
MySQL Cluster Manager 1.4 Release Notes

Abstract

This document contains information about changes in successive versions of MySQL Cluster Manager 1.4, starting with the most recent release (MySQL Cluster Manager 1.4.0). Changes in previous MySQL Cluster Manager releases can be found afterwards, ordered from newest to oldest.

For additional MySQL Cluster Manager 1.4 documentation, see <http://dev.mysql.com/doc/mysql-cluster-manager/1.4/en/>.

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: 2018-04-21 (revision: 14782)

Table of Contents

Preface and Legal Notices	1
Changes in MySQL Cluster Manager 1.4.5 (2017-01-25)	3
Changes in MySQL Cluster Manager 1.4.4 (2017-10-05)	4
Changes in MySQL Cluster Manager 1.4.3 (2017-07-10)	5
Changes in MySQL Cluster Manager 1.4.2 (2017-03-07)	8
Changes in MySQL Cluster Manager 1.4.1 (2016-11-24)	9
Changes in MySQL Cluster Manager 1.4.0 (2015-12-07)	15
Index	17

Preface and Legal Notices

This document contains information about changes in successive versions of MySQL Cluster Manager 1.4, starting with the most recent release (MySQL Cluster Manager 1.4.0). Changes in previous MySQL Cluster Manager releases can be found afterwards, ordered from newest to oldest.

Legal Notices

Copyright © 2009, 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 Cluster Manager 1.4.5 (2017-01-25)

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

Functionality Added or Changed

- **Agent:** When the `rotate log` command is run, the new file name for the old `mcmd` log now has the timestamp put right before the file extension and preceded by an underscore; for example, `mcmd.log` becomes `mcmd_2017-07-15T22-45-28.log` after log rotation. (Bug #26696293)
- **Client:** When an **Internal Error (Error 1003)** is returned, the error message now indicates the host on which the error occurs. (Bug #27082475)

Bugs Fixed

- **Agent:** The reported status of a node was sometimes ignored (for example, a node reported to have stopped was next considered running). This was because log sequence numbers (LSNs) on nodes (by which `mcmd` tells whether the status report is out-of-date) were wrongly assigned sometimes. The fix corrects the LSN mechanism, so that status reports can be properly handled. (Bug #27338269)
- **Agent:** A `stop cluster` command timed out while waiting for a failed management node to respond. With this fix, the command errors out immediately in the situation. (Bug #27253476)
- **Agent:** `mcmd` hung on a `restore cluster` command when it failed to get the process handle of a node that was newly spawned and then failed. This fix adds proper error handling for the situation. (Bug #27225001)
- **Agent:** After bootstrapping an NDB Cluster with the `start cluster --bootstrap` command, connecting to the cluster with a `mysql` client using the command suggested by MySQL Cluster Manager failed, because the command referred to the local host by its host name, which could be unresolvable in some situations. With this fix, the IP address `127.0.0.1` is used instead in the command. (Bug #27169533)
- **Agent:** An `mcmd` agent exited unexpectedly when a `ndbmta` node failed during a rolling restart, because the failure disallowed `mcmd` to identify the node's mirror. With this fix, `mcmd` throws the proper error in the situation and then continues with its operations. (Bug #27105336)
- **Agent:** A `start cluster` command failed when there was an unmanaged node running on one of the cluster hosts. (Bug #27060818)
- **Agent:** A stopping of a single node sometimes caused a large number of repeated exit reports to be generated to all `mcmd` agents. (Bug #27049206)
- **Agent:** `mcmd` reported success for a `restore cluster` command even if it had failed. (Bug #25117413)
- **Agent:** On Windows platforms, `mcmd` gave an excessive number of warnings to its log when it "Could not get process handle" for a cluster node. (Bug #23570996)

- **Agent:** When a `stop cluster` command was running, sometimes a data node was not shut down completely and the command failed with an `ERROR 7006` for `nodegroup failure of the data nodes`. This fix stops the unnecessary error from being thrown in the situation and lets the command continue. (Bug #21508316)
- **Agent:** The `import cluster` command failed when the cluster to be imported used connection pooling for its `mysqld` nodes. (Bug #19357267)
- **Agent:** On Windows platforms, after a cluster reconfiguration, when a management node failed to be restarted, it was still reported by `mcmd` as `starting`. It was because on Windows platforms, `mcmd` continued its attempt to connect to the stopped `ndb_mgmd` node in the situation until timeout. With this fix, status of the `ndb_mgmd` nodes are first checked before attempts were made to connect to them. (Bug #14255038)
- **Agent:** When one of the NDB Cluster's `ndb_mgmd` nodes was already down, using a `stop cluster --force` to shut down the whole cluster was going to take a long time. It was because the `stop process` command's timeout, after being triggered, was not made known to all relevant processes on the cluster. With this fix, all `mcmd` agents are now informed of the timeout, so they will just kill processes that are still running after the timeout period. (Bug #13778840)
- **Agent:** If a data node was added but could not be started at first, nodegroup creation failed even after the node was started successfully at the end. (Bug #13261667)
- **Client:** When the `show status` command reported that a node had failed, if the `mcmd` agent that reported the failure also failed afterwards and the user then tried to restart the node, an `ERROR 7007` was returned, saying wrongly that the node was already running. With this fix, a proper `Error 1003` is returned with the message "Agent on host `hostname` is unavailable." (Bug #27082586)
- **Client:** Using the `--removedirs` option with the `delete cluster` command did not remove a data node's file system (the `ndb_nodeID_fs` directory) if it was located outside of the node's data directory. (Bug #25591980)
- **Client:** While the `set` command cannot be used to set attributes for an unmanaged node that does not appear in the `config.ini` file (for example, the `--port` option for an unmanaged `mysqld` node), such attempts were not rejected by the `mcm` client. With this fix, `mcm` throws an error in the situation. (Bug #25512477)
- **Client:** In some situations, after a new process was added, the `start process --added` command failed. It was because the added process was misidentified with a process that had stopped earlier. With this fix, the handling of stopped and added processes has been improved to prevent such cases of misidentification. (Bug #25160942)
- **Client:** A `stop cluster` command failed if all data nodes of the cluster were in the `failed` status when the command was issued. (Bug #24931681)

Changes in MySQL Cluster Manager 1.4.4 (2017-10-05)

This section documents all changes and bug fixes that have been applied in MySQL Cluster Manager 1.4.4 since the release of MySQL Cluster Manager version 1.4.3.

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

Functionality Added or Changed

- **Agent:** On Windows platforms, when the `collect logs` command was run and the resulting file path for any of the collected files was longer than 255 characters, the command failed with `ERROR 101`

(00MGR) Access denied. A new type of error, `PATH_TOO_LONG (106) Path max length exceeded for file filename`, has been added and is now returned at the situation. (Bug #26619163)

Bugs Fixed

- **Agent:** On Solaris platforms, the `collect logs` command created lots of zombie threads after it finished. It was because a number of thread joins were missing in the code, which have been added by this fix. (Bug #26665347)
- **Agent:** For a large NDB Cluster site, the `collect logs` command frequently failed due to errors with the network connections during file transfers. This was because some error codes related to the network connections were not properly handled, and this fix corrects that. (Bug #26642616)
- **Agent:** The `import config` command sometimes failed, as the `mcmd` agent updated its configuration store before all the values to be imported arrived. This fix makes `mcmd` wait for all the values before starting the updates. (Bug #26610874)
- **Agent:** The `collect logs` command reported success even when an agent failed and the file transfer from it was incomplete. With this fix, when connection to an agent fails, the command allows running file transfers to complete, and then throws an error. (Bug #26594948)
- **Agent:** The `collect logs` command sometimes failed with `ERROR 1003 Internal error: No clients connected`, even though clients were connecting and files were being transferred. That was due to a faulty mechanism for tracking client connections, which has been corrected by this fix. (Bug #26575768)
- **Agent:** The `mcm` client hung when running the `collect logs` command if the host names for the cluster nodes could not be resolved, even if the cluster was actually defined using IP addresses. (Bug #26564176)
- **Agent:** A `start cluster` command timed out when its timeout value, being scaled for handling a large `DataMemory` size for a cluster, was accidentally reset. This fix makes sure the timeout for the execution step of a `start cluster` command is never lowered while the data nodes are starting. (Bug #26543383)
- **Agent:** After a node had failed and an `update process` command on it also failed, a `stop cluster` or `stop process` command that followed ran into an assertion error. It was because the node, automatically restarted after its failure, had its PID file modified by the `update process` command that was still running. This fix adds a check on the status of the target process for the `update process` command to prevent the restarted node's PID file from being modified. (Bug #26388603)
- **Agent:** The `collect logs` command hung for 30 seconds after completing all file transfers before returning the success message. (Bug #26362910)
- **Agent:** When a management node's status went from `starting` to `stopped`, it was not reported by its `mcmd` agent, causing commands like `delete cluster` to fail because the management node was taken to be still running. (Bug #26288009)

Changes in MySQL Cluster Manager 1.4.3 (2017-07-10)

This section documents all changes and bug fixes that have been applied in MySQL Cluster Manager 1.4.3 since the release of MySQL Cluster Manager version 1.4.2.

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

Functionality Added or Changed

- **Agent:** CPU usage during idle time for the `mcmd` agents has been significantly reduced. (Bug #26227736)
- **Agent:** A new error code, `Error 7030`, has been created for failed `ndb_mgmd` commands and `mysqld` queries. (Bug #26160968)
- **Agent:** Added support for the `--skip-networking` option for `mysqld` nodes, allowing `mysqld` nodes of a managed cluster to communicate with client applications using named pipes or shared memory on Windows platforms, and socket files on Unix-like platforms.

Notice that, however, communication between `mcmd` agents and `mcm` clients using named pipes, shared memory, or socket files remain unsupported. (Bug #25992390, Bug #25974499)

- **Client:** The `start cluster --initial` command now reinitializes the SQL nodes (if their data directories are empty) as well as the data nodes of an NDB Cluster. A new option, `--skip-init`, has been introduced, for specifying a comma-separated list of the SQL nodes for which reinitialization is to be skipped. (Bug #25856285, Bug #85713)
- **Client:** Checksum verification has been added for all cluster reconfiguration plans created by the `mcmd` agents. Checksums for plans created locally are shared among all agents, and when the checksums do not match, the reconfiguration is aborted. This prevents agents from executing different plans. (Bug #23225839)
- Files have been removed from the MySQL Cluster Manager + NDB Cluster bundled package, in order to reduce the package size significantly. (Bug #25916635)

Bugs Fixed

- **Agent:** When the `list nextnodeids` command was run against a cluster with the maximum number of nodes allowed, the `mcmd` agent quit unexpectedly. With this fix, the situation is properly handled. (Bug #26286531)
- **Agent:** For a cluster with `NoOfReplicas=1`, trying to stop a data node with the `stop process` command would cause the agent to quit unexpectedly. (Bug #26259780)
- **Agent:** When a data node was killed by an arbitrator in a situation of network partitioning, an `mcmd` failed to handle the exit report from the node and quit unexpectedly. It was due to a mishandling of the `nodegroup` information, which this fix corrects. (Bug #26192412)
- **Agent:** A cluster could not be started if a relative path had been used for the `--manager-directory` option to set the location of the agent repository. (Bug #26172299)
- **Agent:** When executing a user command, the `mcmd` agent could hang if the expected reply from another agent never arrived. This fix improves the timeout handling to avoid such hangs. (Bug #26168339)
- **Agent:** While running the `import config` command, the `mcmd` agents that were present during the earlier dryrun for the import would become silent and then unavailable. This was due to some hostname resolution issues, which has been addressed by this fix. (Bug #26089906)
- **Agent:** A `collect log` command sometimes failed at the middle with an `ERROR 1003 Internal error: No clients connected`. It was because the `mcmd` agent reset the copy completion marker prematurely; the behavior has been stopped by this fix. (Bug #26086958)
- **Agent:** When the `mcmd` agents' clocks ran out of sync due to time drifts on virtual machines running Windows operations systems and then the clocks ran in sync again, communications among the

agents failed. This fix prevents the problem by making the agents use a monotonic timer for their communication. (Bug #26084090)

- **Agent:** The dropping or recreating of a node group that took place when adding data nodes could sometimes fail with an assertion error (“Polled nodegroup info is inconsistent”). This fix relaxes the assertion, which allows the node group reconfiguration to be completed. (Bug #26051753)

References: See also: Bug #20104357.

- **Agent:** During an execution of a `set` command, if no `mysqld` node is available for querying cluster information, an `mcmd` agent timed out while waiting for the “prepared” message from another agent, even after the message was already sent. This was due to the fact that the two agents had inconsistent plans of execution for the `set` command. This fix prevents the inconsistency. (Bug #26021616)

References: This issue is a regression of: Bug #14230789, Bug #23148061.

- **Agent:** A `backup cluster --waitcompleted` timed out sometimes right before the backup was completed when there were a lot of tables to be backed up. This was because the logical backup for the tables' metadata was taking too long in that case. With this fix, the `mcmd` agent is now sent progress reports of the logical backup, and the backup does not time out unless no more progress reports are received. (Bug #26000482)
- **Agent:** When a `set` command involved a restart of data nodes of a cluster but one of the data nodes had been stopped, the `set` command failed with a timeout. With this fix, the `set` command is carried out successfully with a rolling restart for the data nodes. (Bug #25869325)
- **Agent:** If a `mysqld` node was configured with the `--skip-name-resolve` option, attempts for `mcmd` to connect to the `mysqld` node would fail with the error message `Host '127.0.0.1' is not allowed to connect to this MySQL server`. This was because the MySQL account used by `mcmd` had `127.0.0.1` as its host name, which is not allowed when the `--skip-name-resolve` option is used with the `mysqld` node. This fix corrects the account host name to `localhost`. (Bug #25831764, Bug #85620)
- **Agent:** When a host and its `mcmd` agent were restarted, `mcmd` might fail to restart a management or `mysqld` node on the host, and the `show status` command continuously returned the status of the node to be `unknown`. (Bug #25822822)
- **Agent:** When an `mcmd` agent was in the process of shutting down, a user command issued then might cause the agent to quit unexpectedly. With this fix, an error message “Agent is shutting down” is returned, and the agent continues with its shutdown. (Bug #25055338)
- **Agent:** When a `set` command involved a restart of data nodes of a cluster but one of the data nodes was in the `failed` state, `mcmd` restarted the data node, and then restarted it once more as part of a rolling restart, which was unnecessary. This fix eliminates the second restart. (Bug #23586651)
- **Client:** After an agent was started and a few commands had been executed from the `mcm` client, the `show settings` command started returning the wrong value for the `--log-level` option. (Bug #26189795)
- **Client:** Trying to set the `mysqld` node option `--validate-password` resulted in an error complaining that the parameter did not exist, even if the `Password Validation Plugin` had already been installed on the `mysqld` node. It was due to some errors with the plug-in activation for `mysqld` nodes, which have now been corrected. (Bug #25797125)

Changes in MySQL Cluster Manager 1.4.2 (2017-03-07)

This section documents all changes and bug fixes that have been applied in MySQL Cluster Manager 1.4.2 since the release of MySQL Cluster Manager version 1.4.1.

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

Functionality Added or Changed

- **Packaging:** MySQL Cluster Manager is now also shipped in RPM packages for installation on Oracle Linux 7 and Red Hat Enterprise Linux 7. (Bug #25368708)
- **Agent:** To allow easy detection of an incomplete agent backup, an empty file named INCOMPLETE is created in the folder in which the backup is created when the `backup agents` command begins, and is deleted after the backup is finished. The continuous existence of the file after the backup process is over indicates that the backup is incomplete. (Bug #25126866)
- **Agent:** A new `update process` command updates a cluster process' status when it is no longer reflected correctly in the output of the `show status --process`. It imports the process into the control of the `mcmd` agent again. See description of the `update process` command for details. (Bug #25098588)
- **Agent:** MySQL Cluster Manager can now recover automatically a failed `mysqld` node, as long as the data directory of the node is empty when recovery is attempted; if that is not the case, after cleaning up the data directory manually, users can now manually run `start process --initial` to rebuild the `mysqld` node's data directory. (Bug #18415446)
- **Agent:** The `show status` command now reports progress when the new `--progress` or `--progressbar` option is used.
- **Agent:** A new command, `update process`, imports a process back into the control of `mcmd` after it has lost track of the process's status due to different reasons (for example, it has been restarted manually outside of MySQL Cluster Manager). For more details, see the description for the command.

Bugs Fixed

- **Agent:** `mcmd` sometimes quit unexpectedly when running on SLES 11.4 with ReiserFS due to sporadic occurrences of checksum errors with the MySQL Cluster Manager data repository. (Bug #25596300)
- **Agent:** When a custom `FileSystemPath` value was used for a data node, the `list backups` and `restore cluster` commands failed, as the backup directory could not be found. (Bug #25549903)
- **Agent:** In some situations, a certain `mcmd` agent took too long to process event messages that a synchronization timeout occurred among the agents. This was because the agent went into a mutex contention for file access, which this fix removes. (Bug #25462861)
- **Agent:** The `collect logs` command reported success even if file transfers were incomplete. This fix adds checks for file transfer completion and reports any errors. (Bug #25436057)
- **Agent:** An `ndbmtid` node sometimes (for example, at a rolling restart of the cluster) sent out a large amount of event messages, and it might take too long for an `mcmd` agent to process them that the agent lagged behind on its readiness for the next command, resulting in a synchronization timeout among the `mcmd` agents. This fix drastically reduced the amount of event messages sent out about an `ndbmtid` node, thus reducing the chance of a synchronization timeout under the situation. (Bug #25358050)

- **Agent:** A management node failure might trigger `mcmd` to quit unexpectedly on Windows platforms. (Bug #25336594)
- **Agent:** Multiple errors thrown by the `backup agents`, `rotate log`, and `change log-level` commands could potentially overwrite each other, causing a lost of error information. (Bug #25134452)
- **Agent:** The `collect logs` command hung when TCP connections could not be established between the agent that initiated the command and the other agents. This fix makes the command timeout after the situation persists for more than 30s. Also, a new `mcmd` option, `--copy-port`, has been added, by which users can specify the TCP port number to be used for log copying. (Bug #25064313)
- **Agent:** The `.mcm` file created by the `import config --dryrun` command sometimes have certain configuration settings missing from it. (Bug #24962848)
- **Agent:** A `restore cluster` command would fail if MySQL Cluster Manager did not have write access to the `BackupDataDir` of each data node. The unnecessary requirement has now been removed. (Bug #24763936)
- **Agent:** If a `stop cluster` or a `stop process` command had failed, a restart on some of the processes might fail with the complaint from `mcmd` that those processes were already stopped, even if they were actually running. That also made it impossible to reconfigure those processes when `StopOnError` was true. This happened because the failed `stop` command had left those processes' metadata in an incorrect state. With this fix, the process restart is allowed despite the value of `StopOnError`. (Bug #24712504)
- **Agent:** Hostnames referenced in the error messages returned by `mcmd` were always in lower case. With this fix, the hostname is always referred to as it is; moreover, `mcmd` now always refers to a hostname or the IP address used in creating the cluster. (Bug #21375132)
- **Agent:** A `restore cluster` command hung, when an `mcmd` agent failed and the other agents kept waiting to receive messages from it. With the fix, the other agents detect the failure and return an error to the user. (Bug #16907088)
- **Agent:** When a cluster was being started, if a data node failed shortly after it was started and `mcmd` was still in the process of starting an SQL node, even if the SQL node was started successfully at the end, `mcmd` might forever lose connection to the SQL node. It happened when the user `mcmd` required for the `mcmd` agent did not get created on the SQL node. With this fix, the user `mcmd` is always created on the SQL node despite a failure of the `start cluster` command. (Bug #13436550)

Changes in MySQL Cluster Manager 1.4.1 (2016-11-24)

This section documents all changes and bug fixes that have been applied in MySQL Cluster Manager 1.4.1 since the release of MySQL Cluster Manager version 1.4.0.

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

Functionality Added or Changed

- **Microsoft Windows:** A 64-bit Windows installer package for MySQL Cluster Manager is now available.
- **Agent:** Quotes used in specifying an `mcmd option` were interpreted literally as parts of the option value. With this fix, `mcmd` now handles the quotes properly for the following options by stripping the quotes before interpreting the option values:
 - `--basedir`

- `--pid-file`
- `--plugin-dir`
- `--log-file`
- `--manager-username`
- `--manager-password`
- `--manager-directory`

(Bug #24529142)

- **Agent:** When an `mcmd` agent failed to shutdown a `mysqld` node, it kept on retrying until the process timed out, at which point the error code from `mysqld` was returned to the client. With this fix, the error is returned immediately after the first shutdown attempt has failed. (Bug #24418005)
- **Agent:** After a `start process --added` command was executed, if a `show status --operational` command was then run, in the `Command` column of the output, the ID for only one of the added nodes was shown even if multiple nodes had been added and started. With this fix, nodeids are no longer shown in the command output in that case, to avoid any possible confusions. (Bug #24372861)
- **Agent:** The `restore cluster` command now supports a new option, `--promote-attributes`, which allows attributes to be promoted when MySQL Cluster Manager restores data from a backup. (Bug #24295481)
- **Agent:** The `collect logs` command now also collects trace files, to provide more information in the event of a cluster failure. (Bug #23746394)
- **Agent:** Information that is unrelated to debugging is no longer included into the agent log when `--log-level` is set to `debug`, but is included when `--log-level` is set to `info`. (Bug #23716969)
- **Agent:** The capability for MySQL Cluster Manager to set the cluster parameter `ndbinfo_table_prefix` has been removed, as setting it might cause a timeout when `mcmd` tries to start a `mysqld` node. (Bug #23632067)
- **Agent:** The `import cluster` command now supports the import of cluster nodes that have been started with options specified on the command line using underscores (`_`) instead of dashes (`-`) in the option names (for example, `--ndb-connectstring` instead of `--ndb-connectstring`). (Bug #23535372)
- **Agent:** The `--initial` option can now be used with the `start cluster` command for starting clusters in the “created” status. (Bug #23138442)
- **Agent:** The timer in `mcmd` is now monotonic, so interval measures are no longer affected by clock changes on the host. (Bug #22699245)
- **Agent:** The version number for the `mcmd` agent and the name of the host it is running on are now always recorded in the agent log (`mcmd.log`) at startups or log rotations, irrespective of the `--log-level` setting. (Bug #22616530)
- **Agent:** The `backup cluster` command now creates a logical backup for the metadata of the NDB tables in a MySQL Cluster, allowing more flexibility for cluster reconfiguration during a restore. See the description for `backup cluster` for more details. (Bug #21200829)

- **Agent:** The `--backupid` option can now be used with the `list backups` command to specify the ID of the backup to be listed.
- **Client:** Execution of the `.mcm` file created by the `import config --dryrun` command failed when a relative file path was given for a file-path attribute to be imported for a `mysqld` node. The `mcm` client now accepts relative file paths for those attributes, except for the `socket` value, and for any paths for a directory value (for example, `plugin_dir`). (Bug #18650848)
- **Client:** Creating a single-host site with `localhost` as the argument for `--hosts` now results in a site that cannot be scaled up by the `add hosts` command; a warning for that is given by the `mcm` client at the creation of the site. See the description for the `create site` command for details. (Bug #18389510)
- MySQL Cluster Manager now supports MySQL Cluster 7.5. At the same time, support for MySQL Cluster 7.2 has been deprecated. (Bug #24940329)

Bugs Fixed

- **Agent:** When the `backup agents` command failed to copy some files, it finished without erroring out, but caused the subsequent `mcmd` agent restart to fail. This fix adds proper error-handling function to the `backup agents` command. (Bug #25057056)
- **Agent:** When a `start cluster` or `stop cluster` command was run, if it took more than two minutes for any data node to be started or stopped, the `mcmd` agent threw a timeout for the command, and sometimes did that more than once, even though the data node was still in a normal process of being started or stopped. It was because the timeout extension needed was never granted within two minutes. With this fix, a timeout extension is granted at a much earlier time, and the amount of the extension is calculated based on the cluster's `DataMemory` size. (Bug #24749459)
- **Agent:** When `mcmd` failed to create the users it needs on a `mysqld` node, it just kept retrying until a timeout was reached, and without notifying the `mcm` client. With this fix, a failure is declared immediately to the client in the situation, and there are no more retries. (Bug #24511041)
- **Agent:** After a failure by `mcmd` to configure a cluster, the progress message from the process caused the next attempt to configure the cluster to fail. With this fix, the progress message from the last attempt only causes a warning to be issued to the agent log, and is then ignored. (Bug #24499097)
- **Agent:** Attempts to upgrade MySQL Cluster Manager to version 1.4.0 by starting the 1.4.0 executables on pre-1.4 data failed, as the automatic upgrades for backup configurations were not successful. (Bug #24433400)
- **Agent:** `mcmd` exited unexpectedly after failing to parse correctly a message of an unexpected format from a `mysqld` node. With this fix, sanity check is performed on the received message first, in order to avoid parsing the wrong information out of it. (Bug #24430610)
- **Agent:** When attempting to set a global system variable on a `mysqld` node, the `mcmd` agent always returned a success message, even if it failed to contact the `mysqld` node actually. (Bug #24417856)
- **Agent:** When running the `list backups` command, if one of the `mcmd` agents ran into permission errors on some of the backups' file paths, the returned results for the command became incomplete. With this fix, complete results are returned in those situations. (Bug #24414682)
- **Agent:** When the `mcm` client and the `mcmd` agent it connected to were on the same host whose loopback address was not 127.0.0.1, if the host's name or loopback address was not explicitly contained in the host list in the `create site` command, attempts to restart the cluster later would fail with the complaint that the host name for the agent could not be resolved in the host list. With this fix, a warning is given

in the situation when the `create site` command is issued, asking the user to include the proper host name. (Bug #24411504)

- **Agent:** `mcmd` sometimes failed to report errors for its executions of SQL statements to the `mcmd` log. This fix tries to ensure that those errors are captured. (Bug #24375344)
- **Agent:** When an `add package` command failed due to an error in loading the cluster configuration, `mcmd` exited unexpectedly on Windows platforms, and falsely reported that the package was added successfully on Unix-like platforms. (Bug #24361901)
- **Agent:** When `mcmd` failed to write to the configuration file `config.ini` or `my.cnf`, `mcmd` either quit unexpectedly (for most Unix-like platforms), or returned a success message without actually writing to the file (on Solaris). (Bug #23632067)
- **Agent:** When a configuration parameter was imported from a `mysqld` node using the `import config` command, the value got truncated when it was a quoted value containing a hash sign and was followed by a comment when listed inside the configuration file. (Bug #23591849)
- **Agent:** On Windows platforms, the value for the system variable `wait_timeout` for `mysqld` nodes cannot be made larger than “2147483” using the `set` command. (Bug #23503256)
- **Agent:** On Windows platforms, the `import cluster` command failed when a 32-bit `mcmd` was working a 64-bit MySQL Cluster, or vice versa. (Bug #23503256)
- **Agent:** The `import config --dryrun` command produced `set` commands of wrong syntax for instance-level TCP configuration settings specified inside the `config.ini` file. (Bug #23341146)
- **Agent:** When `NoOfReplicas` was “4” for a cluster and only one fourth of all the data nodes were available, `mcmd` reported that the cluster was “non-operational” when the `show status --cluster` command was run. (Bug #23330032)
- **Agent:** `mcmd` ignored status update from an `ndbapi` (or unmanaged `mysqld`) node when it reported an LSN lower than one already reported by some other agent. This was due to an unnecessary check on the reported LSNs, which has been removed by this fix. (Bug #23320387)
- **Agent:** When working with MySQL Cluster 7.5, resetting the `Datadir` parameter for a `mysqld` node using a `reset` command caused an error when the default data directory had contents inside. (Bug #23283577)
- **Agent:** The command `start process --added` failed with the error “Unable to create nodegroup ...” when the parameter `NoOfReplicas` was set to “3” or “4.” (Bug #23257723)
- **Agent:** Some `ndbmt` processes were being excluded from rolling restarts when the parameter `NoOfReplicas` was set to “3” or “4.” (Bug #23251630)
- **Agent:** With the parameter `NoOfReplicas` set to “3” or “4,” even when there was still one mirror node left in a nodegroup for a `ndbmt` process, the `stop process` command issued for the `ndbmt` node was rejected by `mcmd`. (Bug #23250053)
- **Agent:** When a quorum of majority of `mcmd` agents no longer existed (due to, for example, a network failure), an `mcmd` agent reported wrong statuses of failed processes even when those processes were local to its own host and were accessible by the agent. (Bug #23222658)
- **Agent:** When network connection to a host with a data node running was lost and then reestablished, the `show status` command reported falsely that the data node was running again even though it had actually stopped. This fix makes sure the current status of the data node is properly reflected. (Bug #23220981)

- **Agent:** When there was a failure in setting the value of the parameter `LogDestination`, `mcmd`, when returning an error, also issued sometimes a warning that an earlier runtime error had been overwritten. This is now prevented by putting in proper checking and handling for any existing error as a new error is being thrown. (Bug #23211849)
- **Agent:** If a warning was returned by `mcmd` after it executed a `SET GLOBAL` statement to a `mysqld` node, the same warning was issued to the agent log (`mcmd.log`) again and again each time a new `SET GLOBAL` statement was executed. (Bug #23211783)
- **Agent:** When the parameter `ArbitrationRank` was being set for a `mysqld` node, the data nodes and management nodes were not restarted, so they had no knowledge that the arbitration rank for the `mysqld` node has been changed. (Bug #23148368)
- **Agent:** After a failed execution of the `import config` command using an older version of MySQL Cluster Manager, if MySQL Cluster Manager was then upgraded, restarting would fail afterward for some of the `mcmd` agents. (Bug #23148061)
- **Agent:** In the case where there was only one `ndb_mgmd` node in a cluster, during a rolling restart, after the `ndb_mgmd` node was just restarted and a first data node was stopped, a second data node might fail before it got restarted, complaining that there was no arbitrator for the cluster. It was because the management node was still in the process of asserting itself as an arbitrator. With this fix, the restarting of data nodes only begin after they have all seen an arbitrator established. (Bug #23148061)
- **Agent:** A new `ndbd` or `ndbmtbd` node could not be added to a cluster if `BackupDataDir` specified a non-default location for cluster backups, and there were backups existing at that location. (Bug #23123364)
- **Agent:** When already existing nodes were added again to a cluster without being started and the cluster was stopped, a subsequent `start cluster --initial` command failed. This was because `mcmd` attempted to recreate the already existing node groups, which this fix prevents. (Bug #23024367)
- **Agent:** While the `my.cnf` configuration files for added `mysqld` processes generated by `mcmd` used group suffixes for group titles in the files (for example `[mysqld.50]`), the `mysqld` nodes were started without using the `--defaults-group-suffix` option, causing the generated configurations to be unread by the `mysqld` node. With this fix, the group suffixes are no longer used. (Bug #22931198)
- **Agent:** On Windows platforms, an `mcmd` agent quit unexpectedly after it was unable to open a process handler to a PID. This fix makes `mcmd` able to handle the situation and not quit. (Bug #22886512)
- **Agent:** During the time when an `upgrade cluster` command is being performed, if a `mysqld` node that had not yet been upgraded failed somehow, it was mistakenly restarted with the newer instead of the older binary package. With this fix, the node is restarted with the original binary package it was running on. (Bug #22880634)
- **Agent:** Under some conditions, setting a configuration attribute for a cluster using the `set` command made `mcmd` quit unexpectedly, as more than one thread in MySQL Cluster Manager tried to change the cluster configurations together within a very short period of time, with one thread running into a checksum error for the repository directory and causing the agent to quit. This fix adds error checks, retries, and also exponential backoffs to handle the situation, in order to allow proper execution of the `set` command. (Bug #22865068)
- **Agent:** A `restore cluster` command failed because one of the management nodes in the cluster was in the “starting” status, which should not prevent a cluster from being restored. (Bug #22755257)
- **Agent:** When the `start process --added` command was used to start data nodes that were newly added, the command failed if there was an `ndbapi` node that was in the “connected” status, with the complaint that the `ndbapi` node was “already running.” (Bug #22726592)

- **Agent:** If an earlier attempt to stop a `mysqld` node failed and left the node in a status of “stopping,” a subsequent `stop cluster` command failed with a timeout. With this fix, the `stop cluster` command actually reattempts twice the shutdown of the “stopping” process, and throws a proper error if the attempts fail. The timeout period for each attempt is also reduced to 5 seconds. (Bug #22682222, Bug #24735542)

References: See also: Bug #19805950.

- **Agent:** When performing a `create cluster` command for a large cluster with many processes, `mcmd` tried to resolve the host name and check the validity of the package name for every single process, which could take a long time and cause the `create cluster` command to fail with a timeout. `mcmd` now avoids repeated lookups for the same host name or package name, so the timeouts are prevented. (Bug #22671177)
- **Agent:** When a backup was actually completed after an `abort backup` command had just been issued through the `mcmd` agent, the agent failed to respond to the backup completion, resulting in a timeout for the `abort backup` command. With this fix, the `abort backup` command will error out under the situation. (Bug #22655696)
- **Agent:** When setting a value for the `log-error` attribute for a `mysqld` node using the `set` command, `mcmd` failed to, as expected, append the extension `.err` to the supplied file name when it did not have an extension. (Bug #22588267)
- **Agent:** The `mysql` client hung as a `status` command was issued through it to a `mysqld` node, due to some communication packages remaining unhandled by `mcmd`. (Bug #22539167)
- **Agent:** The `autotune --dryrun` command did not write to the `.mcm` script file the TCP connection attributes (for example, `SendBufferMemory` and `ReceiveBufferMemory`) it would set for the cluster when the `--dryrun` option was not used. (Bug #22517603)
- **Agent:** With two clusters running separately on two different hosts in the same site, the `autotune --dryrun` command for a cluster failed with an internal error, complaining that a dump file [for the other cluster running on the other host] could not be opened because it did not exist on the host—which was to be expected. (Bug #22465053, Bug #79586)
- **Agent:** If a cluster log rotation happened while MySQL Cluster Manager was starting the cluster, an `mcmd` agent failure would occur. This fix makes sure log rotation is properly detected and handled by `mcmd`. (Bug #22296243)
- **Agent:** If distributed privileges for `mysqld` nodes were used and the `root` password has been set on the cluster's `mysqld` nodes, after a new `mysqld` process was added, MySQL Cluster Manager failed to connect to it, as `mcmd` tried in vain to log on to the new `mysqld` node as `root` with an empty password. With this fix, `mcmd` then attempts in that situation to log on to the node as the user `mcmd`. Also, even if the `root` account is not secured, when distributed privileges are used, `mcmd` now attempts first to connect to a new `mysqld` node as `mcmd` before it tries to create the `mcmd` user on the new node. (Bug #22274785, Bug #81391)
- **Agent:** For MySQL Cluster 7.4.8 and later, the `import config` command imported the deprecated `PortNumber` attribute under the [tcp] section of the cluster configuration file as “0.”. The attribute is now skipped during a configuration import. (Bug #22274785)
- **Agent:** A `create cluster` command failed with an error when one of the `mcmd` agents was down, even if that agent was not needed for the process. (Bug #22245706)
- **Agent:** When the utility `mysql_install_db` was run by the `mcmd` agent at the creation of a `mysqld` node, it was not run asynchronously, resulting sometimes in unnecessary delay for other `mcmd` processes. With this fix, the utility is now run asynchronously. (Bug #22238508)

- **Agent:** After a new host has been added to a cluster without a package being added for it, `mcmd` returned an error message to any `get` and `set` command, saying that the parameter to be get or set did not exist. With this fix, a proper error message is returned. (Bug #21894353)
- **Agent:** On Windows platform, setting the `datadir` option for a `mysqld` node in the Windows file path format caused the `mcmd` agent to stop unexpectedly after it failed to restart the `mysqld` node. It was due to a mishandling of the Windows format file path, which has now been fixed. (Bug #19209870)
- **Agent:** A rolling restart of the cluster performed by `mcmd` timed out while waiting for GCP and LCP takeover events to complete among the data nodes. With this fix, the timeout is avoided by having `mcmd` check for the status of the takeover events and wait until the involved data nodes are ready before trying to stop them. (Bug #14230789)
- **Client:** Spaces in quoted option values for MySQL Cluster Manager were lost (for example, with “--prompt='mcm1.4.1> ’,” the prompt for the `mcm` client became “mcm1.4.1>” [with no space at the end]). (Bug #24528495)
- **Client:** When the `mcm` client was being started with the `--debug` option, if the `mysql` client could not be found at the expected location, the `mcm` client failed with a segmentation fault. (Bug #24522244)
- **Client:** The success messages returned by the `autotune --dryrun` and `import config --dryrun` commands referred users to the agent log file for the proposed settings to be applied to the cluster, but the settings were not actually in the file. The success messages now give the path to the `.mcm` script file that contains the settings. (Bug #22280689)
- **Client:** After the `mcmd` agent on a certain host failed, the `mcm` client continued to report the statuses of the processes on the host to “running” while the `mcmd` agent log and the `ndb_mgmd` queries already showed their statuses to be “unknown.”. (Bug #22174415)

Changes in MySQL Cluster Manager 1.4.0 (2015-12-07)

This section documents all changes and bug fixes that have been applied in MySQL Cluster Manager 1.4.0 since the release of MySQL Cluster Manager version 1.3.6.

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

Functionality Added or Changed

- **Packaging:** MySQL Cluster Manager is now built and shipped with GLib-2.44.0, OpenSSL 1.0.1p, and the MySQL 5.6 client library. (Bug #22202878)
- **Agent:** When using the `import cluster` command, if a `mysqld` node was started on the command line with options outside of a special, pre-defined set, the import failed with the complaint that those options were unsupported. Now, the import continues, as long as those options and their values are also included in the node's configuration created by MySQL Cluster Manager for import. (Bug #21943518)
- **Agent:** A warning is now logged (if `log-level=warning`) when a failed process is not restarted because the parameter `StopOnError` is set to `true`. (Bug #21575241)
- **Agent:** Two new options have been introduced for the `upgrade cluster` command: `--retry` and `--nodeid`. They, together with the `--force` option, allow a retry after an initial attempt to upgrade a cluster has failed. See the description for `upgrade cluster` for details. (Bug #20469067, Bug #16932006, Bug #21200698)

- **Client:** The `get` command now returns attributes in the same order as the MySQL Cluster `ndb_mgmd` command does when the `--print-full-config` option is used, with the non-data nodes being listed first and other nodes listed in increasing order of their node IDs. (Bug #22202973)
- **Client:** A new `autotune` command has been introduced, which tunes a number of parameters of the cluster to optimize its performance. (Bug #22202855)
- **Client:** The `show settings` command has a new `--hostinfo` option, with which the command prints out information on the host the `mcm` client is connected to. (Bug #21923561)
- **Client:** You can now use the wildcard `*` (asterisk character) to match attribute names in a `get` command. See [The `get` Command](#) for examples. (Bug #18069656)

Bugs Fixed

- **Agent:** On Windows platform, after a cluster import, the subsequent cluster restart timed out if a non-default value of the option `--pid-file` had been imported for a `mysqld` node. (Bug #21943518)

References: This issue is a regression of: Bug #21111944.

- **Agent:** When a data node could not be restarted after a `set` command because some attributes were set wrongly, another `set` command could not be used to correct the attributes, because the `set` command required the data node to be running. With this fix, the second `set` command can now be executed even when the data node is not running, as long as the `--force` option is used. The failed node is then restarted, followed by a rolling restart of the cluster. (Bug #21943518)
- **Agent:** `restore cluster` timed out when the number of tables in the cluster was huge (>1000). It was because a timeout extension was blocked. This fix unblocks the extension. (Bug #21393857)
- **Agent:** At the initial startup of a large cluster (with memory size on the order of 10GB), the process might time out while waiting for a data node to start. This fix makes the transaction timeout longer for data node initiation. (Bug #21355383)
- **Agent:** Under some conditions, a `show status` command might report negative node group ID values for processes after an `add process` command was completed. That was because the agent reported the node group IDs before their proper values had arrived, after the creation of new node groups. This fix makes the agent wait for the correct node group IDs before reporting them. (Bug #21346804)
- **Agent:** After successful the execution of an `add process` and a subsequent `start process --added` command, a third command that was issued very shortly afterward might fail. This was due to the way the updates for the processes' statuses were handled after the new nodes were added, which has now been corrected. (Bug #21138604)

References: See also: Bug #21346804.

- **Agent:** Setting a value for a “key-only” option for a MySQL node (that is, an option that does not take a value—for example, `skip_show_database`) with the `set` command and restarting the cluster afterward caused `mcmd` to attempt a cluster upgrade and back up the cluster. (Bug #21098403)
- **Agent:** The `create site` command sometimes failed with the error message “Lost connection to MySQL server during query.” It was due to an error in the code that handled the socket, which has now been fixed. (Bug #21027818)
- **Agent:** Parameters listed under the `[mysqld default]` or `[tcp default]` section of the `config.ini` file were not imported as configuration parameters for unmanaged API nodes. (Bug #20889471)

- **Client:** Output of the `get` command used with the `--include-defaults (-d)` option did not include matching TCP attributes that had default values. (Bug #21895322)

Index

A

add hosts, 9
add package, 9
add process, 3, 9, 9, 15
Agent, 3, 4, 5, 8, 9, 15
ArbitrationRank, 9
autotune, 9

B

backup agents, 8, 9
backup cluster, 9
backupdatadir, 9
backupid, 9
bootstrap, 3

C

checksum, 5
Client, 3, 5, 9, 15
cluster import, 5
collect logs, 4, 5, 8, 9
config.ini, 15
CPU usage, 5
create cluster, 9
create root user, 9
create site, 4, 15

D

datadir, 9
delete cluster, 3, 4

E

error, 8
error codes, 5
exit report, 3

F

failed process, 3

G

get, 9, 15

H

hostname, 8
hosts, 9

I

import cluster, 9

import config, 4, 5, 8, 9
INCOMPLETE, 8
installing, 9

L

list backups, 9
list nextnodeid, 5
log, 9
log rotation, 9
log-level, 9
LogDestination, 9
log_error, 9
LSN, 9

M

mcmd upgrade, 9
Microsoft Windows, 9
my.cnf, 9
mysql client, 9
mysqld, 9
mysqld recovery, 8
mysqldump, 9
mysql_install_db, 9

N

ndbinfo_table_prefix, 9
ndbmtd, 9
ndb_config, 9
ndb_mgmd, 9
ndb_mgmt, 8
network, 9
network partitioning, 5
node status, 4
nodegroup, 3

O

OL7, 8
option vluaes, 9

P

Packaging, 8, 15
process handle, 3
process handler, 9
progress, 8
progressbar, 8

Q

quotes, 9

R

reset, 9
restart, 3, 5

restart cluster, 9
restore cluster, 3, 8, 9
rolling restart, 9
rotate log, 3

S

set, 3, 5, 9, 15
set global, 9
set wait_timeout, 9
show settings, 5
show status, 8, 9
shut down, 5
skip-init option, 5
start cluster, 3, 4, 5, 9
start cluster initial, 9
start process, 9
status, 9
stop cluster, 3, 8, 9
stop process, 3, 8, 9
StopOnError, 8
sync, 8

T

timeout, 4, 5, 9, 15

U

update process, 4, 8
upgrade mcm, 9

V

validate_password, 5

Z

zombie threads, 4

