #9435) TIMESTAMP columns do not update
- (Bug #9052) Uninitialized data during unique index build, potential cluster crash
- (Bug #8876) Timeout when committing aborted transaction after node failure
- (Bug #8786) ndb_autodiscover, drop index can fail, wait 2 minutes timeout
- (Bug #8853) Transaction aborted after long time during node failure (4012)
- (Bug #8753) Invalid schema object version after dropping index (crash fixed, currently retry required)
- (Bug #8645) Assertion failure with multiple management servers
- (Bug #8557) ndbd does not get same nodeid on restart
- (Bug #8556) corrupt ndb_mgm show printout for certain configurations

- (Bug #8167) cluster shared memory and mysqld signal usage clash

Changes in MySQL Cluster-4.1.10 (12 February 2005)

Bugs fixed:

- (Bug #8284) Out of fragment memory in DBACC
- (Bug #8262) Node crash due to bug in DBLQH
- (Bug #8208) node restart fails on Aix 5.2
- (Bug #8167) cluster shared memory and mysqld signal usage clash
- (Bug #8101) unique index and error 4209 while selecting
- (Bug #8070) (Bug #7937) (Bug #6716) various ndb_restore core dumps on HP-UX
- (Bug #8010) 4006 forces MySQL Node Restart
- (Bug #7928) out of connection objects
- (Bug #7898) mysqld crash with ndb (solaris)
- (Bug #7864) Not possible to have more than 4.5G data memory

Changes in MySQL Cluster-4.1.9 (13 January 2005)

Functionality added or changed:

- New implementation of shared memory transporter.
- Cluster automatically configures shared memory transporter if possible.
- Cluster prioritizes usage of transporters with shared memory and localhost TCP
- Added switches to control the above functions, `ndb-shm` and `ndb-optimized-node-selection`.

Bugs fixed:

- (Bug #7805) config.ini parsing error
- (Bug #7798) Running range scan after alter table in different thread causes node failure
- (Bug #7761) Alter table does not autocommit
- (Bug #7725) Indexed DATETIME Columns Return Random Results
- (Bug #7660) START BACKUP does not increment BACKUP-ID (Big Endian machines)
- (Bug #7593) Cannot Create A Large NDB Data Warehouse
- (Bug #7480) Mysqld crash in ha_ndbcluster using Query Browser
- (Bug #7470) shared memory transporter does not connect
- (Bug #7396) Primary Key not working in NDB Mysql Clustered table (solaris)
- (Bug #7379) ndb restore fails to handle blobs and multiple databases
- (Bug #7346) ndb_restore enters infinite loop
- (Bug #7340) Problem for inserting data into the Text field on utf8

- (Bug #7124) ndb_mgmd is aborted on startup when using SHM connection

Changes in MySQL Cluster-4.1.8 (14 December 2004)

Functionality added or changed:

- Default port for `ndb_mgmd` was changed to 1186 (from 2200) as this port number was officially assigned to MySQL Cluster by IANA.
- New command in `ndb_mgm`, PURGE STALE SESSIONS, as a workaround for cases where nodes fail to allocate a node id even if it is free to use.
- New command in `ndb_mgm`, CONNECT.
- The ndb executables have been changed to make use of the regular MySQL command-line option parsing features. For notes on on changes, see Options Common to MySQL Cluster Programs.
- As bonus of the above you can now specify all command line options in `my.cnf` using the executable names as sections, that is, `[ndbd]`, `[ndb_mgmd]`, `[ndb_mgm]`, `[ndb_restore]`, and so forth.

```
[ndbd]
ndb-connectstring=myhost.domain.com:1234
[ndb_mgm]
ndb-connectstring=myhost.domain.com:1234
```

- Added use of section `[mysql_cluster]` in `my.cnf`. All cluster executables, including mysqld, parse this section. For example, this is a convenient place to put `ndb-connectstring` so that it need be specified only once.
- Added cluster log info events on allocation and deallocation of nodeid's.
- Added cluster log info events on connection refuse as a result of version mismatch.
- Extended connectstring syntax to allow for leaving the port number out. For example, `ndb-connectstring|connect-string=myhost1,myhost2,myhost3` is a valid connectstring and connect occurs on default port 1186.
- Clear text ndb error messages provided also for error codes that are mapped to corresponding mysql error codes, by executing `SHOW WARNINGS` after an error has occurred which relates to the ndb storage engine.
- Significant performance improvements done for read performance, especially for blobs.
- Added some variables for performance tuning, `ndb_force_send` and `ndb_use_exact_count`. Do `show variables like 'ndb%';` in mysql client for listing. Use `set` command to alter variables.
- Added variables to set some options, `ndb_use_transactions` and `ndb_autoincrement_prefetch_sz`.

Bugs fixed:

- (Bug #7303) ndb_mgm: Trying to set CLUSTERLOG for a specific node id core dumps
- (Bug #7193) start backup gives false error printout
- (Bug #7153) Cluster nodes do not report error on endianness mismatch
- (Bug #7152) ndb_mgmd segmentation fault on incorrect HostName in configuration
- (Bug #7104) clusterlog filtering and level setting broken
- (Bug #6995) ndb_recover on varchar fields results in changing case of data

- (Bug #6919) all status only shows 2 nodes on a 8-node cluster
- (Bug #6871) DBD execute failed: Got error 897 'Unknown error code' from ndbcluster
- (Bug #6794) Wrong outcome of update operation of ndb table
- (Bug #6791) Segmentation fault when config.ini is not correctly set
- (Bug #6775) failure in acc when running many mysql clients
- (Bug #6696) ndb_mgm command-line options inconsistent with behavior
- (Bug #6684) ndb_restore doesn't give error messages if improper command given
- (Bug #6677) ndb_mgm can crash on "ALL CLUSTERLOG"
- (Bug #6538) Error code returned when select max() on empty table with index
- (Bug #6451) failing create table givers "ghost" tables which are impossible to remove
- (Bug #6435) strange behavior of left join
- (Bug #6426) update with long pk fails
- (Bug #6398) update of primary key fails
- (Bug #6354) mysql does not complain about --ndbcluster option when NDB is not compiled in
- (Bug #6331) INSERT IGNORE .. SELECT breaks subsequent inserts
- (Bug #6288) cluster nodes crash on data import
- (Bug #6031) To drop database you have to execute DROP DATABASE command twice
- (Bug #6020) LOCK TABLE + delete returns error 208
- (Bug #6018) REPLACE does not work for BLOBs + NDB
- (Bug #6016) Strange crash with blobs + different DATABASES
- (Bug #5973) ndb table belonging to different database shows up in show tables
- (Bug #5872) ALTER TABLE with blob from ndb table to myisam fails
- (Bug #5844) Failing mysql-test-run leaves stray NDB processes behind
- (Bug #5824) HELP text messed up in ndb_mgm
- (Bug #5786) Duplicate key error after restore
- (Bug #5785) lock timeout during concurrent update
- (Bug #5782) Unknown error when using LIMIT with ndb table
- (Bug #5756) RESTART node from ndb_mgm fails
- A few more not reported bugs fixed

Changes in MySQL Cluster-4.1.7 (23 October 2004)

Functionality added or changed:

- Optimization 1: Improved performance on index scans. Measured 30% performance increase on query which do large amounts of index scans.

- Optimization 2: Improved performance on primary key lookups. Around double performance for autocommitted primary key lookups.
- Optimization 3: Improved performance when using blobs by avoiding usage of exclusive locks for blobs.

Bugs fixed:

- A few bugs fixed.

Changes in MySQL Cluster-4.1.6 (10 October 2004)

Functionality added or changed:

- Limited character set support for storage engine NDBCLUSTER:

Character Set	Collation
big5	big5_chinese_ci
	big5_bin
binary	binary
euckr	euckr_korean_ci
	euckr_bin
gb2312	gb2312_chinese_ci
	gb2312_bin
gbk	gbk_chinese_ci
	gbk_bin
latin1	latin1_swedish_ci
	latin1_bin
sjis	sjis_japanese_ci
	sjis_bin
tis620	tis620_bin
ucs2	ucs2_general_ci
	ucs2_bin
ujis	ujis_japanese_ci
	ujis_bin
utf8	utf8_general_ci
	utf8_bin

- The SCI Transporter has been brought up-to-date with all changes and now works and has been documented as well.
- Optimizations when several clients to a MySQL Server access ndb tables.
- Added more checks and warnings for erroneous and inappropriate cluster configurations.
- `SHOW TABLES` now directly shows ndb tables created on a different MySQL server, that is, without a prior table access.
- Enhanced support for starting MySQL Server independently of ndbd and ndb_mgmd.
- Clear text ndb error messages provided by executing `SHOW WARNINGS` after an error has occurred which relates to the ndb storage engine.

Bugs fixed:

- Quite a few bugs fixed.

Changes in MySQL Cluster-4.1.5 (16 September 2004)

Functionality added or changed:

- Many queries in MySQL Cluster are executed as range scans or full table scans. All queries that do not use a unique hash index or the primary hash index use this access method. In a distributed system it is crucial that batching is properly performed.

  In previous versions, the batch size was fixed to 16 per data node. In this version it is configurable per MySQL Server. So for queries using lots of large scans it is appropriate to set this parameter rather large and for queries using many small scans only fetching a small amount of records it is appropriate to set it low.

  The performance of queries can easily change as much as 40% based on how this variable is set.

  In future versions more logic will be implemented for assessing the batch size on a per-query basis. Thus, the semantics of the new configuration variable `ScanBatchSize` are likely to change.

- The fixed size overhead of the ndbd process has been greatly decreased. This is also true for the overhead per operation record as well as overhead per table and index.

  A number of new configuration variables have been introduced to enable configuration of system buffers. Configuration variables for specifying the numbers of tables, unique hash indexes, and ordered indexes have also been introduced.

  New configuration variables: `MaxNoOfOrderedIndexes`, `MaxNoOfUniqueHashIndexes`

  Configuration variables no longer used: `MaxNoOfIndexes` (split into the two above).

- In previous versions `ALTER TABLE`, `TRUNCATE TABLE`, and `LOAD DATA` were performed as one big transaction. In this version, all of these statements are automatically separated into several distinct transactions.

  This removes the limitation that one could not change very large tables due to the `MaxNoOfConcurrentOperations` parameter.

- MySQL CLuster's online backup feature now backs up indexes so that both data and indexes are restored.

- In previous versions it was not possible to use `NULL` in indexes. This is now possible for all supported index types.

- Much work has been put onto making `AUTO_INCREMENT` features work as for other table handlers. Autoincrements as a partial key is still only supported by `MyISAM`.

- In earlier versions, `mysqld` would crash if the cluster wasn't started with the `--ndbcluster` option. Now `mysqld` handles cluster crashes and starts without crashing.

- The `-i` option for initial startup of `ndbd` has been removed. Initial startup still can be specified by using the `--initial` option. The reason for this is to ensure that it is clear what takes place when using `--initial`: this option completely removes all data from the disk and should only be used at initial start, in certain software upgrade cases, and in some cases as a workaround when nodes cannot be restarted successfully.

- The management client (`ndb_mgm`) now has additional commands and more information is printed for some commands such as `show`.

- In previous versions, the files were called `ndb_0..` when it wasn't possible to allocate a node ID when starting the node. To ensure that files are not so easily overwritten, these files are now named `ndb_pid..`, where pid is the process ID assigned by the OS.

- The default parameters have changed for `ndb_mgmd` and `ndbd`. In particular, they are now started as daemons by default. The `-n` option has been removed since it could cause confusion as to its meaning (nostart or nodaemon).
- In the configuration file, you can now use `[NDBD]` as an alias for `[DB]`, `[MYSQLD]` as an alias for `[API]`, and `[NDB_MGMD]` as an alias for `[MGM]`.

Note

In fact, `[NDBD]`, `[MYSQLD]`, and `[NDB_MGMD]` are now the preferred designations, although the older ones will continue to be supported for some time to come to maintain backward compatibility.

- Many more checks for consistency in configuration have been introduced to to provide quicker feedback on configuration errors.
- In the connect string, it is now possible to use both "`;`" and "`,`" as the separator between entries. Thus, "nodeid=2,host=localhost:2200" is equivalent to "nodeid=2;host=localhost:2200".

  In the configuration file, it is also possible to use "`:`" or "`=`" for assignment values. For example, `MaxNoOfOrderedIndexes : 128` and `MaxNoOfOrderedIndexes = 128` are equivalent expressions.
- The configuration variable names are now case insensitive, so `MaxNoOfOrderedIndexes: 128` is equivalent to `MAXNOOFORDEREDINDEXES = 128`.
- It is possible now to set the backup directory separately from the `FileSystemPath` by using the `BackupDir` configuration variable.

  Log files and trace files can now be placed in any directory by setting the `DataDir` configuration variable.

  `FileSystemPath` is no longer mandatory and defaults to `DataDir`.
- Queries involving tables from different databases are now supported.
- It is now possible to update the primary key.
- The performance of ordered indexes has been greatly improved, particularly the maintenance of indexes on updates, inserts and deletes.

Bugs fixed:

- Quite a few bugs fixed.

Changes in MySQL Cluster-4.1.4 (31 August 2004)

Functionality added or changed:

- The names of the log files and trace files created by the `ndbd` and `ndb_mgmd` processes have changed.
- Support for the many `BLOB` data types was introduced in this version.

Bugs fixed:

- Quite a few bugs were fixed in the 4.1.4 release.

Changes in MySQL Cluster-4.1.3 (28 June 2004)

Functionality added or changed:

- This was the first MySQL Cluster release so all functionality was new.

Bugs fixed:

- Various bugs fixed in the development process leading up to 4.1.3.