

MySQL Cluster Manager 1.4.8 User Manual

Abstract

This is the User Manual for the *MySQL Cluster Manager*, version 1.4.8. It documents the *MySQL Cluster Manager Agent* and *MySQL Cluster Manager Client* software applications, which can be used to administer *MySQL NDB Cluster*, a version of the *MySQL Database System* (referred to hereafter as “MySQL Server” or simply “MySQL”) that incorporates the *NDB* storage engine for high availability and data redundancy in a distributed computing environment.

This Manual applies to MySQL Cluster Manager 1.4.8 and contains information that may not apply to older versions of the MySQL Cluster Manager software. For documentation covering previous MySQL Cluster Manager releases, see [MySQL Documentation: MySQL NDB Cluster](#), on the MySQL website.

MySQL Cluster Manager features. This manual describes features that may not be included in every version of MySQL Cluster Manager, and such features may not be included in the version of MySQL Cluster Manager licensed to you. If you have any questions about the features included in your version of MySQL Cluster Manager, refer to your MySQL Cluster Manager license agreement or contact your Oracle sales representative.

MySQL Cluster Manager, MySQL Server, and MySQL NDB Cluster features. This manual contains certain basic information about MySQL Server and MySQL NDB Cluster; however, it is not in any way intended as an exhaustive reference for either of these products.

Current versions of MySQL NDB Cluster compatible with MySQL Cluster Manager are based on versions 7.5 and 7.6 of the *NDB* storage engine; these versions of MySQL NDB Cluster are known as “MySQL NDB Cluster 7.5” and “MySQL NDB Cluster 7.6,” respectively. MySQL NDB Cluster functionality varies between MySQL NDB Cluster releases; MySQL Cluster Manager cannot supply or emulate MySQL NDB Cluster features that are not present in the version of the MySQL NDB Cluster software in use.

For information about MySQL NDB Cluster 7.6 and 7.5, see [MySQL NDB Cluster 7.5 and NDB Cluster 7.6](#).

MySQL Cluster Manager 1.4 supports MySQL NDB Cluster 8.0 on an experimental basis only, and should not be used with MySQL NDB Cluster 8.0 in a production environment. Some examples of MySQL 8.0 features that are **not** supported are listed here:

- The `default_auth_plugin` system variable.
- The `--upgrade` option for `mysqld`.
- The new `SET PERSIST` statement for persisting system variables.

If you do not have the MySQL Server and MySQL NDB Cluster documentation, you can obtain it free of charge from the [MySQL Documentation Library](#), on the MySQL website.

For notes detailing the changes in each release of MySQL Cluster Manager 1.4, see [MySQL Cluster Manager 1.4 Release Notes](#).

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.

Licensing information. This product may include third-party software, used under license. See the [MySQL Cluster Manager 1.4 License Information User Manual](#) for licensing information, including licensing information relating to third-party software that may be included in this release.

Document generated on: 2023-03-29 (revision: 75245)

Table of Contents

Preface and Legal Notices	v
1 Overview of MySQL Cluster Manager	1
1.1 MySQL Cluster Manager Terminology	1
1.2 MySQL Cluster Manager Architecture	2
1.3 Basic Operational Concepts for MySQL Cluster Manager	4
1.3.1 Eventual Consistency	4
2 MySQL Cluster Manager Installation, Configuration, Cluster Setup	5
2.1 Obtaining MySQL Cluster Manager	5
2.2 Supported Platforms and MySQL NDB Cluster Versions	5
2.3 MySQL Cluster Manager Installation	6
2.3.1 Installing MySQL Cluster Manager on Unix Platforms	6
2.3.2 Installing MySQL Cluster Manager on Windows Platforms	9
2.3.3 Setting the MySQL Cluster Manager Agent User Name and Password	15
2.4 MySQL Cluster Manager Configuration File	16
2.5 Upgrading MySQL Cluster Manager	17
3 Using MySQL Cluster Manager	21
3.1 <code>mcmd</code> , the MySQL Cluster Manager Agent	21
3.2 Starting and Stopping the MySQL Cluster Manager Agent	30
3.2.1 Starting and Stopping the Agent on Linux	30
3.2.2 Starting and Stopping the MySQL Cluster Manager Agent on Windows	31
3.3 Starting the MySQL Cluster Manager Client	33
3.4 Setting Up MySQL NDB Clusters with MySQL Cluster Manager	36
3.4.1 Creating a MySQL NDB Cluster with MySQL Cluster Manager	36
3.5 Importing MySQL NDB Clusters into MySQL Cluster Manager	38
3.5.1 Importing a Cluster Into MySQL Cluster Manager: Basic Procedure	38
3.5.2 Importing a Cluster Into MySQL Cluster Manager: Example	40
3.6 MySQL NDB Cluster Backup and Restore Using MySQL Cluster Manager	50
3.6.1 Requirements for Backup and Restore	50
3.6.2 Basic MySQL NDB Cluster Backup and Restore Using MySQL Cluster Manager	51
3.7 Backing Up and Restoring MySQL Cluster Manager Agents	57
3.8 Restoring a MySQL Cluster Manager Agent with Data from Other Agents	59
3.9 Setting Up MySQL NDB Cluster Replication with MySQL Cluster Manager	59
3.10 NDB Cluster Downgrades Involving Logical Downgrades of MySQL Nodes	62
4 MySQL Cluster Manager Client Commands	65
4.1 Online Help and Information Commands	71
4.2 MySQL Cluster Manager Site and Agent Commands	77
4.2.1 The <code>add hosts</code> Command	77
4.2.2 The <code>remove hosts</code> Command	78
4.2.3 The <code>change log-level</code> Command	79
4.2.4 The <code>rotate log</code> Command	79
4.2.5 The <code>collect logs</code> Command	80
4.2.6 The <code>create site</code> Command	81
4.2.7 The <code>delete site</code> Command	82
4.2.8 The <code>list sites</code> Command	82
4.2.9 The <code>list hosts</code> Command	83
4.2.10 The <code>show settings</code> Command	83
4.2.11 The <code>stop agents</code> Command	84
4.2.12 The <code>version</code> Command	84
4.2.13 The <code>show warnings</code> Command	85
4.3 MySQL Cluster Manager Package Commands	85
4.3.1 The <code>add package</code> Command	85
4.3.2 The <code>delete package</code> Command	87
4.3.3 The <code>list packages</code> Command	88
4.4 MySQL Cluster Manager Cluster Commands	89
4.4.1 The <code>create cluster</code> Command	89

4.4.2	The <code>delete cluster</code> Command	93
4.4.3	The <code>list clusters</code> Command	94
4.4.4	The <code>list nextnodeids</code> Command	94
4.4.5	The <code>restart cluster</code> Command	94
4.4.6	The <code>show status</code> Command	95
4.4.7	The <code>start cluster</code> Command	99
4.4.8	The <code>stop cluster</code> Command	101
4.4.9	The <code>autotune</code> Command	101
4.4.10	The <code>upgrade cluster</code> Command	102
4.5	MySQL Cluster Manager Configuration Commands	105
4.5.1	The <code>get</code> Command	108
4.5.2	The <code>reset</code> Command	119
4.5.3	The <code>set</code> Command	125
4.6	MySQL Cluster Manager Process Commands	131
4.6.1	The <code>add process</code> Command	132
4.6.2	The <code>change process</code> Command	135
4.6.3	The <code>list processes</code> Command	137
4.6.4	The <code>start process</code> Command	138
4.6.5	The <code>stop process</code> Command	138
4.6.6	The <code>update process</code> Command	139
4.6.7	The <code>remove process</code> Command	140
4.7	MySQL Cluster Manager Backup and Restore Commands	141
4.7.1	The <code>abort backup</code> Command	141
4.7.2	The <code>backup cluster</code> Command	141
4.7.3	The <code>list backups</code> Command	143
4.7.4	The <code>restore cluster</code> Command	144
4.7.5	The <code>backup agents</code> Command	146
4.8	MySQL Cluster Manager Cluster Importation Commands	147
4.8.1	The <code>import cluster</code> Command	147
4.8.2	The <code>import config</code> Command	147
5	MySQL Cluster Manager Limitations and Known Issues	149
5.1	MySQL Cluster Manager Usage and Design Limitations	149
5.2	MySQL Cluster Manager 1.4.8 Limitations Relating to the MySQL Server	149
5.3	MySQL Cluster Manager Limitations Relating to MySQL NDB Cluster	150
5.4	Syntax and Related Issues in MySQL Cluster Manager	151
A	Attribute Summary Tables	153
A.1	Management Node Configuration Parameters	153
A.2	Data Node Configuration Parameters	154
A.3	API Node Configuration Parameters	160
A.4	Other Node Configuration Parameters	160
A.5	MySQL Server Option and Variable Reference for MySQL Cluster	162
	Index	171

Preface and Legal Notices

This is the User Manual for the *MySQL Cluster Manager*, version 1.4.8. It documents the *MySQL Cluster Manager Agent* and *MySQL Cluster Manager Client* software applications, which can be used to administer *MySQL NDB Cluster*, a version of the *MySQL Database System* (referred to hereafter as “MySQL Server” or simply “MySQL”) that incorporates the [NDB](#) storage engine for high availability and data redundancy in a distributed computing environment.

This manual does contain certain basic information about MySQL and MySQL NDB Cluster; however, it is not in any way intended as an exhaustive reference for either of these products.

Versions of MySQL NDB Cluster compatible with MySQL Cluster Manager are based on versions 7.5 and 7.6 of the [NDB](#) storage engine; these versions of MySQL NDB Cluster are known as “MySQL NDB Cluster 7.5,” and “MySQL NDB Cluster 7.6”, respectively. MySQL NDB Cluster functionality varies between MySQL NDB Cluster releases; MySQL Cluster Manager cannot supply or emulate MySQL NDB Cluster features that are not present in the version of the MySQL NDB Cluster software in use.

For complete information about MySQL Server and MySQL NDB Cluster, see [MySQL NDB Cluster 7.5 and NDB Cluster 7.6](#).



Warning

MySQL Cluster Manager 1.4 supports MySQL NDB Cluster 8.0 on an experimental basis only and should not be used with MySQL NDB Cluster 8.0 in a production environment. Some examples of MySQL 8.0 features that are **not** supported are listed here:

- The `default_auth_plugin` system variable.
- The `--upgrade` option for `mysqld`.
- The new `SET PERSIST` statement for persisting system variables.

Licensing information. This product may include third-party software, used under license. See the [MySQL Cluster Manager 1.4 License Information User Manual](#) for licensing information, including licensing information relating to third-party software that may be included in this release.

Legal Notices

Copyright © 2009, 2023, Oracle and/or its affiliates.

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 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" or "commercial computer software documentation" 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.

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, Java, and MySQL 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.

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

<https://www.oracle.com/corporate/accessibility/>.

Access to Oracle Support for Accessibility

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

<https://www.oracle.com/corporate/accessibility/learning-support.html#support-tab>.

Chapter 1 Overview of MySQL Cluster Manager

Table of Contents

1.1 MySQL Cluster Manager Terminology	1
1.2 MySQL Cluster Manager Architecture	2
1.3 Basic Operational Concepts for MySQL Cluster Manager	4
1.3.1 Eventual Consistency	4

This chapter provides a overview of MySQL Cluster Manager, as well as its architecture, purpose, and capabilities.

1.1 MySQL Cluster Manager Terminology

This section provides definitions of key terms used to describe MySQL Cluster Manager and its components in this manual and in other documentation relating to MySQL Cluster Manager and MySQL NDB Cluster.

Site.

A set of hosts on which MySQL NDB Cluster processes to be managed by MySQL Cluster Manager are located. A site can include one or more clusters.

Cluster.

A MySQL NDB Cluster deployment. A cluster consists of a set of MySQL NDB Cluster processes running on one or more hosts. A minimal cluster is usually considered to include one management node, two data nodes, and one SQL node. A typical production cluster may have one or two management nodes, several SQL nodes, and 4 or more data nodes. The exact numbers of data and SQL nodes can vary according to data size, type and rating of hardware used on the hosts, expected throughput, network characteristics, and other factors; the particulars are beyond the scope of this document, and you should consult [MySQL NDB Cluster 7.5](#) and [NDB Cluster 7.6](#), for more specific information and guidelines.

Host.

A computer. The exact meaning depends on the context:

- A computer where one or more MySQL NDB Cluster processes are run. In this context, we sometimes refer more specifically to a *cluster host*.

The number of cluster processes and number of cluster hosts may be, but are not necessarily, the same.

- A computer where an instance of the MySQL Cluster Manager agent runs.

In order to run a MySQL NDB Cluster using MySQL Cluster Manager, the MySQL Cluster Manager agent must be running on each host where cluster processes are to be run. In other words, when using MySQL Cluster Manager, all cluster hosts must also be MySQL Cluster Manager agent hosts (although the reverse is not necessarily true). Therefore, you should understand that anytime we use the term *host*, we are referring to a host computer in both of the senses just given.

Process.

In the context of MySQL NDB Cluster, a process (more specifically, a cluster process) is a MySQL NDB Cluster node, of one of the following 3 types: management node (`ndb_mgmd`), data node (`ndbd` or `ndbmt`), or SQL node (`mysqld`). For more information about these node types and their functions in a cluster, see [NDB Cluster Core Concepts](#), and [NDB Cluster Nodes, Node Groups, Fragment Replicas, and Partitions](#).

Package.

A copy of the MySQL NDB Cluster software. This includes the binary executables needed to run the cluster processes of the desired types on a given host. The simplest way to make sure that this is done is to place a copy of the entire MySQL NDB Cluster distribution on each computer that you intend to use as a cluster host.

Configuration attribute.

A value whose setting affects cluster operations in a clearly defined and measurable way. When running MySQL NDB Cluster manually, configuration is accomplished using cluster configuration parameters, MySQL server options, and MySQL system and status variables; MySQL Cluster Manager masks the differences between these, providing a unified view of them; see [Configuration attributes](#), for more information.

Agent.

A MySQL Cluster Manager process that runs on each cluster host, responsible for managing the cluster processes running on that host.

Client.

The MySQL Cluster Manager client is a software application that allows a user to connect to MySQL Cluster Manager and perform administrative tasks, such as (but not limited to): creating, starting, and stopping clusters; obtaining cluster and cluster process status reports; getting cluster configuration information and setting cluster configuration attributes.

1.2 MySQL Cluster Manager Architecture

This section provides an architectural overview of MySQL Cluster Manager, its components, and their deployment.

MySQL Cluster Manager is a distributed client/server application consisting of two main components. The MySQL Cluster Manager agent is a set of one or more agent processes that manage NDB Cluster nodes, and the MySQL Cluster Manager client provides a command-line interface to the agent's management functions.

Agent.

The MySQL Cluster Manager agent is comprised of the set of all MySQL Cluster Manager agent processes running on the hosts making up a given management site. A MySQL Cluster Manager agent process is a daemon process that runs on each host to be used in the cluster. In MySQL Cluster Manager, there is no single central server or process; all agents collaborate in managing a cluster as a whole. This means that any connected agent can be used to carry out tasks that effect the entire cluster.

Each agent process is responsible for managing the MySQL NDB Cluster nodes running on the host where the agent is located. MySQL NDB Cluster management and SQL nodes are managed directly by the MySQL Cluster Manager agent; cluster data nodes are managed indirectly, using the cluster management nodes.

Management responsibilities handled by the MySQL Cluster Manager agent include the following:

- Starting, stopping, and restarting cluster nodes
- Cluster configuration changes
- Cluster software upgrades
- Host and node status reporting
- Recovery of failed cluster nodes

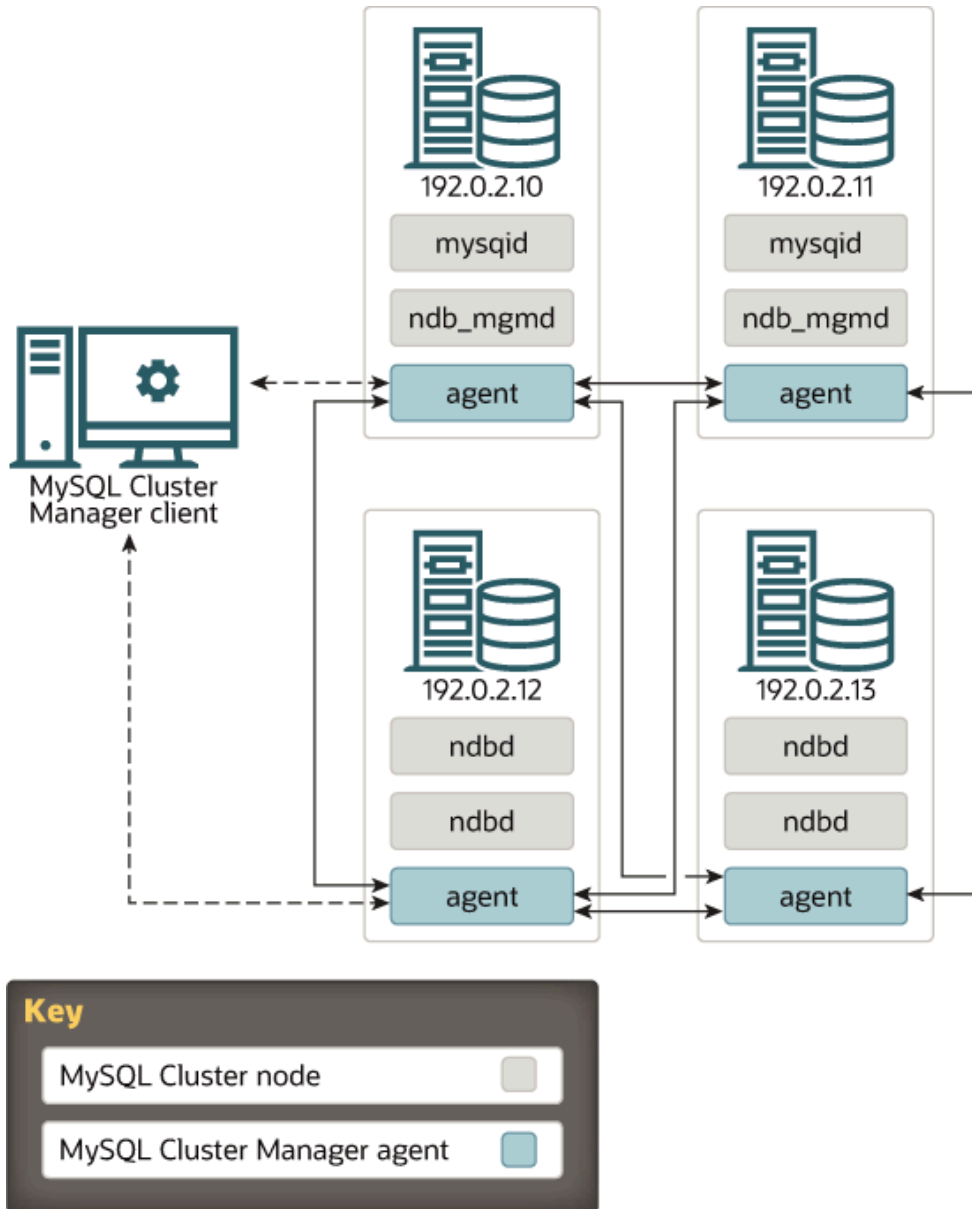
Creating, performing initial configuration of, or starting a cluster, requires that agent processes be running on all cluster hosts. Once the cluster has been started, it continues to run even if one or more agent processes fail. However, any failed agent processes must be restarted before you can perform additional cluster management functions.

Client.

A MySQL Cluster Manager client is a software application used to access an MySQL Cluster Manager agent. In MySQL Cluster Manager, the client is actually nothing more than the `mysql` command-line client, started with the options that are necessary for it to connect to an MySQL Cluster Manager agent. MySQL Cluster Manager 1.4.8 and later releases include an `mcm` client for ease of use; this client consists of a script that acts as a wrapper for the `mysql` client with these options.

By way of example, we show how MySQL Cluster Manager would be deployed for use with a MySQL NDB Cluster running on 4 host computers. This is illustrated in the following diagram:

Figure 1.1 MySQL Cluster Manager Deployment



In this example cluster, two of the hosts each house a management server (`ndb_mgmd`) and an SQL node (`mysqld`); the other 2 hosts each house two data nodes (`ndbd`). However, regardless of the distribution of cluster nodes among the hosts, a MySQL Cluster Manager agent process must be running on each host.

A MySQL Cluster Manager client can be used to access the agent from any of the hosts making up the management site to which the cluster belongs. In addition, the client can be used on any computer that has a network connection to at least 1 of the hosts where an agent process is running. The computer

where the client itself runs is not required to be one of these hosts. The client can connect to and use different agent processes on different hosts within the management site, at different times, to perform cluster management functions.

1.3 Basic Operational Concepts for MySQL Cluster Manager

This section explains some basic operational concepts for MySQL Cluster Manager.

1.3.1 Eventual Consistency

MySQL Cluster Manager guarantees *eventual consistency* among agents, meaning that:

- Any message communicated among agents is either delivered or not delivered to ALL agents (instead of delivered to some and missed by others).
- Order of delivery for any sequence of messages is always identical for all agents (that is, messages cannot get out of order for some agents).

Beyond that, there is no guarantee for message synchronization: a message is not guaranteed to be received and executed within a specific window of time for all agents. The result is that any agent may lag behind in processing messages for any reasons such as network traffic, machine loading, or thread scheduling. Situations like the following might then occur: while Agent C is lagging behind Agent A and B, the two agents have completed some cluster reconfiguration that does not involve any local actions for Agent C; a client connected to Agent A or B might have received a success message for the reconfiguration, while a client querying Agent C is told that the reconfiguration is still in process, since the completion message has not yet reached Agent C.

Such temporary inconsistency among the agents should not be a concern. While an agent might be lagging behind, the guaranteed eventual consistency means that, unless prevented by a network error or some other issues, any lagging agent will eventually catch up with the other agents, and all agents will eventually get a consistent view of the site.

Chapter 2 MySQL Cluster Manager Installation, Configuration, Cluster Setup

Table of Contents

2.1 Obtaining MySQL Cluster Manager	5
2.2 Supported Platforms and MySQL NDB Cluster Versions	5
2.3 MySQL Cluster Manager Installation	6
2.3.1 Installing MySQL Cluster Manager on Unix Platforms	6
2.3.2 Installing MySQL Cluster Manager on Windows Platforms	9
2.3.3 Setting the MySQL Cluster Manager Agent User Name and Password	15
2.4 MySQL Cluster Manager Configuration File	16
2.5 Upgrading MySQL Cluster Manager	17

This chapter discusses basic installation and configuration of the MySQL Cluster Manager Management Agent, connecting to the agent with the MySQL Cluster Manager client, and the basics of creating or importing a cluster using MySQL Cluster Manager.

2.1 Obtaining MySQL Cluster Manager

MySQL Cluster Manager is available only through commercial license. To learn more about licensing terms, and to obtain information about where and how to download MySQL Cluster Manager, visit <https://www.mysql.com/products/cluster/mcm/>, or contact your Oracle representative.

2.2 Supported Platforms and MySQL NDB Cluster Versions

For a list of platforms supported by MySQL Cluster Manager 1.4.8, see *Supported Platforms: MySQL Cluster Manager* at <https://www.mysql.com/support/supportedplatforms/cluster-manager.html>, or contact your Oracle representative.

MySQL Cluster Manager 1.4.8 supports the following MySQL NDB Cluster release versions:

- MySQL NDB Cluster 7.5 (see [MySQL NDB Cluster 7.5 and NDB Cluster 7.6](#)), beginning with MySQL NDB Cluster 7.5.4
- MySQL NDB Cluster 7.6 (see [MySQL NDB Cluster 7.5 and NDB Cluster 7.6](#)), beginning with MySQL NDB Cluster 7.6.6.



Warning

MySQL Cluster Manager 1.4.8 supports MySQL NDB Cluster 8.0 on an experimental basis only and should not be used with MySQL NDB Cluster 8.0 in a production environment. Some examples of MySQL 8.0 features that are **not** supported are listed here:

- Any authentication plugins other than `mysql_native_password` when connecting with `mcmd`.
- The `--upgrade` option for `mysqld`.
- The `SET PERSIST` statement for persisting system variables.

Prior to installation, you must obtain the correct build of MySQL Cluster Manager for your operating system and hardware platform. For Unix platforms, MySQL Cluster Manager is delivered as a Unix `.tar.gz` archive, whose name is in the format of `mcm-1.4.8-cluster-7.6.13-`

[linux-distro-arch.tar.gz](#). For Windows platforms, an MSI installer file is provided. All MySQL Cluster Manager 1.4.8 packages include MySQL NDB Cluster 7.6.13.

When selecting the appropriate installer for your operating system and hardware, keep in mind that 32-bit programs can normally be run on 64-bit platforms, but that the reverse is not true.

2.3 MySQL Cluster Manager Installation

Installation of the MySQL Cluster Manager agent and client programs varies according to platform. On Unix platforms, you must extract the binaries and other files from a `.tar.gz` archive and copy them to the correct locations manually. For Windows platforms, an MSI installer is provided which largely automates this process. Installation of MySQL Cluster Manager using each of these methods is covered in the next two sections.



Note

License keys were required in order to use some very early releases of MySQL Cluster Manager. They are not needed in order to use recent versions, including MySQL Cluster Manager 1.4.8.

2.3.1 Installing MySQL Cluster Manager on Unix Platforms

Install the MySQL Cluster Manager agent on Linux and similar platforms by following the instructions below.

Extract the MySQL Cluster Manager 1.4.8 program and other files from the distribution archive.

You must install a copy of MySQL Cluster Manager on each computer that you intend to use as a MySQL NDB Cluster host. In other words, you need to install MySQL Cluster Manager on each host that is a member of a MySQL Cluster Manager management site. For each host, you should use the MySQL Cluster Manager build that matches that computer's operating system and processor architecture.

On Linux systems, you can unpack the archive using the following command, using `mcm-1.4.8-cluster-7.6.13-linux-sles11-x86-64bit.tar.gz` as an example (the actual filename will vary according to the MySQL Cluster Manager build that you intend to deploy):

```
$> tar -zxvf mcm-1.4.8-cluster-7.6.13-linux-sles11-x86-64bit.tar.gz
```

This command unpacks the archive into a directory having the same name as the archive, less the `.tar.gz` extension. The top-level directory within the archive is named `mcm-1.4.8`.



Important

Because the Solaris version of `tar` cannot handle long filenames correctly, the MySQL Cluster Manager program files may be corrupted if you try to use it to unpack the MySQL Cluster Manager archive. To get around this issue on Solaris operating systems, you should use GNU `tar` (`gtar`) rather than the default `tar` supplied with Solaris. On Solaris 10, `gtar` is often already installed in the `/usr/sfw/bin` directory, although the `gtar` executable may not be included in your path. If `gtar` is not present on your system, please consult the [Solaris 10 system documentation](#) for information on how to obtain and install it.

In general, the location where you place the unpacked MySQL Cluster Manager directory and the name of this directory can be arbitrary. However, we recommend that you use a standard location for optional software, such as `/opt` on Linux systems, and that you name the directory using the 1.4.8 version number. (This facilitates subsequent upgrades.) On a typical Linux system you can accomplish this task like this:

```
$> cd mcm-1.4.8-cluster-7.6.13-linux-sles11-x86-64bit
```

```
$> mv mcm-1.4.8 /opt/mcm-1.4.8
```

For ease of use, we recommend that you put the MySQL Cluster Manager files in the same directory on each host where you intend to run it.

Contents of the MySQL Cluster Manager Unix Distribution Archive.

If you change to the directory where you placed the extracted MySQL Cluster Manager archive and list the contents, you should see something similar to what is shown here:

```
$> cd /opt/mcm-1.4.8
$> ls
bin cluster etc lib libexec licenses share var
```

These directories are described in the following table:

Table 2.1 Contents of the MySQL Cluster Manager Unix distribution archive, by directory

Directory	Contents
<code>bin</code>	MySQL Cluster Manager agent startup scripts
<code>cluster</code>	Contains the MySQL NDB Cluster 7.6.13 binary distribution
<code>etc</code>	Contains the agent configuration file (<code>mcmd.ini</code>)
<code>etc/init.d</code>	Init scripts
<code>lib</code> and subdirectories	Libraries needed to run the MySQL Cluster Manager agent
<code>libexec</code>	MySQL Cluster Manager agent and client executables
<code>licenses/lgpl</code>	An archive containing source code (including licensing and documentation), for <code>glib 2.1</code>
<code>share/doc/mcmd</code>	<code>README.txt</code> file
<code>var</code>	XML files containing information needed by MySQL Cluster Manager about processes, attributes, and command syntax

Normally, the only directories of those shown in the preceding table that you need be concerned with are the `bin` and `etc` directories.

For MySQL Cluster Manager 1.4.8 distributions that include MySQL NDB Cluster, the complete MySQL NDB Cluster 7.6.13 binary distribution is included in the `cluster` directory. Within this directory, the layout of the MySQL NDB Cluster distribution is the same as that of the standalone MySQL NDB Cluster binary distribution. For example, MySQL NDB Cluster binary programs such as `ndb_mgmd`, `ndbd`, `ndbmtd`, and `ndb_mgm` can be found in `cluster/bin`. For more information, see [MySQL Installation Layout for Generic Unix/Linux Binary Package](#), and [Installing an NDB Cluster Binary Release on Linux](#), in the *MySQL Manual*.

If you wish to use the included MySQL NDB Cluster software, it is recommended that you move the `cluster` directory and all its contents to a location outside the MySQL Cluster Manager installation directory, such as `/opt/ndb-version`. For example, on a Linux system, you can move the MySQL NDB Cluster NDB 7.6.13 software that is bundled with MySQL Cluster Manager 1.4.8 to a suitable location by first navigating to the MySQL Cluster Manager installation directory and then using a shell command similar to what is shown here:

```
$> mv cluster /opt/ndb-7.6.13
```



Note

The `mcmd --bootstrap` option uses the included MySQL NDB Cluster binaries in the installation directory's `cluster` directory and does not work if

they cannot be found there. To work around this issue, create a symbolic link to the correct directory in its place, like this:

```
$> ln -s /opt/ndb-7.6.13 cluster
```

After doing this, you can use the `mcm` client commands `add package` and `upgrade cluster` to upgrade any desired cluster or clusters to the new MySQL NDB Cluster software version.



Important

On Linux platforms, do not attempt to install the MySQL NDB Cluster software by the RPM, Debian, or other installation packages for any package management systems. They install MySQL NDB Cluster differently than the binary distribution that comes with the MySQL Cluster Manager archive, and that will cause issue in the future when you try to upgrade your cluster with MySQL Cluster Manager.

The MySQL Cluster Manager agent by default writes its log file as `mcmd.log` in the installation directory. When the agent runs for the first time, it creates a directory where the agent stores its own configuration data; by default, that is `../mcm_data` (relative to the installation directory of MySQL Cluster Manager). The configuration data, log files, and data node file systems for a given MySQL NDB Cluster under MySQL Cluster Manager control, and named `cluster_name`, can be found in `clusters/cluster_name` under this data directory (sometimes also known as the MySQL Cluster Manager *data repository*).

The location of the MySQL Cluster Manager agent configuration file, log file, and data directory can be controlled with `mcmd` startup options or by making changes in the agent configuration file. To simplify upgrades of MySQL Cluster Manager, we recommend that you change the data repository to a directory outside the MySQL Cluster Manager installation directory, such as `/var/opt/mcm`. See [Section 2.4, “MySQL Cluster Manager Configuration File”](#), and [Section 3.2, “Starting and Stopping the MySQL Cluster Manager Agent”](#), for more information.

On Linux and other Unix-like systems, you can set up the MySQL Cluster Manager agent to run as a daemon, using the init script that is supplied with the MySQL Cluster Manager distribution. To do this, follow the steps listed here:

1. Copy the file `/etc/init.d/mcmd` under the MySQL Cluster Manager installation directory to your system's `/etc/init.d/` directory (or equivalent). On a typical Linux system, you can do this using the following command in the system shell, where `mcmdir` is the MySQL Cluster Manager installation directory:

```
$> cd mcmdir/etc/init.d
$> cp mcmd /etc/init.d/mcmd
```

2. Make sure that this file has appropriate permissions and is executable by the user account that runs MySQL Cluster Manager. On a typical Linux system, this can be done by executing commands in your system shell similar to those shown here:

```
$> chown mcmuser /etc/init.d/mcmd
$> chmod 755 /etc/init.d/mcmd
```

Be sure to refer to your operating system documentation for exact information concerning the commands needed to perform these operations, as they may vary between platforms.

3. Open the file `/etc/init.d/mcmd` in a text editor. Here, we show a portion of this file, in which we have highlighted the two lines that need to be updated:

```
MCMD_SERVICE="mcmd"
MCMD_PSERVICE="MySQL Cluster Manager"
MCMD_ROOTDIR=@@MCMD_ROOTDIR@@
MCMD_BIN=" $MCMD_ROOTDIR/bin/mcmd"
MCMD_CONFIG=" $MCMD_ROOTDIR/etc/mcmd.ini"
```

```
# Run service as non-root user
MCMD_USER=@@MCMD_USER@@
SU="su --login $MCMD_USER --command"
```

In the first of these lines, replace the placeholder `@@MCMD_ROOTDIR@@` with the complete path to the MySQL Cluster Manager installation directory. In the second of these lines, replace the placeholder `@@MCMD_USER@@` with the name of the system user that runs the MySQL Cluster Manager agent (note that this must *not* be the system `root` account). Save the edited file.

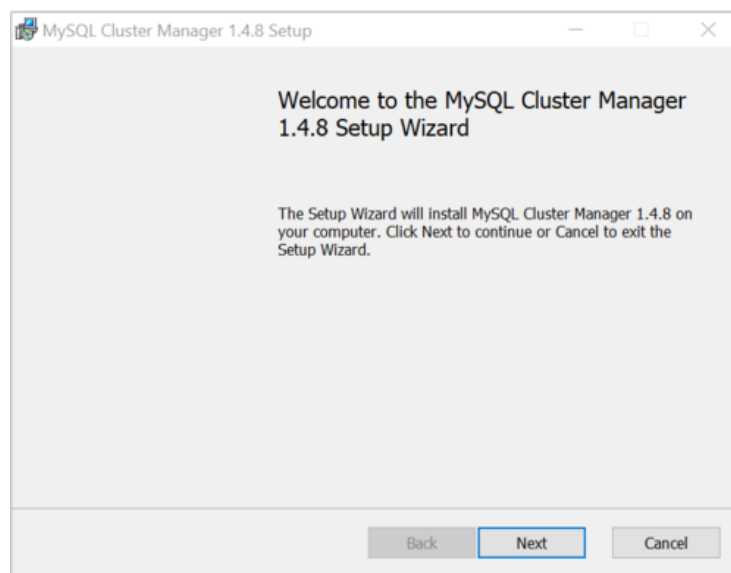
The MySQL Cluster Manager agent should now be started automatically whenever the system is restarted.

When the agent is configured as a daemon, cluster processes are started automatically when the agent is restarted, as long as the cluster was running when the agent shut down. *Note that `StopOnError` must be disabled (set to 0) for all data nodes in order for this to work.* If the cluster was stopped when the agent shut down, it is necessary to supply a script that waits for the agent to complete its startup and recovery phases, and then, when the agent is ready, starts the cluster using a command such as `mcmdir/bin/mcm -e 'start cluster --background cluster_name!'`.

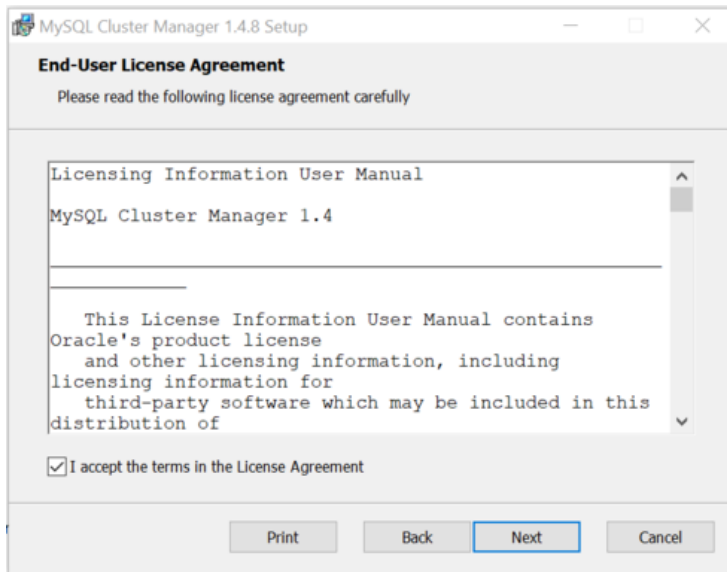
2.3.2 Installing MySQL Cluster Manager on Windows Platforms

To install, for example, MySQL Cluster Manager 1.4.8 on Windows platforms, obtain the MSI installer for it (see [Section 2.1, “Obtaining MySQL Cluster Manager”](#)). Run the installer by double-clicking it in Windows Explorer (some versions of Windows also provide an **Install** item in the Windows Explorer menu that can be used to run the installer). When you start the installer, you may see a Windows **Security Warning** screen. If you obtained the installer from a trusted source and know that it has not been tampered with, choose **Run** from the dialog, which allows the installer to continue to the **Welcome** screen, as shown here:

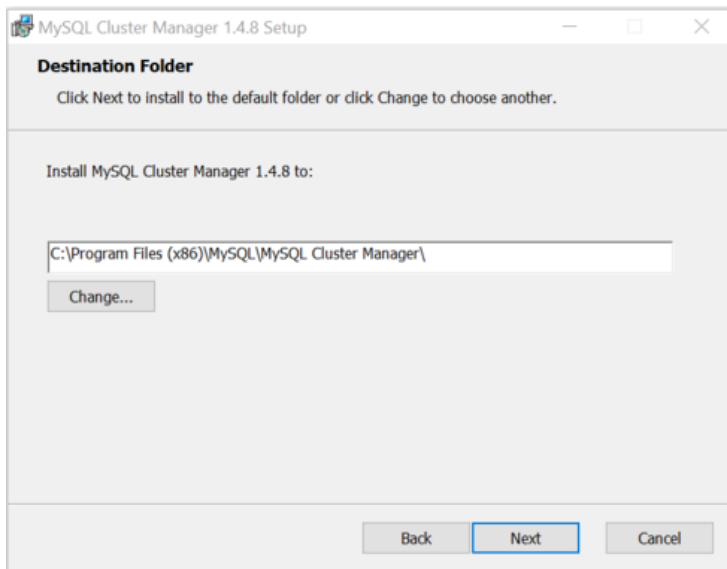
Figure 2.1 MySQL Cluster Manager MSI Installer: Welcome Screen




Click the **Next** button to continue to the **License Agreement** screen. You should read the license text in the text area, and when you have done so, check the box labelled **I accept the terms in the License Agreement**. Until you have checked the box, you cannot complete the MySQL Cluster Manager installation; it is possible only to print the license, go back to the previous screen, or cancel the installation (using the buttons labelled **Print**, **Back**, and **Cancel**, respectively). Checking the box enables the **Next** button, as shown here:

Figure 2.2 MySQL Cluster Manager MSI Installer: Accepting License Agreement


Click the **Next** button to continue to the **Destination Folder** screen, where you can choose the installation directory. The next figure shows the **Destination Folder** screen with the default location. On English-language 64-bit Windows systems, this is `C:\Program Files\MySQL\MySQL Cluster Manager\`; on English 32-bit Windows systems, the default is `C:\Program Files (x86)\MySQL\MySQL Cluster Manager\`. You can click the **Change** button to change the directory where MySQL Cluster Manager should be installed; the default directory is adequate for most cases.

Figure 2.3 MySQL Cluster Manager MSI Installer: Destination Folder

Once you have selected the destination directory, the installer has gathered all the information that it requires to perform the installation. Click **Next** to continue to the **Ready** screen, shown here:

Figure 2.4 MySQL Cluster Manager MSI Installer: Ready to Install

Click the **Install** button to install MySQL Cluster Manager. As the installer begins to copy files and perform other tasks affecting the system, you may see a warning dialog from Windows User Access Control. If this occurs, click the **Yes** button to allow the installation to continue. A **Setup Wizard** screen with a progress bar is displayed while the installer runs, as shown in the next figure:

Figure 2.5 MySQL Cluster Manager MSI Installer: Setup Wizard


The **Setup Wizard** may require several minutes to copy all of the necessary files for MySQL Cluster Manager and MySQL NDB Cluster to the installation directory and to perform other required changes.

**Note**

The MySQL Cluster Manager installer places MySQL NDB Cluster in the `cluster` directory under the installation directory (by default, that is `C:\Program Files\MySQL\MySQL Cluster Manager\cluster` or `C:\Program Files (x86)\MySQL\MySQL Cluster Manager\cluster`). The location of the MySQL NDB Cluster binaries is not separately configurable when using the MySQL Cluster Manager installer.

When the **Setup Wizard** finishes, the installer displays the **Installation Completed** screen, as shown here:

Figure 2.6 MySQL Cluster Manager MSI Installer: Installation Completed



MySQL Cluster Manager 1.4.8 has now been installed to the destination directory; click the **Finish** button to exit the installer.

As mentioned elsewhere (see, for example, [Section 3.4.1, “Creating a MySQL NDB Cluster with MySQL Cluster Manager”](#)), you must install a copy of MySQL Cluster Manager on each computer where you intend to host a MySQL NDB Cluster node. Therefore, the above procedure must be performed separately on each host computer. For ease of installations and upgrades on multiple machines, it is recommended that you install MySQL Cluster Manager to the same location on each host. `C:\Program Files\MySQL\MySQL Cluster Manager\` (for 64-bit systems) or `C:\Program Files (x86)\MySQL\MySQL Cluster Manager\` (for 32-bit systems) is the default location for installation, but it is possible to install MySQL Cluster Manager to an alternate location such as `C:\mcm\`.

2.3.2.1 Installing the MySQL Cluster Manager Agent as a Windows Service

After installing the MySQL Cluster Manager Agent as a Windows service, you can start and stop the agent using the Windows Service Manager. The installation also configures the agent to start automatically whenever Windows starts, and to shut down safely whenever Windows shuts down.

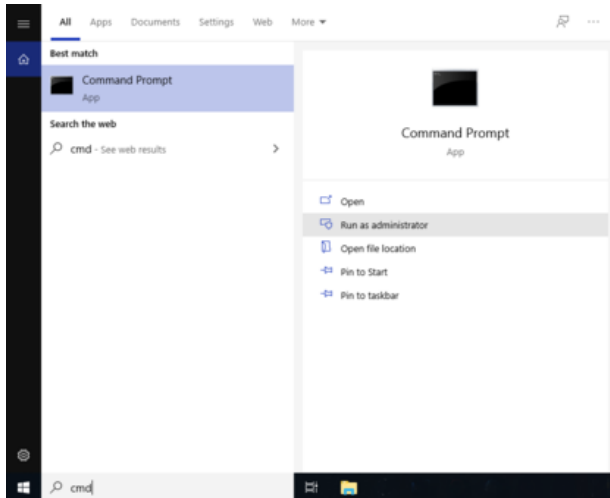


Note

The Windows service can be used to control the running of MySQL Cluster Manager agents on a single host only. To shut down agents on multiple hosts, you can use the `stop agents` command in the MySQL Cluster Manager client.

The installation is performed using the command prompt (`cmd.exe`); as with installing or removing any Windows service, it must also be done as a user having sufficient permissions, such as the system Administrator account.

If the account you are currently using has Administrator privileges, you can simply start `cmd.exe`. Otherwise, you must run the command prompt program as the Administrator. To do this, first locate a shortcut to the command prompt. You can do this by typing `cmd` into the search box in the Windows **Taskbar**, and then select from the search results **Command Prompt > Run as Administrator**. You can see how this looks on a typical Windows system in the next figure.

Figure 2.7 Installing the MySQL Cluster Manager Agent as a Windows Service: Command Prompt

If a Windows UAC dialog referring to `cmd.exe` appears, click **Yes** to allow the command prompt to run as Administrator and thus to continue. You should now have a command prompt window open on your desktop, running a session with Administrator privileges.

To install the MySQL Cluster Manager agent as a service, we use the `SC CREATE` command. This command allows us to specify a name for the service (for use with `SC START` and `SC STOP` or `NET START` and `NET STOP` commands), a display name (to be shown in the Service Manager), a startup mode (automatic or manual start), and a path to the executable to be run as a service (use `mcmd-svc.exe` rather than `mcmd.exe` as the executable). The path must also include any arguments needed by the program; in the case of MySQL Cluster Manager, `mcmd-svc.exe` must be told where to find its configuration file by the `--defaults-file` option. Both of these paths must be absolute.




Important

Installation of the MySQL Cluster Manager agent as a service is recommended. However, you should *not* install MySQL NDB Cluster processes (`ndb_mgmd.exe`, `ndbd.exe`, `ndbmttd.exe`, `mysqld.exe`) as services on Windows hosts to be used as MySQL NDB Cluster nodes under management by MySQL Cluster Manager, since the MySQL Cluster Manager agent itself controls MySQL NDB Cluster nodes independently of the Windows Service Manager.

Assume that you have installed MySQL Cluster Manager to the default location for 32-bit Windows systems `C:\Program Files (x86)\MySQL\MySQL Cluster Manager` (`C:\Program Files\MySQL\MySQL Cluster Manager\` on 64-bit Windows systems), and that its configuration file is located in `C:\Program Files (x86)\MySQL\MySQL Cluster Manager\mcml.4.8\etc\`; then, the following command installs MySQL Cluster Manager as a service named `MCM`, with the display name “MySQL Cluster Manager 1.4.8”:

```
SC CREATE "MCM" DisplayName= "MySQL Cluster Manager 1.4.8" Start= "auto"
  BinPath= "C:\Program Files (x86)\MySQL\MySQL Cluster Manager\mcml.4.8\bin\mcmd-svc.exe
  --defaults-file=\"C:\Program Files (x86)\MySQL\MySQL Cluster Manager\mcml.4.8\etc\mcmd.ini\" "
```

This command can be quite long. For enhanced legibility, we have broken it across several lines, *but you should always enter it on a single line*, allowing it to wrap naturally, similar to what is shown here:

Figure 2.8 Installing the MySQL Cluster Manager Agent as a Windows Service: Command


```

Administrator: Command Prompt
(c) 2018 Microsoft Corporation. All rights reserved.

C:\WINDOWS\system32>SC CREATE "MCM" DisplayName= "MySQL Cluster Manager 1.4.8" Start=
"auto" BinPath= "C:\Program Files (x86)\MySQL\MySQL Cluster Manager\mcm1.4.8\bin\mcmd-
svc.exe --defaults-file="C:\Program Files (x86)\MySQL\MySQL Cluster Manager\mcm1.4.8\
etc\mcmd.ini\"
[SC] CreateService SUCCESS

```

In addition, you should keep in mind that the spaces after the equal signs following the `DisplayName`, `Start`, and `BinPath` arguments are required.

Starting and stopping the MySQL Cluster Manager agent Windows service. After installing the service successfully, you can start and stop the service manually, if the need arises, with the `SC START` and `SC STOP` commands.

```

C:\>SC START MCM
C:\>SC STOP MCM

```


Alternatively, use the `NET START` and `NET STOP` commands:

```

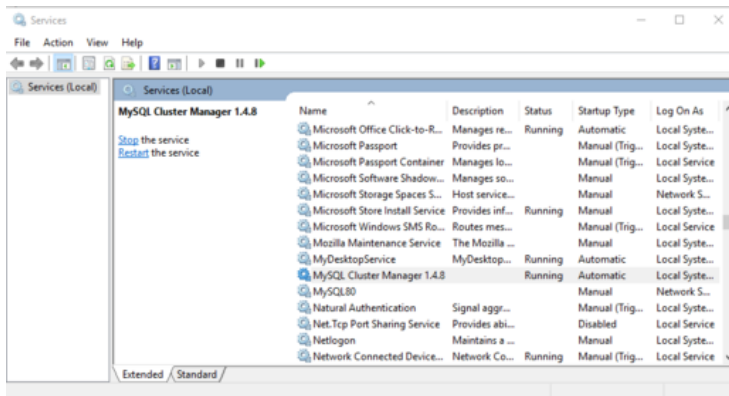
C:\Windows\system32>NET START MCM
C:\Windows\system32>NET STOP MCM

```

Once the service is installed, the MySQL Cluster Manager agent starts automatically whenever Windows is started. You can verify that the service is running with the Windows Task Manager. Open the Task Manager, and switch to the **Services** tab if it is not already displayed. If the MySQL Cluster Manager agent is running, you can find it in the list of services under `MCM` in the `Name` column and `MySQL Cluster Manager 1.4.8` in the `Description` column, as shown here:

Figure 2.9 Installing the MySQL Cluster Manager Agent as a Windows Service: Services Tab

You can also verify if the service is running using the Windows Service Manager, as shown here:

Figure 2.10 Installing the MySQL Cluster Manager Agent as a Windows Service: Service Manager

The Service Manager also allows you to start, stop, or pause the MySQL Cluster Manager agent service manually using a GUI.

**Note**

When first installing the MySQL Cluster Manager agent as a service, the service is not started automatically until Windows is started. If you do not wish to restart Windows, then you must start the service manually using either `SC START` or `NET START` on the command line or the graphical control provided in the Windows Service Manager.

You can remove the service using the `SC DELETE` command and the name of the service—in this case `MCM`—that was used in the `SC CREATE` command, as shown here:

Figure 2.11 Removing the MySQL Cluster Manager Agent from Windows Services

```
Administrator: Command Prompt
Microsoft Windows [Version 10.0.17763.973]
(c) 2018 Microsoft Corporation. All rights reserved.

C:\WINDOWS\system32>SC STOP MCM

SERVICE_NAME: MCM
        TYPE               : 10  WIN32_OWN_PROCESS
        STATE                : 3   STOP_PENDING
                        (STOPPABLE, NOT_PAUSABLE, ACCEPTS_SHUTDOWN)
        WIN32_EXIT_CODE      : 0    (0x0)
        SERVICE_EXIT_CODE  : 0    (0x0)
        CHECKPOINT          : 0x0
        WAIT_HINT           : 0x0

C:\WINDOWS\system32>SC DELETE MCM
[SC] DeleteService SUCCESS

C:\WINDOWS\system32>
```

If the service is running at the time that `SC DELETE` is executed, the removal of the service takes effect the next time the service is stopped. In such a case, you must stop the previous instance of the service manually, and allow it to be removed, before you can reinstall the service.

Once you have installed the MySQL Cluster Manager agent and the service is running correctly, you are ready to connect to it using the MySQL Cluster Manager client. See [Section 3.3, “Starting the MySQL Cluster Manager Client”](#), for information about how to do this.

2.3.3 Setting the MySQL Cluster Manager Agent User Name and Password

Normally it is not necessary to alter the user name or password used by the user agent to administer `mysqld` processes. However, if you should wish to do so, you can change either or both of these, using the procedure outlined here:

1. Stop all agents. (You can use `stop agents` for this purpose.)
2. Update the `etc/mcmd.ini` agent configuration file. Set a new password by uncommenting the line containing `manager-password=` and adding the new password as its value; set a new administrative user account name by uncommenting the line containing `manager-username=` and setting the value to the new user name. See [Section 3.1, “mcmd, the MySQL Cluster Manager Agent”](#), for more information about these options.
3. For each `mysqld` do the following:
 - a. Log in (using the `mysql` client) as the MySQL `root` user
 - b. If you are changing the user name, do this first, using the following statement, where `olduser` is the current user name and `newuser` is the new `manager-user` that you set previously in `etc/mcmd.ini`:


```
RENAME USER 'olduser'@'127.0.0.1' TO 'newuser'@'127.0.0.1';
```

If you are changing the user name for the first time, use `mcmd` for `olduser`. In addition, you should use `127.0.0.1` for the host name (and not `localhost`).
 - c. Execute the following statement, where `newuser` is the new user name, and `newpass` is the new password:


```
SET PASSWORD FOR 'newuser'@'127.0.0.1' = PASSWORD('newpass');
```

Use `mcmd` for the user name if you have not changed it—that is, if `manager-user` has been left unset in `mcmd.ini`. Use `127.0.0.1` for the host name (and not `localhost`).
 - d. Issue a `FLUSH PRIVILEGES` statement.
4. Restart the agents. All agents should now be using the new password for the `mcmd` accounts on the MySQL servers (`mysqld` processes).

2.4 MySQL Cluster Manager Configuration File

Before starting the MySQL Cluster Manager agent, you should make any necessary changes to the `[mcmd]` section of the agent configuration file, supplied with MySQL Cluster Manager as `etc/mcmd.ini` in the installation directory. On Linux and similar operating systems, you can edit this file in place; on Windows, it is recommended that you save a copy of this file to a convenient location for which the path does not contain any spaces, such as `C:\mcm\data`.

A `#` character in the MySQL Cluster Manager configuration file indicates the beginning of a comment; the comment continues to the end of the line.

The version of this file that is supplied with MySQL Cluster Manager reads similarly to what is shown here:

```
[mcmd]

plugins=manager

##manager-port=1862

##agent-uuid=
log-file=mcmd.log
log-level=message
##pid-file=

## Username and password for manager plugin
##manager-username=
##manager-password=

## Top-level directory for manager plugins information stored on disk
```

```
##manager-directory=
```

A minimal agent configuration file as used in production might look like this:

```
[mcmd]
plugins=manager
manager-port=1862
manager-directory=/var/opt/mcm-1.4.8
log-file=/var/opt/mcm/mcmd.log
log-level=warning
```

For more information about these and other options that can be set in the agent configuration file, see [Section 3.1, “mcmd, the MySQL Cluster Manager Agent”](#).

2.5 Upgrading MySQL Cluster Manager

This section discusses upgrading MySQL Cluster Manager from a previous release to the latest release (currently 1.4.8), as well as providing basic guidance on upgrading the bundled MySQL NDB Cluster software.



Note

Only upgrades from release 1.4.0 and later are supported.

The basic steps for upgrading a MySQL Cluster Manager installation are listed here:

1. Install the new version of the MySQL Cluster Manager software in the desired location.
2. Create a configuration for the new installation such that it uses the previous installation's data.
3. Stop all running MySQL Cluster Manager agent processes on all hosts.
4. Start the new agent processes, ensuring that they use the new configuration created previously.

A more detailed explanation is provided of each of these steps in the next few paragraphs. For purposes of example, we assume an upgrade from an existing installation of MySQL Cluster Manager 1.4.1 to a new installation of MySQL Cluster Manager 1.4.8. For a Linux or a Unix-like system, we assume that the installation directories are `/opt/mcm-1.4.1` and `/opt/mcm-1.4.8`, respectively; on Windows, we assume the installation directories are `C:\Program Files\MySQL\MySQL Cluster Manager 1.4.1\` and `C:\Program Files\MySQL\MySQL Cluster Manager 1.4.8\`.

Step 1: Install new MySQL Cluster Manager version. You can obtain and install a new version of MySQL Cluster Manager in the same way as for a new installation (see [Section 2.1, “Obtaining MySQL Cluster Manager”](#), and [Section 2.3, “MySQL Cluster Manager Installation”](#)), with the additional requirement that you should not attempt to install the new version in the same location as the version that you are currently using.

Step 2: Configure new installation. In order for the new MySQL Cluster Manager agent binaries to manage the same MySQL NDB Cluster instances, they must be able to find the data contained in the agent repository used by the old installation's binaries, which is `../mcm_data` (relative to the installation directory of MySQL Cluster Manager) by default, but can be set using the `manager-directory` directive in `mcmd.ini`.

It is simplest for MySQL Cluster Manager software upgrades if the agent repository and the agent configuration file are located externally to the agent installation directory. Suppose the old version of the agent is installed to `/opt/mcm-1.4.1`, and that it uses the directory `/var/opt/mcm` for its agent repository and `/etc/mcm/mcmd.ini` for its configuration file. In this case, to make the new binaries

use the same configuration and repository, it is necessary only to start the new `mcmd` processes such that they use the same configuration file (see [Step 4](#)). Otherwise, you should copy the default configuration file to a location outside the installation directory, such as `/etc/mcm/mcmd.ini`, as shown here:

```
$> cp -r /opt/mcm-1.4.1/etc/mcmd.ini /etc/mcm/mcmd.ini
```

Then, add the following line to the new copy of the `mcmd.ini` file:

```
manager-directory=/var/opt/mcm
```

After this, you can save and close the file. See also [Section 2.4, “MySQL Cluster Manager Configuration File”](#).

Step 3: Stop all agents. Stop the agent processes using the old binaries on all hosts making up the management installation. You can stop all agents for a given site, for example `mysite`, using the `stop agents` command in the MySQL Cluster Manager client, as shown here:

```
mcm> stop agents mysite;
```

You should execute a `stop agents` command, similar to the one just shown, for each site listed in the output of `list sites`.

Step 4: Start new MySQL Cluster Manager binaries. Start the new `mcmd` agent binaries with the `--defaults-file` option so that it uses the correct configuration file, like this:

```
$> mcmd --defaults-file=/etc/mcm/mcmd.ini &
```



Note

A majority of the agents (i.e., at least half of the total number plus one) should be started within a period of 10 seconds; otherwise, the lack of a quorum of nodes for decision making might cause the communication among the nodes to break down.

You should now be able to start the `mcm` client from the new installation and perform management tasks as usual. Once the client successfully starts and connects to the agent, you can verify that it is running the correct version of the MySQL Cluster Manager software using the `version` command, as shown here:

```
mcm> version;
+-----+
| Version |
+-----+
| MySQL Cluster Manager 1.4.8 |
+-----+
1 row in set (0.00 sec)
```

Next, check that all hosts, clusters, and processes on all sites are visible to the new `mcm` client, and are operational; for example:

```
mcm> list hosts mysite;
+-----+
| Host      | Status   | Version |
+-----+
| tonfisk   | Available | 1.4.8   |
| flundra   | Available | 1.4.8   |
| alpha     | Available | 1.4.8   |
| beta      | Available | 1.4.8   |
| gamma     | Available | 1.4.8   |
+-----+

5 rows in set (0.16 sec)

mcm> list clusters mysite;
+-----+
```

```

| Cluster      | Package      |
+-----+-----+
| mycluster   | mypackage   |
| yourcluster | mypackage   |
+-----+-----+
2 rows in set (2.07 sec)

mcm> show status --cluster mycluster;
+-----+-----+-----+
| Cluster | Status          | Comment |
+-----+-----+-----+
| mycluster | fully operational |         |
+-----+-----+-----+
1 row in set (0.01 sec)

mcm> show status --cluster yourcluster;
+-----+-----+-----+
| Cluster | Status          | Comment |
+-----+-----+-----+
| yourcluster | fully operational |         |
+-----+-----+-----+
1 row in set (0.01 sec)

mcm> show status -r mycluster;
+-----+-----+-----+-----+-----+-----+
| NodeId | Process | Host      | Status | Nodegroup | Package |
+-----+-----+-----+-----+-----+-----+
| 49     | ndb_mgmd | tonfisk   | running |           | mypackage |
| 1      | ndbd     | tonfisk   | running | 0         | mypackage |
| 2      | ndbd     | flundra   | running | 0         | mypackage |
| 50     | mysqld   | tonfisk   | running |           | mypackage |
| 51     | mysqld   | flundra   | running |           | mypackage |
| 52     | ndbapi   | *tonfisk  | added   |           |           |
| 53     | ndbapi   | *flundra  | added   |           |           |
+-----+-----+-----+-----+-----+-----+
7 rows in set (0.08 sec)

mcm> show status -r yourcluster;
+-----+-----+-----+-----+-----+-----+
| NodeId | Process | Host      | Status | Nodegroup | Package |
+-----+-----+-----+-----+-----+-----+
| 50     | ndb_mgmd | alpha     | running |           | mypackage |
| 5      | ndbd     | beta     | running | n/a       | mypackage |
| 6      | ndbd     | gamma    | running | n/a       | mypackage |
+-----+-----+-----+-----+-----+-----+
3 rows in set (0.01 sec)

```

See [Chapter 3, Using MySQL Cluster Manager](#), for more information about performing common cluster management tasks with the `mcm` client.

Upgrading MySQL NDB Cluster. Although the MySQL NDB Cluster software typically comes bundled with the MySQL Cluster Manager distribution, it is important to keep in mind that upgrading the MySQL Cluster Manager software does *not* upgrade any existing MySQL NDB Cluster installations. Since the new MySQL Cluster Manager installation uses the same configuration as the old one, the clusters under its control remain unchanged. If you wish to upgrade a cluster using the bundled MySQL NDB Cluster software, you should move the `cluster` directory (see [Contents of the MySQL Cluster Manager Unix Distribution Archive](#)) and all of its contents to a location outside the MySQL Cluster Manager installation directory. After this, you can use `add package` and `upgrade cluster` to upgrade one or more clusters to the new version of the MySQL NDB Cluster software.

Chapter 3 Using MySQL Cluster Manager

Table of Contents

3.1 <code>mcmd</code> , the MySQL Cluster Manager Agent	21
3.2 Starting and Stopping the MySQL Cluster Manager Agent	30
3.2.1 Starting and Stopping the Agent on Linux	30
3.2.2 Starting and Stopping the MySQL Cluster Manager Agent on Windows	31
3.3 Starting the MySQL Cluster Manager Client	33
3.4 Setting Up MySQL NDB Clusters with MySQL Cluster Manager	36
3.4.1 Creating a MySQL NDB Cluster with MySQL Cluster Manager	36
3.5 Importing MySQL NDB Clusters into MySQL Cluster Manager	38
3.5.1 Importing a Cluster Into MySQL Cluster Manager: Basic Procedure	38
3.5.2 Importing a Cluster Into MySQL Cluster Manager: Example	40
3.6 MySQL NDB Cluster Backup and Restore Using MySQL Cluster Manager	50
3.6.1 Requirements for Backup and Restore	50
3.6.2 Basic MySQL NDB Cluster Backup and Restore Using MySQL Cluster Manager	51
3.7 Backing Up and Restoring MySQL Cluster Manager Agents	57
3.8 Restoring a MySQL Cluster Manager Agent with Data from Other Agents	59
3.9 Setting Up MySQL NDB Cluster Replication with MySQL Cluster Manager	59
3.10 NDB Cluster Downgrades Involving Logical Downgrades of MySQL Nodes	62

This chapter discusses starting and stopping the MySQL Cluster Manager agent and client, and setting up, backing up, and restoring MySQL NDB Clusters using the MySQL Cluster Manager.

3.1 `mcmd`, the MySQL Cluster Manager Agent

`mcmd` is the MySQL Cluster Manager agent program; invoking this executable starts the MySQL Cluster Manager Agent, to which you can connect using the `mcm` client (see [Section 3.3, “Starting the MySQL Cluster Manager Client”](#), and [Chapter 4, *MySQL Cluster Manager Client Commands*](#), for more information).

You can modify the behavior of the agent in a number of different ways by specifying one or more of the options discussed in this sections. Most of these options can be specified either on the command line or in the agent configuration file (normally `etc/mcmd.ini`). (Some exceptions include the `--defaults-file` and `--bootstrap` options, which, if used, must be specified on the command line, and which are mutually exclusive with one another.) For example, you can set the agent's cluster logging level to `warning` instead than the default `message` in either one of the following two ways:

- Include `--log-level=warning` on the command line when invoking `mcmd`.



Note

When specifying an agent configuration option on the command line, the name of the option is prefixed with two leading dash characters (`--`).

- Include the following line in the agent configuration file:

```
log-level=warning
```



Note

You can change the logging level at runtime using the `mcm` client `change log-level` command.

When used in the configuration file, the name of the option should not be prefixed with any other characters. Each option must be specified on a separate line. You can comment out all of a given line by inserting a leading hash character (`#`), like this:

```
#log-level=warning
```

You can also comment out part of a line in this way; any text following the # character is ignored, to the end of the current line.

The following table contains a summary of agent options that are read on startup by `mcmd`. More detailed information about each of these options, such as allowed range of values, can be found in the list following the table.

Table 3.1 MySQL Cluster Manager Agent (mcmd) Option Summary

Option Name	Description	Introduced
<code>--agent-uuid</code>	Set the agent's UUID; needed only when running multiple agent processes on the same host	
<code>--basedir</code>	Directory to use as prefix for relative paths in the configuration	
<code>--bootstrap</code>	Bootstrap a default cluster on startup	
<code>--copy-port</code>	Specify the port for file copy operations	1.4.2
<code>--daemon</code>	Run in daemon mode. The option applies only to Linux and other Unix-like platforms	
<code>--defaults-file</code>	Configuration file to use	
<code>--event-threads</code>	Number of event handler threads to use	
<code>--help</code>	Show application options	
<code>--help-all</code>	Show all options (application options and manager module options)	
<code>--help-manager</code>	Show manager module options	
<code>--initial</code>	Wipes the contents in the agent's configuration store after making a backup	1.4.7
<code>--keepalive</code>	Try to restart mcmd in the event of a crash. The option applies only to Linux and other Unix-like platforms	
<code>--log-backtrace-on-crash</code>	Attempt to load debugger in case of a crash	
<code>--log-file</code>	Name of the file to write the log to	
<code>--log-level</code>	Set the mcmd logging level	
<code>--log-use-syslog</code>	Log to syslog	
<code>--manager-directory</code>	Directory used for manager data storage	
<code>--manager-password</code>	Password for the mcmd user-account	
<code>--manager-port</code>	Port for client to use when connecting to manager	

Option Name	Description	Introduced
<code>--manager-username</code>	Username for the <code>mcmd</code> user-account	
<code>--max-open-files</code>	Maximum number of open files (ulimit -n)	
<code>--pid-file</code>	Specify PID file (used if running as daemon)	
<code>--plugin-dir</code>	Directory in which to look for plugins	
<code>--plugins</code>	Comma-separated list of plugins to load; must include "manager"	
<code>--verbose-shutdown</code>	Always log the exit code when shutting down	
<code>--version</code>	Show the manager version	
<code>--xcom-port</code>	Specify the XCOM port	

MySQL Cluster Manager Agent (`mcmd`) Option Descriptions

The following list contains descriptions of each startup option available for use with `mcmd`, including allowed and default values. Options with their **Type** unmentioned need only be specified in order to take effect; you should not try to set a value for them.

- `--agent-uuid=uuid`

Command-Line Format	<code>--agent-uuid=uuid</code>
Type	String
Default Value	[set internally]

Set a UUID for this agent. Normally this value is set automatically, and needs to be specified only when running more than one `mcmd` process on the same host.

- `--basedir=dir_name`

Command-Line Format	<code>--basedir=dir_name</code>
Type	Directory name
Default Value	.

Directory with path to use as prefix for relative paths in the configuration.

- `--bootstrap`

Command-Line Format	<code>--bootstrap</code>
---------------------	--------------------------

Start the agent with default configuration values, create a default one-machine cluster named `mycluster`, and start it. This option works only if no clusters have yet been created.

- `--copy-port`

Command-Line Format	<code>--copy-port=#</code>
Introduced	1.4.2
Type	Numeric
Default Value	0
Minimum Value	0

Maximum Value	65535
---------------	-------

Allows you to specify the port for file copy operations. The default is 0.

- `--daemon`

Command-Line Format	<code>--daemon</code>
Platform Specific	Linux

Run `mcmd` as a daemon.

- `--defaults-file=filename`

Command-Line Format	<code>--defaults-file=file_name</code>
Type	File name
Default Value	<code>etc/mcmd.ini</code>

Set the file from which to read configuration options. The default is `etc/mcmd.ini`. See [Section 2.4, “MySQL Cluster Manager Configuration File”](#), for more information.

- `--event-threads=#`

Command-Line Format	<code>--event-threads=#</code>
Type	Numeric
Default Value	1
Minimum Value	1
Maximum Value	[system dependent]

Number of event handler threads to use. The default is 1, which is sufficient for most normal operations.

- `--help, -?`

Command-Line Format	<code>--help</code>
---------------------	---------------------

`mcmd` help output is divided into *Application* and *Manager* sections. When used with `mcmd`, `--help` causes the *Application* options to be shown, as shown here:

```
$> mcmd --help
Usage:
  mcmd [OPTION...] - MySQL Cluster Manager

Help Options:
  -?, --help                Show help options
  --help-all               Show all help options
  --help-manager            Show options for the manager-module

Application Options:
  -V, --version              Show version
  --defaults-file=<file>    configuration file
  --verbose-shutdown        Always log the exit code when shutting down
  --daemon                  Start in daemon-mode
  --basedir=<absolute path> Base directory to prepend to relative paths in the config
  --pid-file=<file>         PID file in case we are started as daemon
  --plugin-dir=<path>       Path to the plugins
  --plugins=<name>          Plugins to load
  --log-level=<string>      Log all messages of level ... or higher
  --log-file=<file>         Log all messages in a file
  --log-use-syslog          Log all messages to syslog
  --log-backtrace-on-crash  Try to invoke debugger on crash
  --keepalive               Try to restart mcmd if it crashed
```

<code>--max-open-files</code>	Maximum number of open files (ulimit -n)
<code>--event-threads</code>	Number of event-handling threads (default: 1)

- `--help-all`

Command-Line Format	<code>--help-all</code>
---------------------	-------------------------

`mcmd` help output is divided into *Application* and *Manager* sections. When used with `--help-all`, `mcmd` displays both the *Application* and the *Manager* options, like this:

```
$> mcmd --help-all
Usage:
  mcmd [OPTION...] - MySQL Cluster Manager

Help Options:
  -h, --help                Show help options
  --help-all               Show all help options
  --help-manager            Show options for the manager-module

manager-module
  --bootstrap                Bootstrap a cluster on localhost on initial startup
  --copy-port=<copy_port>   Port for file copy operations (default: 0)
  --initial                  Wipes the repository files after making a backup
  --manager-directory=<directory> Path to mcmd config information
  --manager-password=<password> Password for the mcmd user-account (default: super)
  --manager-port=<client_port> Port to contact the mcmd (default: 1862)
  --manager-username=<username> Username for the mcmd user-account (default: mcmd)
  --xcom-port=<xcom_port>    Xcom port for mcmds to communicate (default: 18620)

Application Options:
  -V, --version              Show version
  --defaults-file=<file>    configuration file
  --verbose-shutdown        Always log the exit code when shutting down
  --daemon                  Start in daemon-mode
  --basedir=<absolute path> Base directory to prepend to relative paths in the config
  --pid-file=<file>         PID file in case we are started as daemon
  --plugin-dir=<path>       Path to the plugins
  --plugins=<name>          Plugins to load
  --log-level=<string>      Log all messages of level ... or higher
  --log-file=<file>         Log all messages in a file
  --log-use-syslog           Log all messages to syslog
  --log-backtrace-on-crash  Try to invoke debugger on crash
  --keepalive               Try to restart mcmd if it crashed
  --max-open-files          Maximum number of open files (ulimit -n)
  --event-threads           Number of event-handling threads (default: 1)
```

- `--help-manager`

Command-Line Format	<code>--help-manager</code>
---------------------	-----------------------------

`mcmd` help output is divided into *Application* and *Manager* sections. When used with `--help-manager`, `mcmd` displays the *Manager* options, like this:

```
$> mcmd --help-manager
Usage:
  mcmd [OPTION...] - MySQL Cluster Manager

manager-module
  --bootstrap                Bootstrap a cluster on localhost on initial startup
  --copy-port=<copy_port>   Port for file copy operations (default: 0)
  --initial                  Wipes the repository files after making a backup
  --manager-directory=<directory> Path to mcmd config information
  --manager-password=<password> Password for the mcmd user-account (default: super)
  --manager-port=<client_port> Port to contact the mcmd (default: 1862)
  --manager-username=<username> Username for the mcmd user-account (default: mcmd)
  --xcom-port=<xcom_port>    Xcom port for mcmds to communicate (default: 18620)
```

- `--initial`

Command-Line Format	<code>--initial</code>
Introduced	1.4.7

For MySQL Cluster Manager 1.4.7 and later: Make a backup of the agent's configuration store (`mcm_data/rep/`) like the `backup agents` command would do for the local host, and then wipe the configuration store's contents before starting `mcmd`. The agent's configuration is then recovered from other agents. This is useful when an agent has fallen into an inconsistent state and cannot be properly restarted.

- `--keepalive`

Command-Line Format	<code>--keepalive</code>
Platform Specific	Linux

Use this option to cause `mcmd` to attempt to restart in the event of a crash.

- `--log-backtrace-on-crash`

Command-Line Format	<code>--log-backtrace-on-crash</code>
---------------------	---------------------------------------

Attempt to load the debugger in the event of a crash. Not normally used in production.

- `--log-file=filename`

Command-Line Format	<code>--log-file=file</code>
Type	File name
Default Value	<code>mcmd.log</code>

Set the name of the file to write the log to. The default is `mcmd.log` in the installation directory. On Linux and other Unix-like platforms, you can use a relative path; this is in relation to the MySQL Cluster Manager installation directory, and not to the `bin` or `etc` subdirectory. On Windows, you must use an absolute path, and it cannot contain any spaces; in addition, you must replace any backslash (`\`) characters in the path with forward slashes (`/`).

- `--log-level=level`

Command-Line Format	<code>--log-level=level</code>
Type	Enumeration
Default Value	<code>message</code>
Valid Values	<code>critical</code> <code>error</code> <code>warning</code> <code>message</code> <code>info</code> <code>debug</code>

Sets the `mcmd` log severity level. Possible values for this option and their descriptions are listed in Table 3.2, "MySQL Cluster Manager Agent Log Levels" in descending level of severity. When the option is set to a certain severity level, all events of that or higher levels are logged. `message` is the

default log level, and is the recommended setting for a production environment; running on a more severe log level produces fewer messages and makes it harder to trace a problem when it occurs.

Table 3.2 MySQL Cluster Manager Agent Log Levels

Level of Severity	Description
<code>critical</code>	Conditions that should be corrected immediately, such as a corrupted MySQL Cluster Manager data repository
<code>error</code>	Conditions that should be corrected, such as configuration errors
<code>warning</code>	Conditions that do not fail executions, but may require user attention
<code>message</code>	Messages on main events of the site and from command execution
<code>info</code>	Informational messages to provide users with some execution details
<code>debug</code>	Debugging messages that give execution details useful for developers. This causes large log files if used over a long period of time.

You can also change the `mcmd` logging level at runtime using the `change log-level` command with the `mcm` client. While the setting of the `--log-level` option is applied only to the host whose `mcmd` agent uses the option (either at the command line or in a configuration file), the `change log-level` client command can be used to apply the logging level to an entire management site or to specific hosts.

- `--log-use-syslog`

Command-Line Format	<code>--log-use-syslog</code>
---------------------	-------------------------------

Write logging output to syslog.

- `--manager-directory=dir_name`

Command-Line Format	<code>--manager-directory=dir</code>
Type	Directory name
Default Value	<code>../mcm_data</code> (relative to the installation directory of MySQL Cluster Manager)

Set the location of the agent repository, which contains collections of MySQL Cluster Manager data files and MySQL NDB Cluster configuration and data files. The value must be a valid absolute path. On Linux, if the directory does not exist, it is created; on Windows, the directory must be created if it does not exist. additionally on Windows, the path may not contain any spaces or backslash (`\`) characters; backslashes must be replaced with forward slashes (`/`).

The default location is `../mcm_data` (relative to the installation directory of MySQL Cluster Manager). If you change the default, you should use a standard location external to the MySQL Cluster Manager installation directory, such as `/var/opt/mcm` on Linux.

In addition to the data files for all the clusters under the control of MySQL Cluster Manager, the manager-directory also contains a `rep` directory in which `mcmd` configuration and metadata are kept. Normally, there is no need to interact with these directories beyond specifying the location of the `manager-directory` in the agent configuration file (`mcmd.ini`); see exceptions in, for

example, [Section 3.8, “Restoring a MySQL Cluster Manager Agent with Data from Other Agents”](#) and [Section 3.7, “Backing Up and Restoring MySQL Cluster Manager Agents”](#).

- `--manager-port=#`

Command-Line Format	<code>--manager-port=port</code>
Type	Numeric
Default Value	1862

Specify the port used by MySQL Cluster Manager client connections. Any valid TC/IP port number can be used. Normally, there is no need to change it from the default value (1862).

Previously, this option could optionally take a host name in addition to the port number, but in MySQL Cluster Manager 1.1.1 and later the host name is no longer accepted.

- `--manager-username=user_name`

The option serves the following two purposes:

- Sets the user name for an `mcm` client to connect to the `mcmd` agent. If the option is not specified, the default value of `mcmd` is used. If the option is specified with another value, the client must supply it using the `--user` client option when trying to connect to the agent.

The password for using this user name to connect to the agent is set with the `--manager-password` option.

- Sets a user name for the MySQL account to be used by the `mcmd` agent to access the SQL nodes. If the option is not specified, the default value of `mcmd` is used.

When an SQL node is initialized, the `mcmd` agent creates a new MySQL user account on it using the user name set by the option and the password set by the `--manager-password` option. *This account is created with all privileges on the MySQL server including the granting of privileges.* In other words, it is created as if you had executed `GRANT ALL PRIVILEGES ON *.* ... WITH GRANT OPTION` in the `mysql` client. The existing MySQL `root` account is not altered in such cases, and the default `test` database is preserved.

- `--manager-password=password`

The option serves the following two purposes:

- Sets the user password for an `mcm` client to connect to the `mcmd` agent with the user name set by `--manager-username`. If the option is not specified, the default value of `super` is used. If the option is specified with another value, the client must supply it using the `--password` client option when trying to connect to the agent.
- Sets a password for the MySQL account to be used by the `mcmd` agent to access the SQL nodes. If the option is not specified, the default value of `super` is used.

When an SQL node is initialized, the `mcmd` agent creates a new MySQL user account on it using the user name set by the `--manager-username` option and the password set by this option. See descriptions of `--manager-username` for more details about the account.

- `--max-open-files=#`

Command-Line Format	<code>--max-open-files=#</code>
Type	Numeric
Default Value	1
Minimum Value	1

Maximum Value	[system dependent]
---------------	--------------------

Set the maximum number of open files (as with `ulimit -n`).

- `--pid-file=file`

Command-Line Format	<code>--pid-file=file_name</code>
Type	File name
Default Value	<code>mcmd.pid</code>

Set the name and path to a process ID (`.pid`) file. Not normally used or needed. This option is not supported on Windows systems.

- `--plugin-dir`

Command-Line Format	<code>--plugin-dir=dir_name</code>
Type	Directory name
Default Value	<code>lib/mcmd</code>

Set the directory to search for plugins. The default is `lib/mcmd`, in the MySQL Cluster Manager installation directory; normally there is no need to change this.

- `--plugins`

Command-Line Format	<code>--plugins=list</code>
Type	Directory name
Default Value	<code>manager</code>

Specify a list of plugins to be loaded on startup. To enable MySQL Cluster Manager, this list must include `manager` (the default value). Please be aware that we currently do not test MySQL Cluster Manager with any values for `plugins` other than `manager`. Therefore, we recommend using the default value in a production setting.

- `--verbose-shutdown`

Command-Line Format	<code>--verbose-shutdown</code>
---------------------	---------------------------------

Force `mcmd` to log the exit code whenever shutting down, regardless of the reason.

- `--version, -V`

Command-Line Format	<code>--version</code>
---------------------	------------------------

Display version information and exit. Output may vary according to the MySQL Cluster Manager software version, operating platform, and versions of libraries used on your system, but should closely resemble what is shown here, with the first line of output containing the MySQL Cluster Manager release number (emphasized text):

```
$> mcmd --version
MySQL Cluster Manager 1.4.8 (64bit)
  chassis: 0.8.5.15734304
  glib2: 2.44.0
  libevent: 2.1.11-stable
-- modules
  manager: 1.4.8
```

- `--xcom-port`

Command-Line Format	<code>--xcom-port=#</code>
Type	Numeric
Default Value	18620
Minimum Value	1
Maximum Value	65535

Allows you to specify the XCOM port. The default is 18620.

3.2 Starting and Stopping the MySQL Cluster Manager Agent

Before you can start using MySQL Cluster Manager to create and manage a MySQL NDB Cluster, the MySQL Cluster Manager agent must be started on each computer that is intended to host one or more nodes in the MySQL NDB Cluster to be managed.

The MySQL Cluster Manager agent employs a MySQL user account for administrative access to `mysqld` processes. It is possible, but not a requirement, to change the default user name, the default password used for this account, or both. For more information, see [Section 2.3.3, “Setting the MySQL Cluster Manager Agent User Name and Password”](#).

3.2.1 Starting and Stopping the Agent on Linux

To start the MySQL Cluster Manager agent on a given host running a Linux or similar operating system, you should run `mcmd`, found in the `bin` directory within the manager installation directory on that host. Typical options used with `mcmd` are shown here:

```
mcmd [--defaults-file | --bootstrap] [--log-file] [--log-level]
```

See [Section 3.1, “mcmd, the MySQL Cluster Manager Agent”](#), for information about additional options that can be used when invoking `mcmd` from the command line, or in a configuration file.

`mcmd` normally runs in the foreground. If you wish, you can use your platform's usual mechanism for backgrounding a process. On a Linux system, you can do this by appending an ampersand character (`&`), like this (not including any options that might be required):

```
$> ./bin/mcmd &
```

By default, the agent assumes that the agent configuration file is `etc/mcmd.ini`, in the MySQL Cluster Manager installation directory. You can tell the agent to use a different configuration file by passing the path to this file to the `--defaults-file` option, as shown here:

```
$> ./bin/mcmd --defaults-file=/home/mcm/mcm-agent.conf
```

The `--bootstrap` option causes the agent to start with default configuration values, create a default one-machine cluster named `mycluster`, and start it. This option works only if no cluster has yet been created, and is mutually exclusive with the `--defaults-file` option. Currently, any data stored in the default cluster `mycluster` is not preserved between cluster restarts; this is a known issue, which we may address in a future release of MySQL Cluster Manager.

The use of the `--bootstrap` option with `mcmd` is shown here on a system having the host name `torsk`, where MySQL Cluster Manager has been installed to `/home/jon/mcm`:

```
$> ./mcmd --bootstrap
MySQL Cluster Manager 1.4.8 started
Connect to MySQL Cluster Manager by running "/home/jon/mcm/bin/mcm" -a torsk:1862
Configuring default cluster 'mycluster'...
Starting default cluster 'mycluster'...
Cluster 'mycluster' started successfully
      ndb_mgmd      torsk:1186
```

```

ndbd          torsk
ndbd          torsk
mysqld       torsk:3306
mysqld       torsk:3307
ndbapi       *
Connect to the database by running "/home/jon/mcm/cluster/bin/mysql" -h torsk -P 3306 -u root

```

You can then connect to the agent using the `mcm` client (see [Section 3.3, “Starting the MySQL Cluster Manager Client”](#)), and to either of the MySQL Servers running on ports 3306 and 3307 using `mysql` or another MySQL client application.

The `--log-file` option allows you to override the default location for the agent log file (normally `mcmd.log`, in the MySQL Cluster Manager installation directory).

You can use `--log-level` option to override the `log-level` set in the agent configuration file.

See [Section 3.1, “mcmd, the MySQL Cluster Manager Agent”](#), for more information about options that can be used with `mcmd`.

The MySQL Cluster Manager agent must be started on each host in the MySQL NDB Cluster to be managed.

To stop one or more instances of the MySQL Cluster Manager agent, use the `stop agents` command in the MySQL Cluster Manager client. If the client is unavailable, you can stop each agent process using the system's standard method for doing so, such as `^C` or `kill`.

You can also set the agent up as a daemon or service on Linux and other Unix-like systems. (See [Section 2.3.1, “Installing MySQL Cluster Manager on Unix Platforms”](#).) If you also want data node failed processes from a running MySQL NDB Cluster to be started when the agent fails and restarts in such cases, you must make sure that `StopOnError` is set to 0 on each data node (and not to 1, the default).

3.2.2 Starting and Stopping the MySQL Cluster Manager Agent on Windows

To start the MySQL Cluster Manager agent manually on a Windows host, you should invoke `mcmd.exe`, found in the `bin` directory under the manager installation directory on that host. By default, the agent uses `etc/mcmd.ini` in the MySQL Cluster Manager installation directory as its configuration file; this can be overridden by passing the desired file's location as the value of the `--defaults-file` option.

Typical options for `mcmd` are shown here:

```
mcmd[.exe] [--defaults-file | --bootstrap] [--log-file] [--log-level]
```

For information about additional options that can be used with `mcmd` on the command line or in an option file, see [Section 3.1, “mcmd, the MySQL Cluster Manager Agent”](#).

By default, the agent assumes that the agent configuration file is `etc/mcmd.ini`, in the MySQL Cluster Manager installation directory. You can tell the agent to use a different configuration file by passing the path to this file to the `--defaults-file` option, as shown here:

```

C:\Program Files (x86)\MySQL\MySQL Cluster Manager 1.4.8\bin>
mcmd --defaults-file="C:\Program Files (x86)\MySQL\MySQL Cluster Manager 1.4.8\etc\mcmd.ini"

```

The `--bootstrap` option causes the agent to start with default configuration values, create a default one-machine cluster named `mycluster`, and start it. The use of this option with `mcmd` is shown here on a system having the host name `torsk`, where MySQL Cluster Manager has been installed to the default location:

```

C:\Program Files (x86)\MySQL\MySQL Cluster Manager 1.4.8\bin>mcmd --bootstrap
MySQL Cluster Manager 1.4.8 started
Connect to MySQL Cluster Manager by running "C:\Program Files (x86)\MySQL\MySQL Cluster Manager 1.4.8\bin\mcmd"
Configuring default cluster 'mycluster'...

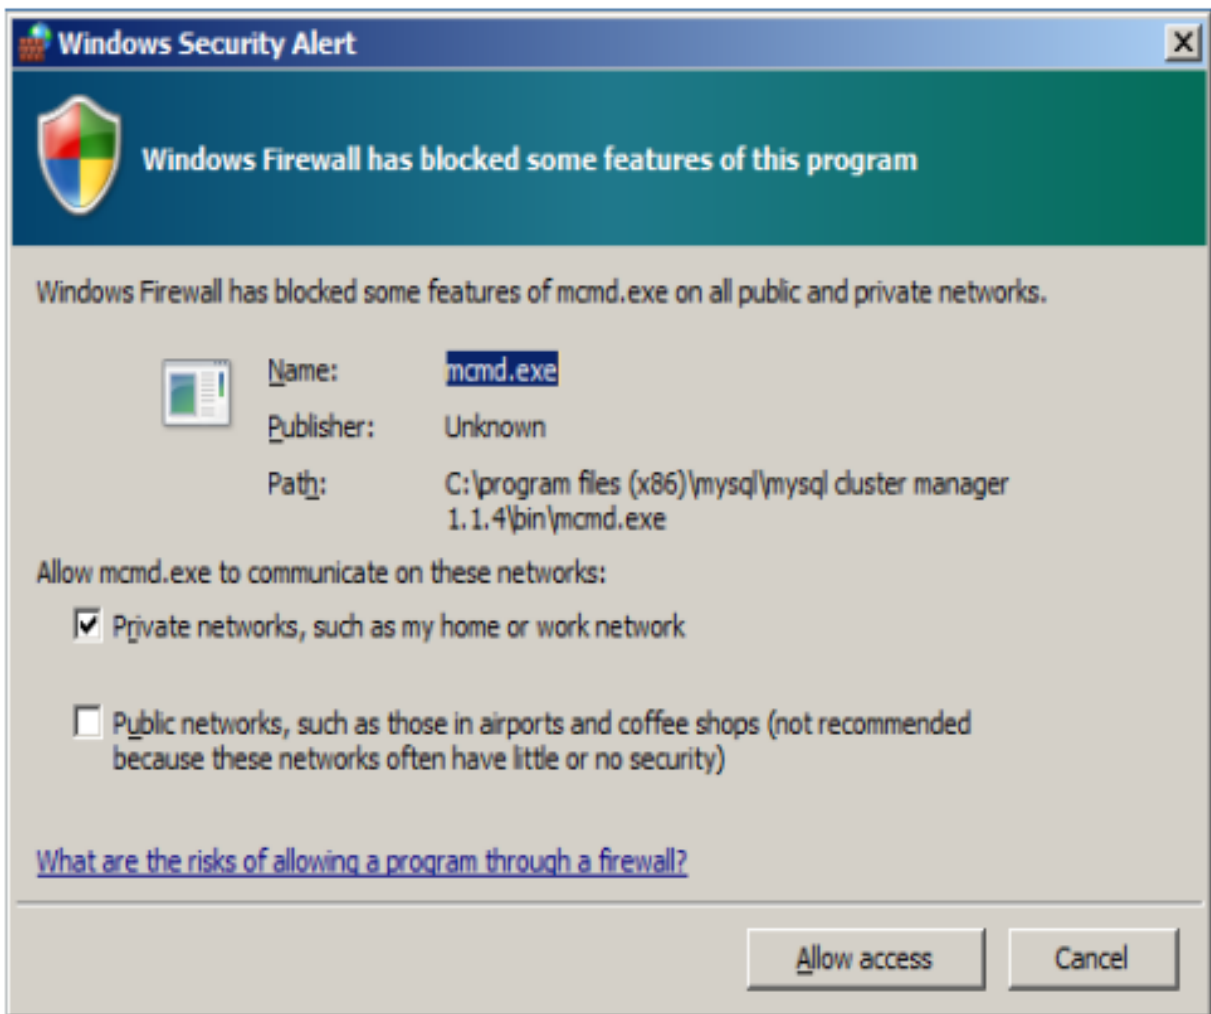
```

```
Starting default cluster 'mycluster'...
Cluster 'mycluster' started successfully
ndb_mgmd      TORSK:1186
ndbd          TORSK
ndbd          TORSK
mysqld        TORSK:3306
mysqld        TORSK:3307
ndbapi        *
Connect to the database by running "C:\Program Files (x86)\MySQL\MySQL Cluster Manager 1.4.8\cluster\bin\my
```

You can then connect to the agent using the `mcm` client (see [Section 3.3, “Starting the MySQL Cluster Manager Client”](#)), and to either of the MySQL Servers running on ports 3306 and 3307 using `mysql` or another MySQL client application.

When starting the MySQL Cluster Manager agent for the first time, you may see one or more Windows **Security Alert** dialogs, such as the one shown here:

Figure 3.1 Starting the MySQL Cluster Manager Agent on Windows: Security Alert



You should grant permission to connect to private networks for any of the programs `mcmd.exe`, `ndb_mgmd.exe`, `ndbd.exe`, `ndbmt.d.exe`, or `mysqld.exe`. To do so, check the **Private Networks...** box and then click the **Allow access** button. It is generally not necessary to grant MySQL Cluster Manager or MySQL NDB Cluster access to public networks such as the Internet.



Note

The `--defaults-file` and `--bootstrap` options are mutually exclusive.

The `--log-file` option allows you to override the default location for the agent log file (normally `mcmd.log`, in the MySQL Cluster Manager installation directory).

You can use `--log-level` option to override the `log-level` set in the agent configuration file.

See [Section 3.1, “mcmd, the MySQL Cluster Manager Agent”](#), for more information about options that can be used with `mcmd`.

The MySQL Cluster Manager agent must be started on each host in the MySQL NDB Cluster to be managed.

It is possible to install MySQL Cluster Manager as a Windows service, so that it is started automatically each time Windows starts. See [Section 2.3.2.1, “Installing the MySQL Cluster Manager Agent as a Windows Service”](#).

To stop one or more instances of the MySQL Cluster Manager agent, use the `stop agents` command in the MySQL Cluster Manager client. You can also stop an agent process using the Windows Task Manager. In addition, if you have installed MySQL Cluster Manager as a Windows service, you can stop (and start) the agent using the Windows Service Manager, `CTRL-C`, or the appropriate `SC STOP` (or `SC START`) or `NET STOP` (or `NET START`) command. See [Starting and stopping the MySQL Cluster Manager agent Windows service](#), for more information about each of these options.

3.3 Starting the MySQL Cluster Manager Client

This section covers starting the MySQL Cluster Manager client and connecting to the MySQL Cluster Manager agent.

MySQL Cluster Manager 1.4.8 includes a command-line client `mcm`, located in the installation `bin` directory. `mcm` can be invoked with any one of the options shown in the following table:

Table 3.3 mcm options

Long form	Short form	Description
<code>--help</code>	<code>-?</code>	Display <code>mcm</code> client options
<code>--version</code>	<code>-V</code>	Shows MySQL Cluster Manager agent/client version.
<code>—</code>	<code>-W</code>	Shows MySQL Cluster Manager agent/client version, with version of <code>mysql</code> used by <code>mcm</code> .
<code>--address</code>	<code>-a</code>	Host and optional port to use when connecting to <code>mcmd</code> , in <code>host[:port]</code> format; default is <code>127.0.0.1:1862</code> .
<code>--mysql-help</code>	<code>-I</code>	Show help for <code>mysql</code> client (see following).

The client/server protocol used by MySQL Cluster Manager is platform-independent. You can connect to any MySQL Cluster Manager agent with an `mcm` client on any platform where it is available. This means, for example, that you can use an `mcm` client on Microsoft Windows to connect to a MySQL Cluster Manager agent that is running on a Linux host.

`mcm` actually acts as a wrapper for the `mysql` client that is included with the bundled MySQL NDB Cluster distribution. Invoking `mcm` with no options specified is equivalent to the following:

```
$> mysql -umcmd -psuper -h 127.0.0.1 -P 1862 --prompt="mcm>"
```

(These `-u` and `-p` options and values are hard-coded and cannot be changed.) This means that you can use the `mysql` client to run MySQL Cluster Manager client sessions on platforms where `mcm` itself (or even `mcmd`) is not available. For more information, see [Connecting to the agent using the `mysql` client](#).

If you experience problems starting an MySQL Cluster Manager client session because the client fails to connect, see [Can't connect to \[local\] MySQL server](#), for some reasons why this might occur, as well as suggestions for some possible solutions.

To end a client session, use the `exit` or `quit` command (short form: `\q`). Neither of these commands requires a separator or terminator character.

For more information, see [Chapter 4, MySQL Cluster Manager Client Commands](#).

Connecting to the agent with the `mcm` client. You can connect to the MySQL Cluster Manager agent by invoking `mcm` (or, on Windows, `mcm.exe`). You may also need to specify one or more of the following command-line options:

- `--host=hostname` or `-h[]hostname`

This option takes the name or IP address of the host to connect to. The default is `localhost` (which may not be recognized on all platforms when starting a `mcm` client session even if it works for starting `mysql` client sessions).

You should keep in mind that the `mcm` client does not perform host name resolution; any name resolution information comes from the operating system on the host where the client is run. For this reason, it is usually best to use a numeric IP address rather than a hostname for this option.

- `--port=portnumber` or `-P[]portnumber`

This option specifies the TCP/IP port for the client to use. This must be the same port that is used by the MySQL Cluster Manager agent. As mentioned elsewhere, if no agent port is specified in the MySQL Cluster Manager agent configuration file (`mcmd.ini`), the default number of the port used by the MySQL Cluster Manager agent is 1862, which is also used by default by `mcm`.

- `--user=username` or `-u[]username`

The option specifies the user name for connecting to the agent. The default value of “`mcmd`” is used if the option is not specified. To connect successfully, the value of the option must match that specified by the `mcmd` configuration option `--manager-username` of the agent you are connecting to, which is also “`mcmd`” by default.

- `--password[=password]` or `-p[password]`

The option specifies the password for connecting to the agent. The default value of “`super`” is used if the option is not specified. To connect successfully, the value of the option must match that specified by the `mcmd` configuration option `--manager-password` of the agent you are connecting to, which is also “`super`” by default.

If you use the short option form (`-p`), you *must not* leave a space between this option and the password. If you omit the `password` value following the `--password` or `-p` option on the command line, the `mcm` client prompts you for one.

Specifying a password on the command line should be considered insecure. It is preferable that you either omit the password when invoking the client and then supply it when prompted, or put the password in a startup script or configuration file.

`mcm` accepts additional `mysql` client options, some of which may possibly be of use for MySQL Cluster Manager client sessions. For example, the `--pager` option might prove helpful when the output of `get` contains too many rows to fit in a single screen. The `--prompt` option can be used to provide

a distinctive prompt to help avoid confusion between multiple client sessions. However, options not shown in the current manual have not been extensively tested with `mcm` and so cannot be guaranteed to work correctly (or even at all). See [mysql Client Options](#), for a complete listing and descriptions of all `mysql` client options.



Note

Like the `mysql` client, `mcm` also supports `\G` as a statement terminator, which causes the output to be formatted vertically. This can be helpful when using a terminal whose width is restricted to some number of (typically 80) characters. See [Chapter 4, MySQL Cluster Manager Client Commands](#), for examples.

Connecting to the agent using the `mysql` client. As mentioned previously, `mcm` actually serves as a wrapper for the `mysql` client. In fact, a `mysql` client from any MySQL distribution should work without any issues for connecting to `mcmd`. In addition, since the client/server protocol used by MySQL Cluster Manager is platform-independent, you can use a `mysql` client on any platform supported by MySQL. (This means, for example, that you can use a `mysql` client on Microsoft Windows to connect to a MySQL Cluster Manager agent that is running on a Linux host.) Connecting to the MySQL Cluster Manager agent using the `mysql` client is accomplished by invoking `mysql` and specifying a hostname, port number, username and password using the following command-line options:

- `--host=hostname` or `-h hostname`

This option takes the name or IP address of the host to connect to. The default is `localhost`. Like the `mcm` client, the `mysql` client does not perform host name resolution, and relies on the host operating system for this task. For this reason, it is usually best to use a numeric IP address rather than a hostname for this option.

- `--port=portnumber` or `-P portnumber`

This option specifies the TCP/IP port for the client to use. This must be the same port that is used by the MySQL Cluster Manager agent. Although the default number of the port used by the MySQL Cluster Manager agent is 1862 (which is also used by default by `mcm`), *this default value is not known to the `mysql` client*, which uses port 3306 (the default port for the MySQL server) if this option is not specified when `mysql` is invoked.

Thus, you *must* use the `--port` or `-P` option to connect to the MySQL Cluster Manager agent using the `mysql` client, *even if the agent process is using the MySQL Cluster Manager default port*, and even if the agent process is running on the same host as the `mysql` client. Unless the correct agent port number is supplied to it on startup, `mysql` is unable to connect to the agent.

- `--user=username` or `-u username`

The option specifies the user name for connecting to the agent. By default, the `mysql` client tries to use the name of the current system user on Unix systems and “ODBC” on Windows, so you *must* supply this option and the username when trying to access the MySQL Cluster Manager agent with the `mysql` client; otherwise, `mysql` cannot connect to the agent.

To connect successfully, the value of the option must match that specified by the `mcmd` configuration option `--manager-username` of the agent you are connecting to, which is “`mcmd`” by default.

- `--password[=password]` or `-p[password]`

The option specifies the password for connecting to the agent. If you do not include the `--password` or `-p` option when invoking `mysql`, it cannot connect to the agent. To connect successfully, the value of the option must match that specified by the `mcmd` configuration option `--manager-password` of the agent you are connecting to, which is “`super`” by default.

If you use the short option form (`-p`), you *must not* leave a space between this option and the password. If you omit the `password` value following the `--password` or `-p` option on the command line, the `mysql` client prompts you for one.

Specifying a password on the command line should be considered insecure. It is preferable that you either omit the password when invoking the client and then supply it when prompted, or put the password in a startup script or configuration file.

In addition, you can use the `--prompt` option to set the `mysql` client's prompt. This is recommended, since allowing the default prompt (`mysql>`) to be used could lead to confusion between a MySQL Cluster Manager client session and a MySQL client session.

Thus, you can connect to a MySQL Cluster Manager agent by invoking the `mysql` client on the same machine from the system shell in a manner similar to what is shown here.

```
$> mysql -h127.0.0.1 -P1862 -umcmd -p --prompt='mcm> '
```

For convenience, on systems where `mcm` itself is not available, you might even want to put this invocation in a startup script. On a Linux or similar system, this script might be named `mcm-client.sh`, with contents similar to what is shown here:

```
#!/bin/sh
/usr/local/mysql/bin/mysql -h127.0.0.1 -P1862 -umcmd -p --prompt='mcm> '
```

In this case, you could then start up a MySQL Cluster Manager client session using something like this in the system shell:

```
$> ./mcm-client
```

On Windows, you can create a batch file with a name such as `mcm-client.bat` containing something like this:

```
C:\mysql\bin\mysql.exe -umcmd -psuper -h localhost -P 1862 --prompt="mcm> "
```

(Adjust the path to the `mysql.exe` client executable as necessary to match its location on your system.)

If you save this file to a convenient location such as the Windows desktop, you can start a MySQL Cluster Manager client session merely by double-clicking the corresponding file icon on the desktop (or in Windows Explorer); the client session opens in a new `cmd.exe` (DOS) window.

3.4 Setting Up MySQL NDB Clusters with MySQL Cluster Manager

This section provides basic information about setting up a new MySQL NDB Cluster with MySQL Cluster Manager. It also supplies guidance on migration of an existing MySQL NDB Cluster to MySQL Cluster Manager.

For more information about obtaining and installing the MySQL Cluster Manager agent and client software, see [Chapter 2, MySQL Cluster Manager Installation, Configuration, Cluster Setup](#).

See [Chapter 4, MySQL Cluster Manager Client Commands](#), for detailed information on the MySQL Cluster Manager client commands shown in this chapter.

3.4.1 Creating a MySQL NDB Cluster with MySQL Cluster Manager

In this section, we discuss the procedure for using MySQL Cluster Manager to create and start a new MySQL NDB Cluster. We assume that you have already obtained the MySQL Cluster Manager and MySQL NDB Cluster software, and that you are already familiar with installing MySQL Cluster Manager (see [Chapter 2, MySQL Cluster Manager Installation, Configuration, Cluster Setup](#)).

MySQL Cluster Manager also supports importing existing, standalone MySQL NDB Clusters; for more information, see [Section 3.5, "Importing MySQL NDB Clusters into MySQL Cluster Manager"](#).

We also assume that you have identified the hosts on which you plan to run the cluster and have decided on the types and distributions of the different types of nodes among these hosts, as well as basic configuration requirements based on these factors and the hardware characteristics of the host machines.

**Note**

You can create and start a MySQL NDB Cluster on a single host for testing or similar purposes, simply by invoking `mcmd` with the `--bootstrap` option. See [Section 3.2, “Starting and Stopping the MySQL Cluster Manager Agent”](#).

Creating a new cluster consists of the following tasks:

- **MySQL Cluster Manager agent installation and startup.** Install the MySQL Cluster Manager software distribution, make any necessary edits of the agent configuration files, and start the agent processes as explained in [Chapter 2, *MySQL Cluster Manager Installation, Configuration, Cluster Setup*](#). Agent processes must be running on all cluster hosts before you can create a cluster. This means that you need to place a complete copy of the MySQL Cluster Manager software distribution on every host. The MySQL Cluster Manager software does not have to be in a specific location, or even the same location on all hosts, but it must be present; you cannot manage any cluster processes hosted on a computer where `mcmd` is not also running.
- **MySQL Cluster Manager client session startup.** Starting the MySQL Cluster Manager client and connect to the MySQL Cluster Manager agent. You can connect to an agent process running on any of the cluster hosts, using the `mcm` client on any computer that can establish a network connection to the desired host. See [Section 3.3, “Starting the MySQL Cluster Manager Client”](#), for details.

On systems where `mcm` is not available, you can use the `mysql` client for this purpose. See [Connecting to the agent using the `mysql` client](#).

- **MySQL NDB Cluster software deployment.** The simplest and easiest way to do this is to copy the complete MySQL NDB Cluster distribution to the same location on every host in the cluster. (If you have installed MySQL Cluster Manager 1.4.8 on each host, the MySQL NDB Cluster 7.6.13 distribution is already included, in `mcm_installation_dir/cluster`.) If you do not use the same location on every host, be sure to note it for each host. Do not yet start any MySQL NDB Cluster processes or edit any configuration files; when creating a new cluster, MySQL Cluster Manager takes care of these tasks automatically.

On Windows hosts, you should *not* install as services any of the MySQL NDB Cluster node process programs, including `ndb_mgmd.exe`, `ndbd.exe`, `ndbmtd.exe`, and `mysqld.exe`. MySQL Cluster Manager manages MySQL NDB Cluster processes independently of the Windows Service Manager and does not interact with the Service Manager or any Windows services when doing so.

**Note**

You can actually perform this step at any time up to the point where the software package is registered (using `add package`). However, we recommend that you have all required software—including the MySQL NDB Cluster software—in place before executing any MySQL Cluster Manager client commands.

- **Management site definition.** Using the `create site` command in the MySQL Cluster Manager client, define a MySQL Cluster Manager management site—that is, the set of hosts to be managed. This command provides a name for the site, and must reference all hosts in the cluster. [Section 4.2.6, “The `create site` Command”](#), provides syntax and other information about this command. To verify that the site was created correctly, use the MySQL Cluster Manager client commands `list sites` and `list hosts`.
- **MySQL NDB Cluster software package registration.** In this step, you provide the location of the MySQL NDB Cluster software on all hosts in the cluster using one or more `add package` commands. To verify that the package was created correctly, use the `list packages` and `list processes` commands.
- **Cluster definition.** Execute a `create cluster` command to define the set of MySQL NDB Cluster nodes (processes) and hosts on which each cluster process runs, making up a the MySQL

NDB Cluster. This command also uses the name of the package registered in the previous step so that MySQL Cluster Manager knows the location of the binary running each cluster process. You can use the `list clusters` and `list processes` commands to determine whether the cluster has been defined as desired.

If you wish to use SQL node connection pooling, see [Setup for `mysqld` connection pooling](#) before creating the cluster.

- **Initial configuration.** Perform any configuration of the cluster that is required or desired prior to starting it. You can set values for MySQL Cluster Manager configuration attributes (MySQL NDB Cluster parameters and MySQL Server options) using the MySQL Cluster Manager client `set` command. You do not need to edit any configuration files directly—in fact, you should *not* do so. Keep in mind that certain attributes are read-only, and that some others cannot be reset after the cluster has been started for the first time. You can use the `get` command to verify that attributes have been set to the correct values.
- **Cluster startup.** Once you have completed the previous steps, including necessary or desired initial configuration, you are ready to start the cluster. The `start cluster` command starts all cluster processes in the correct order. You can verify that the cluster has started and is running normally after this command has completed, using the MySQL Cluster Manager client command `show status`. At this point, the cluster is ready for use by MySQL NDB Cluster applications.

3.5 Importing MySQL NDB Clusters into MySQL Cluster Manager

It is possible to bring a “wild” MySQL NDB Cluster—that is, a cluster not created using MySQL Cluster Manager—under the control of MySQL Cluster Manager. The following sections provide an outline of the procedure required to import such a cluster into MySQL Cluster Manager, followed by a more detailed example.

3.5.1 Importing a Cluster Into MySQL Cluster Manager: Basic Procedure

The importing process consists generally of the steps listed here:

1. Prepare the “wild” cluster for migration.
2. Verify PID files for cluster processes.
3. Create and configure in MySQL Cluster Manager a “target” cluster whose configuration matches that of the “wild” cluster.
4. Perform a test run, and then execute the `import cluster` command.

This expanded listing breaks down each of the tasks just mentioned into smaller steps:

1. *Prepare the “wild” cluster for migration*
 - a. It is highly recommended that you take a complete backup of the “wild” cluster before you make changes to it, using the `ndb_mgm` client. For more information, see [Using The NDB Cluster Management Client to Create a Backup](#).
 - b. Any cluster processes that are under the control of the system's boot-time process management facility, such as `/etc/init.d` on Linux systems or the Services Manager on Windows platforms, should be removed from its control.
 - c. The wild cluster's configuration must meet the following requirements, and it should be reconfigured and restarted if it does not:
 - `NodeID` must be assigned for every node.
 - `DataDir` must be specified for each management and data node, and the data directories for different nodes cannot overlap with each other.

- A “free” API node not bounded to any host must be provisioned, through which the `mcmd` agent can communicate with the cluster.
- d. Create a MySQL user named `mcmd` on each SQL node, and grant root privileges to the user.
 - e. Make sure that the configuration cache is disabled for each management node. Since the configuration cache is enabled by default, unless the management node has been started with the `--config-cache=false` option, you will need to stop and restart it with that option, in addition to other options that it has been started with previously.
 - f. *For MySQL Cluster Manager 1.4.6 and earlier:* Kill each data node angel process using your system's facility for doing so. Do not kill any non-angel data node daemons.

(This step is unnecessary for MySQL Cluster Manager 1.4.7 and later.)

2. *Verify cluster process PID files.*

- a. Verify that each process in the “wild” cluster has a valid PID file.
- b. If a given process does not have a valid PID file, you must create one for it.

See [Section 3.5.2.2, “Verify All Cluster Process PID Files”](#), for a more detailed explanation and examples.

3. *Create and configure “target” cluster under MySQL Cluster Manager control*

- a. [Install MySQL Cluster Manager](#) and start `mcmd` on all hosts with the same system user who started the wild cluster processes.
- b. Create a MySQL Cluster Manager site encompassing these hosts, using the `create site` command.
- c. Add a MySQL Cluster Manager package referencing the MySQL NDB Cluster binaries, using the `add package` command. Use this command's `--basedir` option to point to the location of the MySQL NDB Cluster installation directory.
- d. Create the target cluster using the `create cluster` command, including the same processes and hosts used by the wild cluster. Use the command's `--import` option to specify that the cluster is a target for import.

If the wild cluster adheres to the recommendation for node ID assignments given in the description for the `create cluster` command, you need not specify the node IDs for the processes in the `create cluster` command.

Also, this step may be split into a `create cluster` command followed by one or more `add process` commands (see [Section 3.5.2.3, “Creating and Configuring the Target Cluster”](#)).

- e. Use `import config` to copy the wild cluster's configuration data into the target cluster. Use this command's `--dryrun` option (short form: `-y`) to perform a test run that merely logs the configuration information the command copies when it is executed without the option.

If any `ndb_mgmd` or `mysqld` processes in the wild cluster are running on ports other than the default, you must first perform `set` commands to assign the correct port numbers for them in the target cluster. When all such processes are running on the correct ports and the dry run is successful, you can execute `import config` (without the `--dryrun` option) to copy the wild cluster's configuration data. Following this step, you should check the log as well as the configuration of the target cluster to ensure that all configuration attributes were copied correctly and with the correct scope. Correct any inconsistencies with the wild cluster's configuration using the appropriate `set` commands.

4. *Test and perform migration of wild cluster.*
 - a. Perform a test run of the proposed migration using `import cluster` with the `--dryrun` option, which causes MySQL Cluster Manager to check for errors, but not actually migrate any processes or data.
 - b. Correct any errors found using `--dryrun`. Repeat the dry run from the previous step to ensure that no errors were missed.
 - c. When the dry run no longer reports any errors, you can perform the migration using `import cluster`, but without the `--dryrun` option.

3.5.2 Importing a Cluster Into MySQL Cluster Manager: Example

As discussed previously (see [Section 3.5.1, “Importing a Cluster Into MySQL Cluster Manager: Basic Procedure”](#)), importing a standalone or “wild” cluster that was not created with MySQL Cluster Manager into the manager requires the completion of four major tasks. The example provided over the next few sections shows all the steps required to perform those tasks.

Sample cluster used in example. The “wild” cluster used in this example consists of four nodes—one management node, two data nodes, and one SQL node. Each of these nodes resides on one of three hosts, the IP address for each is shown in the following table:

Table 3.4 Nodes in the example cluster

Node type (executable)	Host name
Management node (<code>ndb_mgmd</code>)	198.51.100.102
Data node (<code>ndbd</code>)	198.51.100.103
Data node (<code>ndbd</code>)	198.51.100.104
SQL node (<code>mysqld</code>)	198.51.100.102

We assume that these hosts are on a dedicated network or subnet, and that each of them is running only the MySQL NDB Cluster binaries and applications providing required system and network services. We assume on each host that the MySQL NDB Cluster software has been installed from a release binary archive (see [Installing an NDB Cluster Binary Release on Linux](#)). We also assume that management node is using `/home/ari/bin/cluster/wild-cluster/config.ini` as the cluster's global configuration file, which is shown here:

```
[ndbd default]
NoOfReplicas= 2

[ndb_mgmd]
HostName= 198.51.100.102
DataDir= /home/ari/bin/cluster/wild-cluster/50/data
NodeId= 50

[ndbd]
HostName= 198.51.100.103
DataDir= /home/ari/bin/cluster/wild-cluster/2/data
NodeId=2

[ndbd]
HostName= 198.51.100.104
DataDir= /home/ari/bin/cluster/wild-cluster/3/data
NodeId=3

[mysqld]
HostName= 198.51.100.102
NodeId= 51

[api]
NodeId= 52
```

Notice that for the import into MySQL Cluster Manager to be successful, the following must be true for the cluster's configuration:

- `NodeID` must be explicitly assigned for every node.
- `DataDir` must be specified for each management and data node, and the data directories for different nodes cannot overlap with each other.
- A “free” API node not bounded to any host must be provisioned, through which the `mcmd` agent can communicate with the cluster.

3.5.2.1 Preparing the Standalone Cluster for Migration

The next step in the import process is to prepare the wild cluster for migration. This requires, among other things, removing cluster processes from control by any system service management facility, making sure all management nodes are running with configuration caching disabled, and, for MySQL Cluster Manager 1.4.6 and earlier, killing any data node angel processes that may be running. More detailed information about performing these tasks is provided in the remainder of this section.

- Before proceeding with any migration, the taking of a backup using the `ndb_mgm` client's `START BACKUP` command is strongly recommended.
- Any cluster processes that are under the control of a system boot process management facility such as `/etc/init.d` on Linux systems or the Services Manager on Windows platforms should be removed from this facility's control. Consult your operating system's documentation for information about how to do this. Be sure not to stop any running cluster processes in the course of doing so.
- Create a MySQL user account on each of the wild cluster's SQL nodes for MySQL Cluster Manager to execute the `import config` and `import cluster` commands in the steps to follow. The account name and password MySQL Cluster Manager uses to access MySQL nodes are specified by the `mcmd` client options `manager-username` and `manager-password` (the default values are `mcmd` and `super`, respectively); use those credentials when creating the account on the wild cluster's SQL nodes, and grant the user all privileges on the server, including the privilege to grant privileges. For example, log in to each of the wild cluster's SQL nodes with the `mysql` client as `root` and execute the SQL statements shown here:

```
CREATE USER 'mcmd'@'localhost' IDENTIFIED BY 'super';
GRANT ALL PRIVILEGES ON *.* TO 'mcmd'@'localhost' WITH GRANT OPTION;
```

Keep in mind that this must be done on all the SQL nodes, unless distributed privileges are enabled on the wild cluster.

- Make sure every node of the wild cluster has been started with its node ID specified with the `--ndb-nodeid` option at the command line, not just in the cluster configuration file. That is required for each process to be correctly identified by `mcmd` during the import. You can check if the requirement is fulfilled by the `ps -ef | grep` command, which shows the options the process has been started with:

```
$> ps -ef | grep ndb_mgmd
ari      8118      1  0 20:51 ?                00:00:04 /home/ari/bin/cluster/bin/ndb_mgmd --config-file=/h
--configdir=/home/ari/bin/cluster/wild-cluster --initial --ndb-nodeid=50
```

(For clarity's sake, in the command output for the `ps -ef | grep` command in this and the upcoming sections, we are skipping the line of output for the `grep` process itself.)

If the requirement is not fulfilled, restart the process with the `--ndb-nodeid` option; the restart can also be performed in step (e) or (f) below for any nodes you are restarting in those steps.

- Make sure that the configuration cache is disabled for each management node. Since the configuration cache is enabled by default, unless the management node has been started with the

`--config-cache=false` option, you will need to stop and restart it with that option, in addition to other options that it has been started with previously.

On Linux, we can once again use `ps` to obtain the information we need to accomplish this step. In a shell on host `198.51.100.102`, on which the management node resides:

```
$> ps -ef | grep ndb_mgmd
ari      8118      1  0 20:51 ?          00:00:04 /home/ari/bin/cluster/bin/ndb_mgmd --config-file=/home/ari/bin/cluster/wild-cluster --configdir=/home/ari/bin/cluster/wild-cluster --initial --ndb-nodeid=50
```

The process ID is 8118. The configuration cache is turned on by default, and a configuration directory has been specified using the `--configdir` option. First, terminate the management node using `kill` as shown here, with the process ID obtained from `ps` previously:

```
$> kill -15 8118
```

Verify that the management node process was stopped—it should no longer appear in the output of another `ps` command.

Now, restart the management node as described previously, with the configuration cache disabled and with the options that it was started with previously. Also, as already stated in step (d) above, make sure that the `--ndb-nodeid` option is specified at the restart:

```
$> /home/ari/bin/cluster/bin/ndb_mgmd --config-file=/home/ari/bin/cluster/wild-cluster/config.ini --config-cache=false
MySQL Cluster Management Server mysql-5.7.29-ndb-7.6.13
2016-11-08 21:29:43 [MgmtSrvr] INFO -- Skipping check of config directory since config cache is disabled
```



Caution

Do not use `0` or `OFF` for the value of the `--config-cache` option when restarting `ndb_mgmd` in this step. Using either of these values instead of `false` at this time causes the migration of the management node process to fail at later point in the import process.

Verify that the process is running as expected, using `ps`:

```
$> ps -ef | grep ndb_mgmd
ari      10221     1  0 19:38 ?          00:00:09 /home/ari/bin/cluster/bin/ndb_mgmd --config-file=/home/ari/bin/cluster/wild-cluster --config-cache=false
```

The management node is now ready for migration.



Important

While our example cluster has only a single management node, it is possible for a MySQL NDB Cluster to have more than one. In such cases, you must make sure the configuration cache is disabled for *each* management with the steps described in this step.

- f. *For MySQL Cluster Manager 1.4.6 and earlier:* Kill each data node angel process using the system's facility. The angel process monitors a data node process during a cluster's operation and, if necessary, attempts to restart the data node process (see [this FAQ](#) for details). Before a cluster can be imported, the angel processes must be stopped first. On a Linux system, you can identify angel processes by the output of the `ps -ef` command executed on the processes' host; here is an example of doing that on the host `198.51.100.103` of the sample cluster :

```
$> ps -ef | grep ndbd
ari      12836     1  0 20:52 ?          00:00:00 ./bin/ndbd --initial --ndb-nodeid=2 --ndb-connectstring=/home/ari/bin/cluster/wild-cluster
ari      12838 12836   2  0 20:52 ?          00:00:00 ./bin/ndbd --initial --ndb-nodeid=2 --ndb-connectstring=/home/ari/bin/cluster/wild-cluster
```

While both the actual data node process and its angel process appear as processes `ndbd`, you can identify each by looking at the process IDs. The process ID of the angel process (*italicized* in the sample output above) appears twice in the command output, once for itself (in the first line of the

output), and once as the ID of the parent process of the actual data node daemon (in the second line). Use the `kill` command to terminate the process with the identified process ID, like this:

```
$> kill -9 12836
```

Verify that the angel process has been killed and the other `ndbd` process (the non-angel data node daemon) is still running by issuing the `ps -ef` command again, as shown here:

```
$> ps -ef | grep ndbd
ari      12838      1  0 20:52 ?                00:00:02 ./bin/ndbd --initial --ndb-nodeid=2 --ndb-connectst
```

Now repeat this process in a shell on host `198.51.100.104`, as shown here:

```
$> ps -ef | grep ndbd
ari      11274      1  0 20:57 ?                00:00:00 ./cluster/bin/ndbd --initial --ndb-nodeid=3 --ndb-
ari      11276 11274  0 20:57 ?                00:00:01 ./cluster/bin/ndbd --initial --ndb-nodeid=3 --ndb-

$> kill -9 11274

$> ps -ef | grep ndbd
ari      11276      1  0 20:57 ?                00:00:01 ./cluster/bin/ndbd --initial --ndb-nodeid=3 --ndb-
```

For MySQL Cluster Manager 1.4.7 and later: There is no need kill the angel processes manually, as they will be taken care of when the `--remove-angel` option is used with the `import cluster` command at the last step of the import process.

The wild cluster's data nodes are now ready for migration.

3.5.2.2 Verify All Cluster Process PID Files

You must verify that each process in the wild cluster has a valid PID file. For purposes of this discussion, a valid PID file has the following characteristics:

- The file name is in the format of `ndb_node_id.pid`, where `node_id` is the node ID used for the process.
 - The file is located in the data directory used by the process.
 - The first line of the file contains the process ID of the node process, and only that process ID (this requirement changes when applied to data nodes for MySQL Cluster Manager 1.4.7 and later; see explanations below).
- a. To check the PID file for the management node process, log in to a system shell on host `198.51.100.102`, change to the management node's data directory as specified by the `Datadir` parameter in the cluster's configuration file, then check to see whether the PID file is present. On Linux, you can use the command shown here:

```
$> ls ndb_*.pid
ndb_50.pid
```

Check the content of the matching `.pid` file using a pager or text editor. We use `more` for this purpose here:

```
$> more ndb_50.pid
10221
```

The number shown should match the `ndb_mgmd` process ID. We can check this on Linux using the `ps` command:

```
$> ps -ef | grep ndb_mgmd
ari      10221      1  0 19:38 ?                00:00:09 /home/ari/bin/cluster/bin/ndb_mgmd --config-file=/h
```

The management node PID file satisfies the requirements listed at the beginning of this section.

- b. Next, we check the PID files for the data nodes, on hosts `198.51.100.103` and `198.51.100.104`. Log in to a system shell on `198.51.100.103`, then obtain the process ID of the `ndbd` process on this host, as shown here:

```
$> ps -ef | grep ndbd
ari      12838      1  0 Nov08 ?                00:10:12 ./bin/ndbd --initial --ndb-nodeid=2 --ndb-connectstring
```

As specified in the cluster's configuration file, the node's `DataDir` is `/home/ari/bin/cluster/wild-cluster/2/data`. Go to that directory to look for a file named `ndb_2.pid`:

```
$> ls ndb*.pid
ndb_2.pid
```

Now check the content of this file, and you are going to see the process ID for angel process for the data node (see [earlier instructions on stopping the angel process](#) for more explanations on the angel process):

```
$> more ndb_2.pid
12836
```

For MySQL Cluster Manager 1.4.6 and earlier: Change the number in the PID file to the data node's own PID:

```
$> sed -i 's/12836/12838/' ndb_2.pid
$> more ndb_2.pid
12838
```

Similarly, we locate and adjust the content of the PID file for the remaining data node (node ID 3, whose data directory is `/home/ari/bin/cluster/wild-cluster/3/data`) on host `198.51.100.104`:

```
$> ps -ef | grep ndbd
ari      11276      1  0 Nov09 ?                00:09:44 ./cluster//bin/ndbd --initial --ndb-nodeid=3 --ndb-connn

$> more /home/ari/bin/cluster/wild-cluster/3/data/ndb_3.pid
11274
```

Edit the `.pid` file, so it contains the data node process's own PID:

```
$> cd /home/ari/bin/cluster/wild-cluster/3/data/
$> sed -i 's/11274/11276/' ndb_3.pid
$> more ndb_3.pid
11276
```

The PID file for this data node also meets our requirements now.

For MySQL Cluster Manager 1.4.7 and later: There is no need to adjust the PID files to contain the data node processes' own PIDs, as that will be taken care of by the `--remove-angel` option used with the `import cluster` command at the last step of the import process. The data nodes are ready for import as long as they have valid PID files containing the PIDs for their angel processes.

We are ready to proceed to the `mysqld` node running on host `198.51.100.102`.

- c. To check the PID file for the `mysqld` node: the default location for it is the data directory of the node, specified by the `datadir` option in either a configuration file or at the command line at the start of the `mysqld` process. Let's go to the data directory `/home/ari/bin/cluster/wild-cluster/51/data` on host `198.51.100.104` and look for the PID file.

```
$> ls *.pid
```

```
localhost.pid
```

Notice that the MySQL Server could have been started with the `--pid-file` option, which puts a PID file at a specified location. In the following case, the same `mysqld` node has been started with the `mysqld_safe` script, and the `ps` command reveals the value for the `--pid-file` used:

```
$> ps -ef | grep mysqld
ari      11999  5667  0 13:15 pts/1    00:00:00 /bin/sh ./bin/mysqld_safe --defaults-file=/home/ari/...
ari      12136 11999  1 13:15 pts/1    00:00:00 /home/ari/bin/cluster/bin/mysqld --defaults-file=/h...
--basedir=/home/ari/bin/cluster/ --datadir=/home/ari/bin/cluster/wild-cluster/51/data/ --plugin-dir...
--ndb-nodeid=51 --log-error=/home/ari/bin/cluster/wild-cluster/51/data//localhost.localdomain.err...
--pid-file=/home/ari/bin/cluster/wild-cluster/51/data//localhost.localdomain.pid
```

As in the example, it is likely that you have a PID file that is not named in the required format for cluster import (`ndb_node_id.pid`); and if the `--pid-file` option was used, the PID file might not be at the required location (the data directory). Let us look into the PID file being referred to in the last example:

```
$> more /home/ari/bin/cluster/wild-cluster/51/data//localhost.localdomain.pid
12136
```

The PID file for the SQL node is at an acceptable location (inside the data directory) and has the correct contents (the right PID), but has the wrong name. Let us just copy the PID file into a correctly named file in the same directory, like this

```
$> cd /home/ari/bin/cluster/wild-cluster/51/data/
$> cp localhost.localdomain.pid ndb_51.pid
```

3.5.2.3 Creating and Configuring the Target Cluster

The next task is to create a “target” cluster. Once this is done, we modify the target cluster's configuration until it matches that of the wild cluster that we want to import. At a later point in the example, we also show how to test the configuration in a dry run before attempting to perform the actual configuration import.

To create and then configure the target cluster, follow these steps:

- a. Install MySQL Cluster Manager and start `mcmd` on all hosts *with the same system user who started the wild cluster processes*. Once you have done this, you can start the `mcm` client (see [Section 3.3, “Starting the MySQL Cluster Manager Client”](#)) on any one of these hosts to perform the next few steps.



Important

The cluster import is going to fail due to insufficient rights for the `mcmd` agents to perform their tasks if the `mcmd` agents are not started by the same system user who started the wild cluster processes.

- b. Create a MySQL Cluster Manager site encompassing all four of the wild cluster's hosts, using the `create site` command, as shown here:

```
mcm> create site --hosts=198.51.100.102,198.51.100.103,198.51.100.104 newsite;
+-----+
| Command result |
+-----+
| Site created successfully |
+-----+
1 row in set (0.15 sec)
```

We have named this site `newsite`. You should be able to see it listed in the output of the `list sites` command, similar to what is shown here:

```
mcm> list sites;
+-----+-----+-----+-----+
| Site   | Port | Local | Hosts |
+-----+-----+-----+-----+
```

```

+-----+-----+-----+-----+-----+
| newsite | 1862 | Local | 198.51.100.102,198.51.100.103,198.51.100.104 |
+-----+-----+-----+-----+-----+
1 row in set (0.01 sec)
    
```

- c. Add a MySQL Cluster Manager package referencing the MySQL NDB Cluster binaries using the `add package` command; use the command's `--basedir` option to point to the correct location of the MySQL NDB Cluster executables. The command shown here creates such a package, named `newpackage`:

```

mcm> add package --basedir=/home/ari/bin/cluster newpackage;
+-----+-----+
| Command result |
+-----+-----+
| Package added successfully |
+-----+-----+
1 row in set (0.70 sec)
    
```

You do not need to include the `bin` directory containing the MySQL NDB Cluster executables in the `--basedir` path. If the executables are in `/home/ari/bin/cluster/bin`, it is sufficient to specify `/home/ari/bin/cluster`; MySQL Cluster Manager automatically checks for the binaries in a `bin` directory within the directory specified by `--basedir`.

- d. Create the target cluster including at least some of the same processes and hosts used by the standalone cluster. *Do not include any processes or hosts that are not part of this cluster.* In order to prevent potentially disruptive process or cluster operations from interfering by accident with the import process, it is strongly recommended that you create the cluster for import using the `--import` option for the `create cluster` command.

You must also take care to preserve the correct node ID (as listed in the `config.ini` file shown previously) for each node.

The following command creates the cluster `newcluster` for import and includes the management and data nodes, but not the SQL or “free” API node (which we add in the next step):

```

mcm> create cluster --import --package=newpackage \
  --processhosts=ndb_mgmd:50@198.51.100.102,ndbd:2@198.51.100.103,ndbd:3@198.51.100.104 \
  newcluster;
+-----+-----+
| Command result |
+-----+-----+
| Cluster created successfully |
+-----+-----+
1 row in set (0.96 sec)
    
```

You can verify that the cluster was created correctly by checking the output of `show status` with the `--process (-r)` option, like this:

```

mcm> show status -r newcluster;
+-----+-----+-----+-----+-----+-----+
| NodeId | Process | Host | Status | Nodegroup | Package |
+-----+-----+-----+-----+-----+-----+
| 50 | ndb_mgmd | 198.51.100.102 | import | | newpackage |
| 2 | ndbd | 198.51.100.103 | import | n/a | newpackage |
| 3 | ndbd | 198.51.100.104 | import | n/a | newpackage |
+-----+-----+-----+-----+-----+-----+
3 rows in set (0.05 sec)
    
```

- e. If necessary, add any remaining processes and hosts from the wild cluster not included in the previous step using one or more `add process` commands. We have not yet accounted for two of the nodes from the wild cluster: the SQL node with node ID 51, on host `198.51.100.102`, and the API node with node ID 52, which is not bound to any specific host. You can use the following command to add both of these processes to `newcluster`:

```

mcm> add process --processhosts=mysqlid:51@198.51.100.102,ndbapi:52@* newcluster;
+-----+-----+
    
```

```

| Command result |
+-----+
| Process added successfully |
+-----+
1 row in set (0.41 sec)
    
```

Once again checking the output from `show status -r`, we see that the `mysqld` and `ndbapi` processes were added as expected:

```

mcm> show status -r newcluster;
+-----+-----+-----+-----+-----+-----+
| NodeId | Process | Host | Status | Nodegroup | Package |
+-----+-----+-----+-----+-----+-----+
| 50 | ndb_mgmd | 198.51.100.102 | import | n/a | newpackage |
| 2 | ndbd | 198.51.100.103 | import | n/a | newpackage |
| 3 | ndbd | 198.51.100.104 | import | n/a | newpackage |
| 51 | mysqld | 198.51.100.102 | import | n/a | newpackage |
| 52 | ndbapi | * | import | n/a | newpackage |
+-----+-----+-----+-----+-----+-----+
5 rows in set (0.06 sec)
    
```

You can also see that since `newcluster` was created using the `create cluster` command's `--import` option, the status of all processes in this cluster—including those we just added—is `import`. This means we cannot yet start `newcluster` or any of its processes. The `import` status and its effects on `newcluster` and its cluster processes persist until we have completed importing another cluster into `newcluster`.

The target `newcluster` cluster now has the same processes, with the same node IDs, and on the same hosts as the original standalone cluster. We are ready to proceed to the next step.

- f. Duplicate the wild cluster's configuration attributes in the target cluster using the `import config` command. Test out first the effects of the command by running it with the `--dryrun` option (the step only works if you have [created the mcmd user on the cluster's mysqld nodes](#)):



Important

Before executing this command it is necessary to set any non-default ports for `ndb_mgmd` and `mysqld` processes using the `set` command in the `mcm` client.

```

mcm> import config --dryrun newcluster;
+-----+-----+-----+-----+-----+-----+
| Command result |
+-----+-----+-----+-----+-----+-----+
| Import checks passed. Please check /home/ari/bin/mcm_data/clusters/newcluster/tmp/import_config.4 |
+-----+-----+-----+-----+-----+-----+
1 row in set (6.87 sec)
    
```

As indicated by the output from `import config --dryrun`, you can see the configuration attributes and values that would be copied to `newcluster` by the command without the `--dryrun` option in the file `/path-to-mcm-data-repository/clusters/clustername/tmp/import_config.message_id.mcm`. If you open this file in a text editor, you will see a series of `set` commands that would accomplish this task, similar to what is shown here:

```

# The following will be applied to the current cluster config:
set NoOfReplicas:ndbd=2 newcluster;
set DataDir:ndb_mgmd:50=/home/ari/bin/cluster/wild-cluster/50/data newcluster;
set DataDir:ndbd:2=/home/ari/bin/cluster/wild-cluster/2/data newcluster;
set DataDir:ndbd:3=/home/ari/bin/cluster/wild-cluster/3/data newcluster;
set basedir:mysqld:51=/home/ari/bin/cluster/ newcluster;
set datadir:mysqld:51=/home/ari/bin/cluster/wild-cluster/51/data/ newcluster;
set sql_mode:mysqld:51="NO_ENGINE_SUBSTITUTION,STRICT_TRANS_TABLES" newcluster;
    
```

```
set ndb_connectstring:mysqld:51=198.51.100.102 newcluster;
```

Options used at the command line instead of in a configuration file to start a node of the standalone cluster are not imported into the target cluster by the `import config` command; moreover, they will cause one of the following to happen when the `import config --dryrun` is run:

- i. For some options, MySQL Cluster Manager will issue a warning that “Option `<param>` may be removed on next restart of process `<type><nodeid>`,” meaning that those options will not be imported into the target cluster, and thus will not be applied when those nodes are restarted after the import. Here are the lists of such options for each node type:
 - For `ndb_mgmd` nodes: `---configdir`, `--initial`, `--log-name`, `--reload`, `--verbose`, `--core-file`
 - For `ndbd` and `ndbmtd` nodes: `--connect-retries`, `--connect-delay`, `--daemon=false`, `--nodaemon`, `--verbose`, `--core-file`
 - For `mysqld` nodes: `--ndbcluster`, the `--ndbinfo-*` options, `--verbose`, `--datadir`, `--defaults-group-suffix`, `--core-file`
- ii. For some other options, while their values will also not be imported into the target cluster, no warnings will be issued for them. Here are lists of such options for each node type:
 - For `ndb_mgmd` nodes: `--config-cache`, `--daemon`, `--ndb-nodeid`, `---nodaemon=false`, `--config-file`, `--skip-config-cache`
 - For `ndbd` and `ndbmtd` nodes: `--daemon`, `--foreground`, `--initial`, `--ndb-connectstring`, `--connect-string`, `--ndb-mgmd-host`, `--ndb-nodeid`, `---nodaemon=false`
 - For `mysqld` nodes: `--ndb-connectstring`, `--ndb-mgmd-host`, `--ndb-nodeid`, `---defaults-file`, `--no-defaults`, `--basedir`
- iii. For options that belong to neither of the two groups described above, having started the standalone cluster's nodes with them at the command line will cause the `import config --dryrun` command to fail with an error, complaining that the options are unsupported.

When you run into the first or third case described above, you have to do one of the following:

- If the options are required for the target cluster and they can be set using the `set` command (see [Command-line-only attributes](#)), set them for the target cluster using the `set` command, and then retry the `import config --dryrun` command.
- If the options are not needed for the target cluster, or it cannot be set using the `set` command, restart the wild cluster's node without those options, and then retry the `import config --dryrun` command.

After the successful dry run, you are now ready to import the wild cluster's configuration into `newcluster`, with the command shown here:

```
mcm> import config newcluster;
+-----+
| Command result
+-----+
| Configuration imported successfully. Please manually verify plugin options, abstraction level and def
+-----+
```

As an alternative, instead of importing all the settings using the `import config` command, you can make changes to the `/path-to-mcm-data-repository/clusters/clustername/tmp/`

`import_config.message_id.mcm` file generated by the dry run as you wish, and then import the settings by executing the file with the `mcm` agent:

```
mcm> source /path-to-mcm-data-repository/clusters/clustername/tmp/import_config.message_id.mcm
```

You should check the resulting configuration of `newcluster` carefully against the configuration of the wild cluster. If you find any inconsistencies, you must correct these in `newcluster` using the appropriate `set` commands.

3.5.2.4 Testing and Migrating the Standalone Cluster

Testing and performing the migration of a standalone MySQL NDB Cluster into MySQL Cluster Manager consists of the following steps:

1. Perform a test run of the proposed import using `import cluster` with the `--dryrun` option. When this option is used, MySQL Cluster Manager checks for mismatched configuration attributes, missing or invalid processes or hosts, missing or invalid PID files, and other errors, and warns of any it finds, without actually performing any migration of processes or data (the step only works if you have [created the mcmd user on the cluster's mysql nodes](#)):

```
mcm> import cluster --dryrun newcluster;
```

2. If errors occur, correct them, and repeat the dry run shown in the previous step until it returns no more errors. The following list contains some common errors you may encounter, and their likely causes:
 - MySQL Cluster Manager requires a specific MySQL user and privileges to manage SQL nodes. If the `mcmd` MySQL user account is not set up properly, you may see `No access for user...`, `Incorrect grants for user...`, or possibly other errors. Follow the instructions given in [this step](#) in [Section 3.5.2.1, "Preparing the Standalone Cluster for Migration"](#) to remedy the issue.
 - As described previously, each cluster process (other than a process whose type is `ndbapi`) being brought under MySQL Cluster Manager control must have a valid PID file. Missing, misnamed, or invalid PID files can produce errors such as `PID file does not exist for process...`, `PID ... is not running ...`, and `PID ... is type ...`. See [Section 3.5.2.2, "Verify All Cluster Process PID Files"](#).
 - Process version mismatches can also produce seemingly random errors whose cause can sometime prove difficult to track down. Ensure that all nodes are supplied with the correct release of the MySQL NDB Cluster software, and that it is the same release and version of the software.
 - Each data node angel process in the standalone cluster must be killed prior to import. A running angel process can cause errors such as `Angel process pid exists ...` or `Process pid is an angel process for ...`. Do the following when you see such errors:
 - For MySQL Cluster Manager 1.4.6 and earlier: See [this step](#) in [Section 3.5.2.1, "Preparing the Standalone Cluster for Migration"](#).
 - For MySQL Cluster Manager 1.4.7 and later: Proceed to the next step if these are the only errors you get. The angel processes and the data node PIDs will be taken care of by the `--remove-angel` option used with the `import cluster` command at the last step of the import process.
 - The number of processes, their types, and the hosts where they reside in the standalone cluster must be reflected accurately when creating the target site, package, and cluster for import. Otherwise, you may get errors such as `Process id reported # processes ...`, `Process id ... does not match configured process ...`, `Process id not configured ...`, and `Process id does not match configured process ...`. See [Section 3.5.2.3, "Creating and Configuring the Target Cluster"](#).

- Other factors that can cause specific errors include processes in the wrong state, processes that were started with unsupported command-line options (see [Section 3.5.2.3, “Creating and Configuring the Target Cluster”](#) for details) or without required options, and processes having the wrong process ID, or using the wrong node ID.
3. When `import cluster --dryrun` no longer warns of any errors, you can perform the import with the `import cluster` command, this time omitting the `--dryrun` option.

For MySQL Cluster Manager 1.4.6 and earlier:

```
mcm> import cluster newcluster;
+-----+
| Command result |
+-----+
| Cluster imported successfully |
+-----+
1 row in set (5.58 sec)
```

For MySQL Cluster Manager 1.4.7 and later: Use the `--remove-angel` option for the `import cluster` command, which kills the angel processes for the data nodes and adjusts the data nodes' PID files to contain the data node processes' own PIDs before importing the cluster:

```
mcm> import cluster --remove-angel newcluster;
+-----+
| Command result |
+-----+
| Cluster imported successfully |
+-----+
1 row in set (5.58 sec)
```

You can check that the wild cluster has now been imported, and is now under management of MySQL Cluster Manager:

```
mcm> show status -r newcluster;
+-----+-----+-----+-----+-----+-----+
| NodeId | Process | Host | Status | Nodegroup | Package |
+-----+-----+-----+-----+-----+-----+
| 50 | ndb_mgmd | 198.51.100.102 | running |  | newpackage |
| 2 | ndbd | 198.51.100.103 | running | 0 | newpackage |
| 3 | ndbd | 198.51.100.104 | running | 0 | newpackage |
| 51 | mysqld | 198.51.100.102 | running |  | newpackage |
| 52 | ndbapi | * | added |  |  |
+-----+-----+-----+-----+-----+-----+
5 rows in set (0.01 sec)
```

3.6 MySQL NDB Cluster Backup and Restore Using MySQL Cluster Manager

This section describes usage of the [NDB](#) native backup and restore functionality implemented in MySQL Cluster Manager, to perform a number of common tasks.

3.6.1 Requirements for Backup and Restore

This section provides information about basic requirements for performing backup and restore operations using MySQL Cluster Manager.

Requirements for MySQL NDB Cluster backup. Basic requirements for performing MySQL backups using MySQL Cluster Manager are minimal. At least one data node in each node group must be running, and there must be sufficient disk space on the node file systems. Partial backups are not supported.

Requirements for MySQL NDB Cluster restore. In general, the following requirements apply when you try to restore a MySQL NDB Cluster using MySQL Cluster Manager:

- A complete restore requires that all data nodes are up and running, and that all files belonging to a given backup are available.
- A partial restore is possible, but must be specified as such. This can be accomplished using the `restore cluster` client command with its `--skip-nodeid` option. See [Section 3.6.2.3, “Partial restore—missing images”](#) for details.
- In the event that data nodes have been added to the cluster since the backup was taken, only those data nodes for which backup files exist are restored. In such cases data is not automatically distributed to the new nodes, and, following the restore, you must redistribute the data manually by issuing an `ALTER ONLINE TABLE ... REORGANIZE PARTITION` statement in the `mysql` client for each NDB table in the cluster. See [Adding NDB Cluster Data Nodes Online: Basic procedure](#), for details.
- To restore a backup created by MySQL Cluster Manager to a cluster with fewer data nodes, you need to restore first the logical backup of the metadata of the NDB tables (available in backups created by MySQL Cluster Manager 1.4.1 and after) using the `mysqldump` utility and then restore the table data using the `ndb_restore` program. See [Section 3.6.2.5, “Restoring a Backup to a Cluster with Fewer Data Nodes”](#) for details.

3.6.2 Basic MySQL NDB Cluster Backup and Restore Using MySQL Cluster Manager

This section describes backing up and restoring a MySQL NDB Cluster, with examples of complete and partial restore operations. Note that the `backup cluster` and `restore cluster` commands work with NDB tables only; tables using other MySQL storage engines (such as InnoDB or MyISAM) are ignored.

For purposes of example, we use a MySQL NDB Cluster named `mycluster` whose processes and status can be seen here:

```
mcm> show status -r mycluster;
```

NodeId	Process	Host	Status	Nodegroup	Package
49	ndb_mgmd	tonfisk	running		mypackage
1	ndbd	tonfisk	running	0	mypackage
2	ndbd	tonfisk	running	0	mypackage
50	mysqld	tonfisk	running		mypackage
51	mysqld	tonfisk	running		mypackage
52	ndbapi	*tonfisk	added		
53	ndbapi	*tonfisk	added		

7 rows in set (0.08 sec)

You can see whether there are any existing backups of `mycluster` using the `list backups` command, as shown here:

```
mcm> list backups mycluster;
```

BackupId	NodeId	Host	Timestamp	Comment
1	1	tonfisk	2012-12-04 12:03:52	
1	2	tonfisk	2012-12-04 12:03:52	
2	1	tonfisk	2012-12-04 12:04:15	
2	2	tonfisk	2012-12-04 12:04:15	
3	1	tonfisk	2012-12-04 12:17:41	
3	2	tonfisk	2012-12-04 12:17:41	

6 rows in set (0.12 sec)

3.6.2.1 Simple backup

To create a backup, use the `backup cluster` command with the name of the cluster as an argument, similar to what is shown here:

```
mcm> backup cluster mycluster;
+-----+
| Command result |
+-----+
| Backup completed successfully |
+-----+
1 row in set (3.31 sec)
```

`backup cluster` requires only the name of the cluster to be backed up as an argument; for information about additional options supported by this command, see [Section 4.7.2, “The backup cluster Command”](#). To verify that a new backup of `mycluster` was created with a unique ID, check the output of `list backups`, as shown here (where the rows corresponding to the new backup files are indicated with emphasized text):

```
mcm> list backups mycluster;
+-----+-----+-----+-----+-----+
| BackupId | NodeId | Host      | Timestamp                | Comment |
+-----+-----+-----+-----+-----+
| 1         | 1       | tonfisk   | 2012-12-04 12:03:52      |         |
| 1         | 2       | tonfisk   | 2012-12-04 12:03:52      |         |
| 2         | 1       | tonfisk   | 2012-12-04 12:04:15      |         |
| 2         | 2       | tonfisk   | 2012-12-04 12:04:15      |         |
| 3         | 1       | tonfisk   | 2012-12-04 12:17:41      |         |
| 3         | 2       | tonfisk   | 2012-12-04 12:17:41      |         |
| 4         | 1       | tonfisk   | 2012-12-12 14:24:35      |         |
| 4         | 2       | tonfisk   | 2012-12-12 14:24:35      |         |
+-----+-----+-----+-----+-----+
8 rows in set (0.04 sec)
```

If you attempt to create a backup of a MySQL NDB Cluster in which each node group does not have at least one data node running, `backup cluster` fails with the error `Backup cannot be performed as processes are stopped in cluster cluster_name`.

3.6.2.2 Simple complete restore

To perform a complete restore of a MySQL NDB Cluster from a backup with a given ID, follow the steps listed here:

1. Identify the backup to be used.

In this example, we use the backup having the ID 4, that was created for `mycluster` previously in this section.

2. Wipe the MySQL NDB Cluster data.

The simplest way to do this is to stop and then perform an initial start of the cluster as shown here, using `mycluster`:

```
mcm> stop cluster mycluster;
+-----+
| Command result |
+-----+
| Cluster stopped successfully |
+-----+
1 row in set (15.24 sec)

mcm> start cluster --initial mycluster;
+-----+
| Command result |
+-----+
| Cluster started successfully |
+-----+
1 row in set (34.47 sec)
```

3. Restore the backup.

This is done using the `restore cluster` command, which requires the backup ID and the name of the cluster as arguments. Thus, you can restore backup 4 to `mycluster` as shown here:

```
mcm> restore cluster --backupid=4 mycluster;
+-----+
| Command result |
+-----+
| Restore completed successfully |
+-----+
1 row in set (16.78 sec)
```

3.6.2.3 Partial restore—missing images

It is possible using MySQL Cluster Manager to perform a partial restore of a MySQL NDB Cluster—that is, to restore from a backup in which backup images from one or more data nodes are not available. This is required if we wish to restore `mycluster` to backup number 6, since an image for this backup is available only for node 1, as can be seen in the output of `list backups` in the `mcm` client :

```
mcm> list backups mycluster;
+-----+-----+-----+-----+-----+
| BackupId | NodeId | Host      | Timestamp          | Comment |
+-----+-----+-----+-----+-----+
| 1         | 1       | tonfisk   | 2012-12-04 12:03:52 |         |
| 1         | 2       | tonfisk   | 2012-12-04 12:03:52 |         |
| 2         | 1       | tonfisk   | 2012-12-04 12:04:15 |         |
| 2         | 2       | tonfisk   | 2012-12-04 12:04:15 |         |
| 3         | 1       | tonfisk   | 2012-12-04 12:17:41 |         |
| 3         | 2       | tonfisk   | 2012-12-04 12:17:41 |         |
| 4         | 1       | tonfisk   | 2012-12-12 14:24:35 |         |
| 4         | 2       | tonfisk   | 2012-12-12 14:24:35 |         |
| 5         | 1       | tonfisk   | 2012-12-12 14:31:31 |         |
| 5         | 2       | tonfisk   | 2012-12-12 14:31:31 |         |
| 6         | 1       | tonfisk   | 2012-12-12 14:32:09 |         |
+-----+-----+-----+-----+-----+
11 rows in set (0.08 sec)
```

To perform a restore of only those nodes for which we have images (in this case, node 1 only), we can use the `--skip-nodeid` option when executing a `restore cluster` command. This option causes one or more nodes to be skipped when performing the restore. Assuming that `mycluster` has been cleared of data (as described earlier in this section), we can perform a restore that skips node 2 as shown here:

```
mcm> restore cluster --backupid=6 --skip-nodeid=2 mycluster;
+-----+
| Command result |
+-----+
| Restore completed successfully |
+-----+
1 row in set (17.06 sec)
```

Because we excluded node 2 from the restore process, no data has been distributed to it. To cause MySQL NDB Cluster data to be distributed to any such excluded or skipped nodes following a partial restore, it is necessary to redistribute the data manually by executing an `ALTER ONLINE TABLE ... REORGANIZE PARTITION` statement in the `mysql` client for each NDB table in the cluster. To obtain a list of NDB tables from the `mysql` client, you can use multiple `SHOW TABLES` statements or a query such as this one:

```
SELECT CONCAT(' ' TABLE_SCHEMA, '. ', TABLE_NAME)
FROM INFORMATION_SCHEMA.TABLES
WHERE ENGINE='ndbcluster';
```

You can generate the necessary SQL statements using a more elaborate version of the query just shown, such the one employed here:

```
mysql> SELECT
->   CONCAT('ALTER ONLINE TABLE `', TABLE_SCHEMA,
->         '`', TABLE_NAME, ' REORGANIZE PARTITION;')
->   AS Statement
-> FROM INFORMATION_SCHEMA.TABLES
-> WHERE ENGINE='ndbcluster';
+-----+
```

```

| Statement
+-----+
| ALTER ONLINE TABLE `mysql`.`ndb_apply_status` REORGANIZE PARTITION;
| ALTER ONLINE TABLE `mysql`.`ndb_index_stat_head` REORGANIZE PARTITION;
| ALTER ONLINE TABLE `mysql`.`ndb_index_stat_sample` REORGANIZE PARTITION;
| ALTER ONLINE TABLE `db1`.`n1` REORGANIZE PARTITION;
| ALTER ONLINE TABLE `db1`.`n2` REORGANIZE PARTITION;
| ALTER ONLINE TABLE `db1`.`n3` REORGANIZE PARTITION;
| ALTER ONLINE TABLE `test`.`n1` REORGANIZE PARTITION;
| ALTER ONLINE TABLE `test`.`n2` REORGANIZE PARTITION;
| ALTER ONLINE TABLE `test`.`n3` REORGANIZE PARTITION;
| ALTER ONLINE TABLE `test`.`n4` REORGANIZE PARTITION;
+-----+
10 rows in set (0.09 sec)

```

3.6.2.4 Partial restore—data nodes added

A partial restore can also be performed when new data nodes have been added to a MySQL NDB Cluster following a backup. In this case, you can exclude the new nodes using `--skip-nodeid` when executing the `restore cluster` command. Consider the MySQL NDB Cluster named `mycluster` as shown in the output of the following `show status` command:

```

mcm> show status -r mycluster;
+-----+-----+-----+-----+-----+-----+
| NodeId | Process | Host      | Status | Nodegroup | Package |
+-----+-----+-----+-----+-----+-----+
| 49     | ndb_mgmd | tonfisk   | stopped |           | mypackage |
| 1      | ndbd     | tonfisk   | stopped | 0         | mypackage |
| 2      | ndbd     | tonfisk   | stopped | 0         | mypackage |
| 50     | mysqld   | tonfisk   | stopped |           | mypackage |
| 51     | mysqld   | tonfisk   | stopped |           | mypackage |
| 52     | ndbapi   | *tonfisk  | added  |           |           |
| 53     | ndbapi   | *tonfisk  | added  |           |           |
+-----+-----+-----+-----+-----+-----+
7 rows in set (0.03 sec)

```

The output of `list backups` shows us the available backup images for this cluster:

```

mcm> list backups mycluster;
+-----+-----+-----+-----+-----+
| BackupId | NodeId | Host      | Timestamp          | Comment |
+-----+-----+-----+-----+-----+
| 1         | 1      | tonfisk   | 2012-12-04 12:03:52 |         |
| 1         | 2      | tonfisk   | 2012-12-04 12:03:52 |         |
| 2         | 1      | tonfisk   | 2012-12-04 12:04:15 |         |
| 2         | 2      | tonfisk   | 2012-12-04 12:04:15 |         |
| 3         | 1      | tonfisk   | 2012-12-04 12:17:41 |         |
| 3         | 2      | tonfisk   | 2012-12-04 12:17:41 |         |
| 4         | 1      | tonfisk   | 2012-12-12 14:24:35 |         |
| 4         | 2      | tonfisk   | 2012-12-12 14:24:35 |         |
+-----+-----+-----+-----+-----+
8 rows in set (0.06 sec)

```

Now suppose that, at a later point in time, 2 data nodes have been added to `mycluster` using an `add process` command. The `show status` output for `mycluster` now looks like this:

```

mcm> show status -r mycluster;
+-----+-----+-----+-----+-----+-----+
| NodeId | Process | Host      | Status | Nodegroup | Package |
+-----+-----+-----+-----+-----+-----+
| 49     | ndb_mgmd | tonfisk   | running |           | mypackage |
| 1      | ndbd     | tonfisk   | running | 0         | mypackage |
| 2      | ndbd     | tonfisk   | running | 0         | mypackage |
| 50     | mysqld   | tonfisk   | running |           | mypackage |
| 51     | mysqld   | tonfisk   | running |           | mypackage |
| 52     | ndbapi   | *tonfisk  | added  |           |           |
| 53     | ndbapi   | *tonfisk  | added  |           |           |
| 3      | ndbd     | tonfisk   | running | 1         | mypackage |
| 4      | ndbd     | tonfisk   | running | 1         | mypackage |
+-----+-----+-----+-----+-----+-----+
9 rows in set (0.01 sec)

```

Since nodes 3 and 4 were not included in the backup, we need to exclude them when performing the restore. You can cause `restore cluster` to skip multiple data nodes by specifying a comma-separated list of node IDs with the `--skip-nodeid` option. Assume that we have just cleared `mycluster` of MySQL NDB Cluster data using the `mcm` client commands `stop cluster` and `start cluster --initial` as described previously in this section; then we can restore `mycluster` (now having 4 data nodes numbered 1, 2, 3, and 4) from backup number 4 (made when `mycluster` had only 2 data nodes numbered 1 and 2) as shown here:

```
mcm> restore cluster --backupid=4 --skip-nodeid=3,4 mycluster;
+-----+
| Command result |
+-----+
| Restore completed successfully |
+-----+
1 row in set (17.61 sec)
```

No data is distributed to the skipped (new) nodes; you must force nodes 3 and 4 to be included in a redistribution of the data using `ALTER ONLINE TABLE ... REORGANIZE PARTITION` as described previously in this section.

For MySQL Cluster Manager 1.4.1 and later: An alternative to generating and running the `ALTER ONLINE TABLE ... REORGANIZE PARTITION` steps is to make use of the logical backup of the NDB tables' metadata, which is part of the cluster backup created by MySQL Cluster Manager. To do this, before you run the `restore cluster` step outlined above:

1. Locate the logical backup for the metadata; see [Locations of backup files](#) in [Section 3.6.2.5, "Restoring a Backup to a Cluster with Fewer Data Nodes"](#) for instructions.
2. Restore the logical backup; see [Restore the Logical Backup of NDB Table Metadata \[57\]](#) in [Section 3.6.2.5, "Restoring a Backup to a Cluster with Fewer Data Nodes"](#) for instructions.

You can then run the `restore cluster` step, and the data is going to be redistributed across all the data nodes, without the need for further manual intervention.

3.6.2.5 Restoring a Backup to a Cluster with Fewer Data Nodes

Sometimes, you want to transfer data from your cluster to another one that has fewer data nodes—for example, when you want to scale down your cluster or prepare a smaller replica cluster for a replication setup. While the methods described in [Section 3.6.2, "Basic MySQL NDB Cluster Backup and Restore Using MySQL Cluster Manager"](#) will not work in that case, starting from MySQL Cluster Manager 1.4.1, you can perform the transfer by just using the `backup cluster` command and the `ndb_restore` program.

The process starts with creating a backup for the original cluster using the `backup cluster` command. Next, create a new cluster with fewer data nodes using the `create cluster` command. Before the NDB table data can be transferred, the metadata for the NDB tables must first be restored to the new cluster. Starting from MySQL Cluster Manager 1.4.1, the `backup cluster` command also creates a logical backup for the metadata of the NDB tables (see [Logical Backup for NDB Table Metadata](#) for details). Use the `--all` option with the `list backups` command to list all backups, including the logical backups for the NDB tables' metadata, which are marked by the comment "Schema":

```
mcm> list backups --all mycluster;
+-----+-----+-----+-----+-----+
| BackupId | NodeId | Host | Timestamp | Comment |
+-----+-----+-----+-----+-----+
| 1 | 1 | tonfisk | 2016-09-21 21:13:09Z | |
| 1 | 2 | tonfisk | 2016-09-21 21:13:09Z | |
| 1 | 3 | tonfisk | 2016-09-21 21:13:09Z | |
| 1 | 4 | tonfisk | 2016-09-21 21:13:09Z | |
| 1 | 50 | tonfisk | 2016-09-21 21:13:12Z | Schema |
| 2 | 1 | tonfisk | 2016-09-21 21:17:50Z | |
| 2 | 2 | tonfisk | 2016-09-21 21:17:50Z | |
| 2 | 3 | tonfisk | 2016-09-21 21:17:50Z | |
```

```

| 2 | 4 | tonfisk | 2016-09-21 21:17:50Z | | |
| 2 | 50 | tonfisk | 2016-09-21 21:17:52Z | Schema | |
+-----+-----+-----+-----+-----+-----+
10 rows in set (0.01 sec)

```

Next, we have to find out the locations of the logical backup file and the backup files for each data node of the original cluster.

Locations of backup files. The backup files for each node are to be found under the folder specified by the cluster parameter `BackupDataDir` for data nodes and the parameter `backupdatadir` for `mysqld` nodes. Because the `get` command is not case sensitive, you can use this single command to check the values of both parameters:

```

mcm> get BackupDataDir mycluster;
+-----+-----+-----+-----+-----+-----+-----+
| Name          | Value          | Process1 | NodeId1 | Process2 | NodeId2 | Level  | Comment |
+-----+-----+-----+-----+-----+-----+-----+
| BackupDataDir | /opt/mcmbackup | ndbmttd  | 1        |           |          | Process |         |
| BackupDataDir | /opt/mcmbackup | ndbmttd  | 2        |           |          | Process |         |
| BackupDataDir | /opt/mcmbackup | ndbmttd  | 3        |           |          | Process |         |
| BackupDataDir | /opt/mcmbackup | ndbmttd  | 4        |           |          | Process |         |
| backupdatadir | /opt/mcmbackup | mysqld   | 50       |           |          | Process | MCM only |
+-----+-----+-----+-----+-----+-----+-----+
5 rows in set (0.18 sec)

```

The backup files for each backup of a specific BackupID are found under `BackupDataDir/BACKUP/BACKUP-ID/` for each data node, and under `backupdatadir/BACKUP/BACKUP-ID/` for each `mysqld` node. The comment "MCM only" in the row returned for the parameter `backupdatadir` indicates that `backupdatadir` is used by MySQL Cluster Manager only, and the directory it specifies contains only backups for the NDB tables' metadata. Notice that if `BackupDataDir` is not specified, the `get` command will return no value for it, and it takes up the value of `DataDir`, so that the image is stored in the directory `DataDir/BACKUP/BACKUP-backup_id`. If `backupdatadir` has not been specified, the `get` command will again return no value for it, and the logical backup files for the `mysqld` node are to be found at the default locations of `/path-to-mcm-data-repository/clusters/clustername/nodeid/BACKUP/BACKUP-Id`.

The process of restoring the backed-up data from the original cluster to the new one consists of the following steps:

1. Stop the original cluster:

```

mcm> stop cluster mycluster;
+-----+-----+
| Command result |
+-----+-----+
| Cluster stopped successfully |
+-----+-----+
1 row in set (19.54 sec)

mcm> show status mycluster;
+-----+-----+-----+
| Cluster | Status | Comment |
+-----+-----+-----+
| mycluster | stopped |         |
+-----+-----+-----+
1 row in set (0.05 sec)

```

2. Start your new cluster. Make sure the new cluster is operational and it has at least one free `ndbapi` slot for the `ndb_restore` utility to connect to the cluster:

```

mcm> start cluster newcluster2nodes;
+-----+-----+
| Command result |
+-----+-----+
| Cluster started successfully |
+-----+-----+
1 row in set (33.68 sec)

```

```
mcm> show status -r newcluster2nodes;
+-----+-----+-----+-----+-----+-----+
| NodeId | Process | Host   | Status | Nodegroup | Package |
+-----+-----+-----+-----+-----+-----+
| 49     | ndb_mgmd | tonfisk | running |           | mypackage |
| 1      | ndbmttd  | tonfisk | running | 0         | mypackage |
| 2      | ndbmttd  | tonfisk | running | 0         | mypackage |
| 50     | mysqld   | tonfisk | running |           | mypackage |
| 51     | ndbapi   | *      | added  |           |           |
+-----+-----+-----+-----+-----+-----+
5 rows in set (0.09 sec)
```

- Restore the logical backup of the metadata of the NDB tables onto the new cluster. See [Reloading SQL-Format Backups](#) for different ways to restore a logical backup. One way to do it is to open a `mysql` client, connect it to a `mysqld` node of the cluster, and then source the logical backup file with the `mysql` client:

```
mysql> source path-to-logical-backup-file/BACKUP-BackupID.mysqld_nodeid.schema.sql
```

See [Locations of backup files](#) above on how to find the path of the logical backup file. For our sample clusters, this is how the command looks like for restoring the NDB table metadata from the backup with the BackupID 2:

```
mysql> source /opt/mcmbackup/BACKUP/BACKUP-2/BACKUP-2.50.schema.sql
```

- Restore one by one the backup for each data node of the original cluster to the new cluster, using the `ndb_restore` program:

```
$> ndb_restore -b BackupID -n nodeID -r --backup_path=backup-folder-for-data_node
```

See [Locations of backup files](#) above on how to find the path of the data node backup files. For our sample clusters, to restore the data from the backup with the BackupID 2 for data node 1 to 4 of `mycluster`, execute the following commands:

```
$> ndb_restore --backupid=2 --nodeid=1 --restore_data --backup_path=/opt/mcmbackup/BACKUP/BACKUP-2/
```

```
$> ndb_restore --backupid=2 --nodeid=2 --restore_data --backup_path=/opt/mcmbackup/BACKUP/BACKUP-2/
```

```
$> ndb_restore --backupid=2 --nodeid=3 --restore_data --backup_path=/opt/mcmbackup/BACKUP/BACKUP-2/
```

```
$> ndb_restore --backupid=2 --nodeid=4 --restore_data --backup_path=/opt/mcmbackup/BACKUP/BACKUP-2/
```

The option `--disable-indexes` is used so indexes are ignored during the restores. This is because if we also try to restore the indexes node by node, they might not be restored in the right order for the foreign keys and unique key constraints to work properly. Therefore, the `--disable-indexes` option is used in the above commands, after the execution of which we rebuild the indexes with the following `ndb_restore` command and the `--rebuild-indexes` option (you only need to run this on one of the data nodes):

```
$> ndb_restore --backupid=2 --nodeid=1 --rebuild-indexes --backup_path=/opt/mcmbackup/BACKUP/BACKUP-2/
```

The data and indexes have now been fully restored to the new cluster.

3.7 Backing Up and Restoring MySQL Cluster Manager Agents

This section explains how to back up configuration data for `mcmd` agents and how to restore the backed-up agent data. Used together with the `backup cluster` command, the `backup agents` command allows you to backup and restore a complete cluster-plus-manager setup.

If no host names are given with the `backup agents` command, backups are created for all agents of the site:

```
mcm> backup agents mysite;
```

```
+-----+
| Command result |
```

```

+-----+
| Agent backup created successfully |
+-----+
1 row in set (0.07 sec)
    
```

To backup one or more specific agents, specify them with the `--hosts` option:

```

mcm> backup agents --hosts=tonfisk mysite;
+-----+
| Command result          |
+-----+
| Agent backup created successfully |
+-----+
1 row in set (0.07 sec)
    
```

If no site name is given, only the agent that the `mcm` client is connected to is backed up.

The backup for each agent includes the following contents from the agent repository (`mcm_data` folder):

- The `rep` subfolder
- The metadata files `high_water_mark` and `repchksum`

The repository is locked while the backup are in progress, to avoid creating an inconsistent backup. The backup for each agent is created in a subfolder named `rep_backup/timestamp` under the agent's `mcm_data` folder, with `timestamp` reflecting the time the backup began. If you want the backup to be at another place, create a soft link from `mcm_data/rep_backup` to your desired storage location.

For MySQL Cluster Manager 1.4.6 and later: You can list agent backups using the `list backups` command with the `--agent` option and the site name:

```

mcm> list backups --agent mysite;
+-----+-----+-----+-----+-----+
| BackupId | Agent | Host   | Timestamp                | Comment |
+-----+-----+-----+-----+-----+
| 1522914101 | 0     | tonfisk | 2018-04-05 07:41:41Z |         |
| 1522914105 | 0     | tonfisk | 2018-04-05 07:41:45Z |         |
| 1522914121 | 0     | tonfisk | 2018-04-05 07:42:01Z |         |
+-----+-----+-----+-----+-----+
3 rows in set (0.00 sec)
    
```

To restore the backup for an agent:

- Wipe the contents of the agent's `mcm_data/rep` folder
- Delete the metadata files `high_water_mark` and `repchksum` from the `mcm_data` folder
- Copy the contents in the `mcm_data/rep_backup/timestamp/rep` folder back into the `mcm_data/rep` folder
- Copy the metadata files `high_water_mark` and `repchksum` from the `mcm_data/rep_backup/timestamp` folder back into the `mcm_data` folder
- Restart the agent

The steps are illustrated below:

```

mysql@tonfisk$ cd mcm_data
mysql@tonfisk$ cp mcm_data/rep_backup/timestamp/rep/* ./rep/
mysql@tonfisk$ cp mcm_data/rep_backup/timestamp/high_water_mark ./
mysql@tonfisk$ cp mcm_data/rep_backup/timestamp/repchksum ./
mysql@tonfisk$ mcm1.4.8/bin/mcmd
    
```

The backup may be manually restored on just one, or more than one agents. If backup is restored for only one agent on, say, host A, host A will contact the other agents of the site to make them recover their repositories from host A using the usual mechanism for agent recovery. If all agents on all hosts are restored and restarted manually, the situation will be similar to the normal restarting all agents after stopping them at slightly different points in time.

If configuration changes has been made to the cluster since the restored backup was created, the same changes must be made again after the agent restores have been completed, to ensure that the agents' configurations match those of the actual running cluster. For example: sometime after a backup was done, a `set MaxNoOfTables:ndbmt=500 mycluster` command was issued and soon afterward, something happened and corrupted the agent repository; after the agent backup was restored, the same `set` command has to be rerun in order to update the `mcmd` agents' configurations. While the command does not effectively change anything on the cluster itself, after it has been run, a rolling restart of the cluster processes using the `restart cluster` command is still required.

3.8 Restoring a MySQL Cluster Manager Agent with Data from Other Agents

Sometimes, an `mcmd` agent can fail to restart after a failure because its configuration store has been corrupted (for example, by an improper shutdown of the host). If there is at least one other `mcmd` agent that is still functioning properly on another host of the cluster, you can restore the failed agent by the following steps:

For MySQL Cluster Manager 1.4.6 and earlier:

- Make sure the `mcmd` agent has really been stopped.
- Go to the agent repository (the agent's `mcm_data` folder).
- Wipe the contents of the `rep` folder.
- Delete the metadata files `high_water_mark` and `repchksum`.
- Delete the `manager.lock` file.
- Restart the agent.

For MySQL Cluster Manager 1.4.7 and later:

- Make sure the `mcmd` agent has really been stopped.
- Restart `mcmd` with the `--initial` option, which wipes the contents of the `rep` folder after backing them up and then starts the agent.

The agent then recovers the configuration store from other agents on the other hosts.

However, if all the `mcmd` agents for the cluster are malfunctioning, you will have to do one of the following:

- Restore one of the agents first using an agent backup (see [Section 3.7, “Backing Up and Restoring MySQL Cluster Manager Agents”](#) for details), and then restore the other agents with it .
- Recreate the whole cluster and restore it using a cluster backup (see [Section 3.6, “MySQL NDB Cluster Backup and Restore Using MySQL Cluster Manager”](#) for details).

3.9 Setting Up MySQL NDB Cluster Replication with MySQL Cluster Manager

This section provides sample steps for setting up a MySQL NDB Cluster replication with a single replication channel using the MySQL Cluster Manager.

Before trying the following steps, it is recommended that you first read [NDB Cluster Replication](#) to familiarize yourself with the concepts, requirements, operations, and limitations of MySQL NDB Cluster replication.

1. Create and start a source cluster:

```
mcm> create site --hosts=tonfisk msite;

mcm> add package --basedir=/usr/local/cluster-mgt/cluster-7.3.2 7.3.2;

mcm> create cluster -P 7.3.2 -R \
    ndb_mgmd@tonfisk,ndbmt@tonfisk,ndbmt@tonfisk,mysql@tonfisk,mysql@tonfisk,ndbapi@*,ndbapi@* \
    source;

mcm> set portnumber:ndb_mgmd=4000 source;

mcm> set port:mysql:51=3307 source;

mcm> set port:mysql:50=3306 source;

mcm> set server_id:mysql:50=100 source;

mcm> set log_bin:mysql:50=binlog source;

mcm> set binlog_format:mysql:50=ROW source;

mcm> set ndb_connectstring:mysql:50=tonfisk:4000 source;

mcm> start cluster source;
```

2. Create and start a replica cluster (we begin with creating a new site called “ssite” just for the replica cluster; you can also skip that and put the source and replica cluster hosts under the same site instead):

```
mcm> create site --hosts=flundra ssite;

mcm> add package --basedir=/usr/local/cluster-mgt/cluster-7.3.2 7.3.2;

mcm> create cluster -P 7.3.2 -R \
    ndb_mgmd@flundra,ndbmt@flundra,ndbmt@flundra,mysql@flundra,mysql@flundra,ndbapi@*,ndbapi@* \
    replica;

mcm> set portnumber:ndb_mgmd=4000 replica;

mcm> set port:mysql:50=3306 replica;

mcm> set port:mysql:51=3307 replica;

mcm> set server_id:mysql:50=101 replica;

mcm> set ndb_connectstring:mysql:50=flundra:4000 replica;

mcm> set slave_skip_errors:mysql=all replica;

mcm> start cluster replica;
```

3. Create a replica account (with the user name “myreplica” and password “mypw”) on the source cluster with the appropriate privilege by logging into the source replication client (`mysqlM`) and issuing the following statements:

```
mysqlM> GRANT REPLICATION SLAVE ON *.* TO 'myreplica'@'flundra'
-> IDENTIFIED BY 'mypw';
```

4. Log in to the replica cluster client (`mysqlS`) and issue the following statements:

```
mysqlS> CHANGE MASTER TO
-> MASTER_HOST='tonfisk',
-> MASTER_PORT=3306,
```

```
-> MASTER_USER='myreplica',
-> MASTER_PASSWORD='mypw';
```

5. Start replication by issuing the following statement with the replica cluster client:

```
mysql> START SLAVE;
```

The above example assumes that the source and replica clusters are created at about the same time, with no data on both before replication starts. If the source cluster has already been operating and has data on it when the replica cluster is created, after step 3 above, follow these steps to transfer the data from the source cluster to the replica cluster and prepare the replica cluster for replication:

1. Back up your source cluster using the `backup cluster` command of MySQL Cluster Manager:

```
mcm> backup cluster source;
```



Note

Only NDB tables are backed up by the command; tables using other MySQL storage engines are ignored.

2. Look up the backup ID of the backup you just made by listing all backups for the source cluster:

```
mcm> list backups source;
```

BackupId	NodeId	Host	Timestamp	Comment
1	1	tonfisk	2014-10-17 20:03:23	
1	2	tonfisk	2014-10-17 20:03:23	
2	1	tonfisk	2014-10-17 20:09:00	
2	2	tonfisk	2014-10-17 20:09:00	

From the output, you can see that the latest backup you created has the backup ID “2”, and backup data exists for node “1” and “2”.

3. Using the backup ID and the related node IDs, identify the backup files just created under `/mcm_data/clusters/cluster_name/node_id/data/BACKUP/BACKUP-backup_id/` in the source cluster's installation directory (in this case, the files under the `/mcm_data/clusters/source/1/data/BACKUP/BACKUP-2` and `/mcm_data/clusters/source/2/data/BACKUP/BACKUP-2`), and copy them over to the equivalent places for the replica cluster (in this case, `/mcm_data/clusters/replica/1/data/BACKUP/BACKUP-2` and `/mcm_data/clusters/replica/2/data/BACKUP/BACKUP-2` under the replica cluster's installation directory). After the copying is finished, use the following command to check that the backup is now available for the replica cluster:

```
mcm> list backups replica;
```

BackupId	NodeId	Host	Timestamp	Comment
2	1	flundra	2014-10-17 21:19:00	
2	2	flundra	2014-10-17 21:19:00	

4. Restore the backed up data to the replica cluster (note that you need an unused `ndbapi` slot for the `restore cluster` command to work):

```
mcm> restore cluster --backupid=2 replica;
```

5. On the source cluster client, use the following command to identify the correct binary log file and position for replication to start:

```
mysqlM> SHOW MASTER STATUS\G;
***** 1. row *****
      File: binlog.000017
      Position: 2857
```

```
Binlog_Do_DB:
Binlog_Ignore_DB:
Executed_Gtid_Set:
```

6. On the replica cluster client, provide to the replica cluster the information of the source cluster, including the binary log file name (with the `MASTER_LOG_FILE` option) and position (with the `MASTER_LOG_POS` option) you just discovered in step 5 above:

```
mysqlS> CHANGE MASTER TO
-> MASTER_HOST='tonfisk',
-> MASTER_PORT=3306,
-> MASTER_USER='myreplica',
-> MASTER_PASSWORD='mypw',
-> MASTER_LOG_FILE='binlog.000017',
-> MASTER_LOG_POS=2857;
```

7. Start replication by issuing the following statement with the replica cluster client:

```
mysqlS> START SLAVE;
```

As an alternative to these steps, you can also follow the steps described in [NDB Cluster Backups With NDB Cluster Replication](#) to copy the data from the source to the replica and to specify the binary log file and position for replication to start.

3.10 NDB Cluster Downgrades Involving Logical Downgrades of MySQL Nodes

While the `upgrade cluster` command can be used to perform a downgrade for an NDB Cluster (see the description for the `upgrade cluster` command for the restrictions that apply), special steps are required when that involves a downgrade for `mysqld` for which in-place downgrade is not supported. For example, downgrading from NDB Cluster 7.5 to 7.4 involves downgrading the MySQL server from 5.7 and 5.6, for which the logical downgrade method must be employed (see [Downgrade Paths](#), for information on downgrading `mysqld`). This section explains the steps to take in the situation. Let us use this sample cluster as an example:

```
mcm> show status -r mycluster;
```

NodeId	Process	Host	Status	Nodegroup	Package
49	ndb_mgmd	tonfisk	running		7.5-DMR1
1	ndbmtdd	flundra	running	0	7.5-DMR1
2	ndbmtdd	grindval	running	0	7.5-DMR1
50	mysqld	flundra	running		7.5-DMR1
51	mysqld	grindval	running		7.5-DMR1
52	ndbapi	*	added		

Assuming the NDB Cluster has been upgraded from version 7.4 to 7.5 and the user then encountered a major problem and now needs to revert the cluster to version 7.4, these are the steps to follow:

1. Backup all the cluster data:

```
mcm> backup cluster mycluster;
```

```
+-----+
| Backup completed successfully |
+-----+
```

2. Use `mysqldump` to backup non-NDB data from each of the `mysqld` nodes.
3. Assuming the package for the lower version of NDB Cluster you want to downgrade to is still available to MySQL Cluster Manager (if it is not the case, use the `add package` command to supply the package), use the `upgrade cluster` command and the `--nodeid` option to downgrade the management and data nodes:

```
mcm> upgrade cluster -P 7.4.10 --nodeid=49,1,2 --retry mycluster;
```

```
+-----+
| Cluster upgraded successfully |
+-----+
```

Remember to include ALL management nodes and data nodes with the `--nodeid` option, but leave out any `mysqld` node. The `show status` command shows which nodes have now been downgraded:

```
mcm> show status -r mycluster;
+-----+-----+-----+-----+-----+-----+
| NodeId | Process | Host      | Status | Nodegroup | Package |
+-----+-----+-----+-----+-----+-----+
| 49     | ndb_mgmd | tonfisk   | running |           | 7.4.10  |
| 1      | ndbmttd  | flundra   | running | 0         | 7.4.10  |
| 2      | ndbmttd  | grindval  | running | 0         | 7.4.10  |
| 50     | mysqld   | flundra   | running |           | 7.5-DMR1|
| 51     | mysqld   | grindval  | running |           | 7.5-DMR1|
| 52     | ndbapi   | *         | added  |           |         |
+-----+-----+-----+-----+-----+-----+
```

4. Recreate the data directory for each `mysqld` node:

a. Stop the `mysqld` node:

```
mcm> stop process 50 mycluster;
+-----+
| Process stopped successfully |
+-----+
```

b. Delete the data directory of the `mysqld` node and recreate it, as an empty folder:

```
user@host$ cd mcm_data/clusters/mycluster/50/
user@host$ rm -rf data
user@host$ mkdir data
```

c. Initialize the recreated data directory using the `mysql_install_db` script in the package of the NDB Cluster version you are downgrading to:

```
user@host$ /clusters/7.4/scripts/mysql_install_db \
--defaults-file=/mcm_data/clusters/mycluster/50/cfg/my.cnf \
--datadir=/mcm_data/clusters/mycluster/50/data --basedir=/clusters/7.4.10 \
> mysql_install_db.log 2>&1
```

(When you are downgrading to NDB Cluster 7.5 or later, the `mysqld` data directory should be initialized with the `mysqld --initialize` command; see [Initializing the Data Directory](#) for details.)

d. Restart the `mysqld` node using MySQL Cluster Manager (due to the downgrade already performed earlier on the management and data nodes, `mcmd` is going to start the `mysqld` node with the binary of the cluster version you are downgrading to):

```
mcm> start process 50 mycluster;
ERROR 9003 (00MGR): Tx {690be0d6 270 0 15} timed out on agent 0 @{19 0 50
Wait for mysqld to start} participants: 0@tonfisk:18620[X] 1@flundra:18620[ ] 2@grindval:18620[ ]
```

There might be a timeout error as shown above, which is due to the fact that the `mcmd` user is still to be created on the `mysqld` node (this happens in the next step). The error can be avoided if you perform the next step soon enough. But even if the error occurs, there is no issue here, as `mcmd` will keep trying to connect to the `mysqld` node.

e. Recreate the `mcmd` user on the `mysqld` node:

```
mysql> CREATE USER 'mcmd'@'127.0.0.1' IDENTIFIED BY '<manager-password>';
mysql> GRANT ALL PRIVILEGES ON *.* TO 'mcmd'@'127.0.0.1' IDENTIFIED BY \
'<manager-password>' WITH GRANT OPTION;
mysql> FLUSH PRIVILEGES;
```

- f. Reload onto the `mysqld` node the non-NDB data you backed up earlier using `mysqldump`:

```
mysql -h 127.0.0.1 -P 3306 -u root -p < dumpfile.sql
```

5. Repeat step 4 above and each of its substeps for every `mysqld` node of your cluster.
6. After all the `mysqld` nodes have been downgraded and their data directories prepared, `mcmd` eventually reconnects with them all:

```
mcm> show status -r mycluster;
```

NodeId	Process	Host	Status	Nodegroup	Package
49	ndb_mgmd	tonfisk	running		7.4.10
1	ndbmt	flundra	running	0	7.4.10
2	ndbmt	grindval	running	0	7.4.10
50	mysqld	flundra	running		7.4.10
51	mysqld	grindval	running		7.4.10
52	ndbapi	*	added		

Chapter 4 MySQL Cluster Manager Client Commands

Table of Contents

4.1 Online Help and Information Commands	71
4.2 MySQL Cluster Manager Site and Agent Commands	77
4.2.1 The <code>add hosts</code> Command	77
4.2.2 The <code>remove hosts</code> Command	78
4.2.3 The <code>change log-level</code> Command	79
4.2.4 The <code>rotate log</code> Command	79
4.2.5 The <code>collect logs</code> Command	80
4.2.6 The <code>create site</code> Command	81
4.2.7 The <code>delete site</code> Command	82
4.2.8 The <code>list sites</code> Command	82
4.2.9 The <code>list hosts</code> Command	83
4.2.10 The <code>show settings</code> Command	83
4.2.11 The <code>stop agents</code> Command	84
4.2.12 The <code>version</code> Command	84
4.2.13 The <code>show warnings</code> Command	85
4.3 MySQL Cluster Manager Package Commands	85
4.3.1 The <code>add package</code> Command	85
4.3.2 The <code>delete package</code> Command	87
4.3.3 The <code>list packages</code> Command	88
4.4 MySQL Cluster Manager Cluster Commands	89
4.4.1 The <code>create cluster</code> Command	89
4.4.2 The <code>delete cluster</code> Command	93
4.4.3 The <code>list clusters</code> Command	94
4.4.4 The <code>list nextnodeids</code> Command	94
4.4.5 The <code>restart cluster</code> Command	94
4.4.6 The <code>show status</code> Command	95
4.4.7 The <code>start cluster</code> Command	99
4.4.8 The <code>stop cluster</code> Command	101
4.4.9 The <code>autotune</code> Command	101
4.4.10 The <code>upgrade cluster</code> Command	102
4.5 MySQL Cluster Manager Configuration Commands	105
4.5.1 The <code>get</code> Command	108
4.5.2 The <code>reset</code> Command	119
4.5.3 The <code>set</code> Command	125
4.6 MySQL Cluster Manager Process Commands	131
4.6.1 The <code>add process</code> Command	132
4.6.2 The <code>change process</code> Command	135
4.6.3 The <code>list processes</code> Command	137
4.6.4 The <code>start process</code> Command	138
4.6.5 The <code>stop process</code> Command	138
4.6.6 The <code>update process</code> Command	139
4.6.7 The <code>remove process</code> Command	140
4.7 MySQL Cluster Manager Backup and Restore Commands	141
4.7.1 The <code>abort backup</code> Command	141
4.7.2 The <code>backup cluster</code> Command	141
4.7.3 The <code>list backups</code> Command	143
4.7.4 The <code>restore cluster</code> Command	144
4.7.5 The <code>backup agents</code> Command	146
4.8 MySQL Cluster Manager Cluster Importation Commands	147
4.8.1 The <code>import cluster</code> Command	147
4.8.2 The <code>import config</code> Command	147

Identifiers in client commands

Case-sensitivity rules for client commands

Options common to client commands

The sections in this chapter describe commands used in the MySQL Cluster Manager 1.4.8 client for tasks such as defining sites, packages, and MySQL NDB Cluster instances (“clusters”); configuring a MySQL NDB Cluster; and getting the status of a running MySQL NDB Cluster. These commands are issued to the management agent using the `mysql` client program included with the MySQL NDB Cluster distribution (for information about the `mysql` client not specific to using MySQL Cluster Manager, see [mysql — The MySQL Command-Line Client](#)). Each MySQL Cluster Manager client command takes the form shown here:

```
instruction [options] [arguments]

options:
  option [option] [...]

option:
  --option-long-name[=value-list]
  | -option-short-name [value-list]

value-list:
  value[,value[,...]]

arguments:
  argument [argument] [...]
```

Consider the following MySQL Cluster Manager command, which starts a MySQL NDB Cluster named `mycluster` and backgrounds the deletion process so that the client can be used to execute other commands in the meantime, without having to wait on the `start cluster` command to complete:

```
start cluster --background mycluster;
```

In this example, the command contains a `start cluster` instruction. An instruction consists of one or two keywords, such as `set`, or `show status`. This instruction is modified by the `--background` option, which follows it; however, this option assigns no values.

Most command options have short forms, consisting of single letters, in addition to their long forms. Using the short form of the `--background` option, the previous example could also be written like this:

```
start cluster -B mycluster;
```

The long form of an option must be preceded by a double dash (`--`), and is not case-sensitive (lower case being the canonical form). The short form of an option must be preceded by a single dash (`-`), and is case-sensitive. In either case, the dash character or characters must come immediately before the option name, and there must be no space characters between them. Otherwise, the MySQL Cluster Manager client cannot parse the command correctly. More information about long and short forms of options is given later in this section.



Important

Do not confuse options given to MySQL Cluster Manager client commands with `mysql` client options. A MySQL Cluster Manager client command option is always employed as part of a MySQL Cluster Manager client command; it is *not* passed to the `mysql` client when invoking it.

In addition, you cannot issue queries or other SQL statements in the MySQL Cluster Manager client. These are not recognized by the client, and are rejected with an error. The converse of this is also true: MySQL Cluster Manager client commands are not recognized by the standard `mysql` client.

The instruction just shown takes the argument `mysite`. The argument is usually an identifier that names the object to be effected; in this case, the command deletes the site whose name matches the argument. (For more information, see [Section 4.2.6, “The `create site` Command”](#).)

An additional `--verbose` option can be used for the `create cluster`, `add process`, and `list hosts` commands. In both cases, using the option causes the command to return a list of the MySQL NDB Cluster processes affected by the command; this includes their node IDs, process types, and the hosts where they are located.

Identifiers in client commands.

A legal MySQL Cluster Manager identifier consists of any sequence of characters from among the following:

- The letters `a` through `z` and `A` through `Z`
- The digits `0` through `9`
- The dash (`-`), period (`.`), and underscore (`_`) characters

A MySQL Cluster Manager identifier must begin with a letter or digit.

Case-sensitivity rules for client commands.

The rules for case-sensitivity of MySQL Cluster Manager identifiers, commands, command options, process names, and configuration attributes are as follows:

- *Identifiers are case-sensitive.* For example, `delete site mycluster` cannot be used to delete a site named `myCluster`.
- *Command keywords and the long forms of command options are case-insensitive.* For example, any of the three commands `delete cluster mycluster`, `DELETE CLUSTER mycluster`, and `DeLeTe cLuStEr mycluster` works to delete the MySQL NDB Cluster instance named `mycluster`.

In this manual, we show command keywords and the long forms of command options in lowercase, but you are not required to follow this convention if you do not wish to do so.

- *The short forms of command options are case-sensitive.* For example, `-b` (lowercase) is the short form of the `--basedir` option, but `-B` (uppercase) is the short form of the `--background` option.
- *Names of MySQL NDB Cluster processes are case-insensitive.* For example, either of the commands `get --include-defaults DataMemory:ndbd mycluster` or `get --include-defaults datamemory:NDBD mycluster` reports the data memory allocated for each `ndbd` process in the cluster named `mycluster`.

In this manual, we show names of MySQL NDB Cluster processes in lowercase. You are not required to follow this convention if you do not wish to do so; however, since the corresponding executables are named and must be invoked in lowercase, we suggest that you use lowercase.

- *Configuration attribute names are case-insensitive.* For example, either of the commands `get --include-defaults DataMemory:ndbd mycluster` or `get --include-defaults datamemory:ndbd mycluster` returns the data memory allocated for each `ndbd` process in the cluster named `mycluster`; either of the commands `set engine-condition-pushdown:mysqlid:4=0 mycluster` or `set Engine-Condition-Pushdown:mysqlid:4=0 mycluster` disables the condition pushdown optimization in the `mysqld` process having the node ID `4` in the MySQL NDB Cluster named `mycluster`.



Note

Configuration attributes in the MySQL Cluster Manager derive from two different sources: MySQL NDB Cluster configuration parameters, and

MySQL Server options. MySQL NDB Cluster configuration parameters are case-insensitive, but their canonical forms use upper camelcase (that is, medial capitalization including the first letter). This means that whether you set a value for data memory using the MySQL Cluster Manager client or in the `config.ini` file, you can refer to it as `DataMemory`, `datamemory`, or `DATAMEMORY` without any negative impact. However, MySQL Server command-line options are case-sensitive and use only lowercase. This means that, for example, `set Engine-Condition-Pushdown:mysqld:4=0 mycluster` in the MySQL Cluster Manager client works to disable condition pushdown in the indicated `mysqld` process, but if you invoke the `mysqld` executable from a system prompt using `--Engine-Condition-Pushdown=0, mysqld` fails to start.

In this manual, for easy recognition, we show configuration attribute names as having the same lettercase used in other MySQL documentation; thus, we always refer to `DataMemory`, rather than `datamemory` or `DATAMEMORY`, and `engine-condition-pushdown`, rather than `Engine-Condition-Pushdown` or `ENGINE-CONDITION-PUSHDOWN`. While you are not required to do this when using MySQL Cluster Manager, we suggest that you also follow this convention.



Note

Values that contain space characters must be quoted using single quote (') characters. For example, if you wish to define a package named `mypackage` for a site named `mysite` using `/usr/local/mysql cluster/7.3` (where a space occurs between `mysql` and `cluster`) as the path to the base directory on all hosts, the correct command would be `add package --basedir='/usr/local/mysql cluster/7.3' mypackage`.

To decrease the possibility of errors in reading and entering MySQL Cluster Manager commands, we recommend avoiding the use of space characters whenever possible.

Each command must end with a terminator character. By default, this is the semicolon (;) character. However, the sequences `\g` and `\G` are also supported as command terminators. The `\G` terminator causes the output to be vertically formatted (the same as in the standard `mysql` client), as shown in this example:

```
mcm> get DataMemory mycluster\G
***** 1. row *****
  Name: DataMemory
  Value: 500M
Process1: ndbd
  Id1: 2
Process2:
  Id2:
  Level: Process
  Comment:
***** 2. row *****
  Name: DataMemory
  Value: 500M
Process1: ndbd
  Id1: 3
Process2:
  Id2:
  Level: Process
  Comment:
2 rows in set (0.22 sec)
```

By convention (for reasons of readability), we do not normally include the command terminator when showing the syntax for a command in Backus-Naur format or when including a MySQL Cluster Manager command inline in this text. However, if you do not use a statement terminator when you enter the command in the MySQL Cluster Manager client, the client displays a special “waiting...” prompt `->` until you supply a terminator, as shown here:

```
mcm> list sites
->
->
->
-> ;
Empty set (1.50 sec)
```

(This is the same as the behavior of the `mysql` client when you fail to end a statement with a terminator.)

A command option can also in many cases accept (or even require) a set of one or more *values*. The next example includes such an option, and also demonstrates setting of multiple values in a single option by passing them to the option as a comma-separated list:

```
mcm> create site --hosts=tonfisk,flundra mysite;
+-----+
| Command result          |
+-----+
| Site created successfully |
+-----+
1 row in set (7.41 sec)
```

The command just shown creates a site named `mysite`, consisting of two hosts named `tonfisk` and `flundra`. (See [Section 4.2.6, “The create site Command”](#), for more information about this command.) Since we used the long form of the `--hosts` option, we were required to use an equals sign (=) to mark the end of the option name and the beginning of the values list. You must not insert any space characters before or after the equal sign; doing so causes an error, as shown here:

```
mcm> create site --hosts =grindval,haj yoursite;
ERROR 7 (00MGR): Option --hosts requires a value
mcm> create site --hosts= grindval,haj yoursite;
ERROR 7 (00MGR): Option --hosts requires a value
```

The short form of an option does not use an equal sign. Instead, the value-list is separated from the option by a space. Using the `-h` option, which is the short form of the `--hosts` option, the previous `create site` command can be entered and executed like this:

```
mcm> create site -h tonfisk,flundra mysite;
+-----+
| Command result          |
+-----+
| Site created successfully |
+-----+
1 row in set (7.41 sec)
```

The short forms of options actually accept multiple spaces between the option name and the values list; however, a single space is sufficient. If you omit the space, or try to use an equal sign, the command fails with an error, as shown here:

```
mcm> create site -htonfisk,flundra mysite;
ERROR 6 (00MGR): Illegal number of operands
mcm> create site -h=tonfisk,flundra mysite;
ERROR 3 (00MGR): Illegal syntax
```

Any option value containing one or more whitespace characters, one or more dash characters (-), or both, must be quoted using single quotation marks. Multiple values should be separated by commas only; do not insert spaces before or after any of the commas. Using spaces before or after the commas in a list of values causes the command to fail with an error, as shown here:

```
mcm> create site --hosts=tonfisk, flundra mysite;
ERROR 6 (00MGR): Illegal number of operands
```

As you can see from the examples just shown, a MySQL Cluster Manager client command returns a result set, just as an SQL statement does in the standard `mysql` client. The result set returned by a MySQL Cluster Manager client command consists of one of the following:

- **A single row that contains a message indicating the outcome of the command.** The `create site` command in the last example returned the result `Site created successfully`, to inform the user that the command succeeded.
- **One or more rows listing requested objects or properties.** An example of such a command is `list processes`, as shown here:

```
mcm> list processes mycluster;
+-----+-----+-----+
| NodeId | Name      | Host      |
+-----+-----+-----+
| 49     | ndb_mgmd  | flundra   |
| 1      | ndbd      | tonfisk   |
| 2      | ndbd      | grindval  |
| 50     | mysqld    | haj       |
| 51     | mysqld    | torsk     |
| 52     | ndbapi    | *         |
+-----+-----+-----+
6 rows in set (0.03 sec)
```

In the case of `list processes`, each row in the result contains the ID and type of a node in the MySQL NDB Cluster named `mycluster`, together with the name of the host on which the process is running.

- **An empty result set.** This can occur with one of the `list` commands when there is nothing to report, such as when `list sites` is used before any sites have been created:

```
mcm> list sites;
Empty set (0.72 sec)
```

Each command must be entered separately; it is not possible to combine multiple commands on a single line.

Options common to client commands.

The following three options are common to most MySQL Cluster Manager client commands:

1. `--help` (short form: `-?`): Common to all client commands. Provides help output specific to the given command. See [Section 4.1, “Online Help and Information Commands”](#), for more information about this option.
2. `--force` (short form `-f`): Causes any safety checks to be bypassed when executing the command. For example, `delete cluster mycluster` normally fails if any of the MySQL NDB Cluster processes in the MySQL NDB Cluster named `mycluster` are running; however, `delete cluster --force mycluster` forces the shutdown of `mycluster`, followed by the deletion of `mycluster` from MySQL Cluster Manager's inventory.

The `--force` option is supported for the following MySQL Cluster Manager client commands:

- `delete site`
- `start cluster`
- `restart cluster`
- `stop cluster`
- `delete cluster`
- `upgrade cluster`
- `add process`
- `start process`
- `stop process`

- `remove process`
 - `set`
 - `reset`
3. `--background` (short form `-B`): Rather than waiting for the command to complete, the MySQL Cluster Manager client immediately returns the command prompt, allowing you to perform additional tasks in the client while that command continues to execute in the background. This can be useful when executing commands that might require some time to complete (such as starting a cluster with a great many nodes).

This option is supported by all client commands except for `create site`, `delete site`, `add hosts`, `add package`, and `delete package`.

4.1 Online Help and Information Commands

Online help is available in the MySQL Cluster Manager client for MySQL Cluster Manager client commands. The client can provide both general and command-specific information. In addition, you can obtain information about `mysql` client commands that are independent of the MySQL server and thus are also available for use when connected to the MySQL Cluster Manager agent.

Listing MySQL Cluster Manager client commands.

For a list of all commands with brief descriptions, use the `list commands` command, as shown here:

```
mcm> list commands;
+-----+
| Help
+-----+
| COMMANDS
|
| abort backup          Abort an ongoing cluster backup.
| add hosts             Add hosts to site.
| add package           Add a package alias.
| add process           Add cluster process.
| autotune              Autotune a cluster to given use-case template.
| backup agents         Backup the agents repository and metadata.
| backup cluster        Backup a cluster.
| change log-level      Change the log-level
| change process        Change process type.
| collect logs          Collect log files.
| create cluster         Create a cluster.
| create site           Create a site.
| delete cluster        Delete a cluster.
| delete package        Delete a package.
| delete site           Delete a site.
| get                   Get configuration variables.
| import cluster         Import a running cluster.
| import config         Import the configuration of a running cluster.
| list backups          List backup images.
| list clusters         List all clusters.
| list commands         List the help text.
| list hosts            List hosts in site.
| list nextnodeids      List next nodeids to be allocated.
| list packages         List all packages.
| list processes        List processes.
| list sites            List all sites.
| remove hosts          Remove hosts from site.
| remove process        Remove a cluster process.
| reset                 Reset configuration variables.
| restart cluster       Restart a cluster.
| restore cluster       Restore a cluster.
| rotate log            Rotate the mcmd log.
| set                   Set configuration variables.
| show settings         Show agent settings.
| show status           Show cluster, process, operation or backup status.
| start cluster         Start a cluster.
```

```

| start process      Start a cluster process.
| stop agents       Stop agents in site.
| stop cluster      Stop a cluster.
| stop process      Stop a cluster process.
| upgrade cluster   Upgrade a cluster.
| version           Print version information.
|
| GLOBAL OPTIONS
| Options that can be used with all commands
|
|     --help|-?          Print detailed help.
|
| Use '<COMMAND> --help' to see verbose help for individual commands.
+-----+
51 rows in set (0.03 sec)

```

Obtaining information about specific MySQL Cluster Manager client commands.

To obtain more detailed help specific to a given command, invoke the command using the `--help` option, as shown in this example:

```

mcm> create site --help;
+-----+
| Help
+-----+
|
| create site [options] <sitename>
|
|     Creates a site from the hosts listed in --hosts.
|
| Required options:
| --hosts|-h          Comma separated list of hostnames.
|                    Format: --hosts = <host>[,<host>]*.
|
+-----+
9 rows in set (0.00 sec)

```

For any MySQL Cluster Manager client command, the `--help` option may be abbreviated to `-?`:

```

mcm> list processes -?;
+-----+
| Help
+-----+
|
| list processes <sitename>
|
|     Lists all processes defined in the specified cluster.
|
+-----+
4 rows in set (0.00 sec)

```

As mentioned elsewhere in this manual (see [Chapter 4, MySQL Cluster Manager Client Commands](#)), many other MySQL Cluster Manager command options have short forms as well. These are included in the documentation for each command. You can also find out what these are for a given command by invoking it with the `--help` or `-?` option.

You can obtain the release version of the MySQL Cluster Manager software in use from the output of the `version` command.

mysql client commands in the MySQL Cluster Manager client.

You can also use most standard `mysql` client commands in the MySQL Cluster Manager client (but *not* SQL statements, which depend on being connected to a MySQL server), such as `prompt`, `quit`, and `status`. For example, the output of the `status` command when connected to the MySQL Cluster Manager agent looks something like this (depending on the exact version of the client and agent that you are using and possibly other factors):

```

mcm> status
-----
/home/jon/bin/mcm/libexec/./cluster/bin/mysql
Ver 14.14 Distrib 5.7.29-ndb-7.6.13, for linux2.6 (x86_64) using EditLine wrapper

```

```

Connection id:          1
Current database:      <n/a>
Current user:          admin
SSL:                   Not in use
Current pager:         less
Using outfile:         ''
Using delimiter:       ;
Server version:        1.4.8 MySQL Cluster Manager
Protocol version:      10
Connection:            127.0.0.1 via TCP/IP
Server characterset:   <n/a>
Db characterset:       <n/a>
Client characterset:   <n/a>
Conn. characterset:    <n/a>
TCP port:              1862
-----
    
```



Note

You may use the command delimiter with `mysql` client commands, but you are not required to do so. For instance, assuming that the delimiter in use was the default semicolon (`;`) character, we could have executed the `status` command like this:

```

mcm> status;
-----
/home/jon/bin/mcm/cluster/bin/mysql Ver 14.14 Distrib 5.7.29-ndb-7.6.13,...
    
```

A particularly useful `mysql` client command that you can also employ with `mcm` is the `source` command (short form: `\.`), which you can use for executing scripts containing MySQL Cluster Manager client commands. On a Linux system, you might have a text file in your home directory named `get-attributes.mcm`, whose contents are shown here:

```

get :ndb_mgmd mycluster\G
get :ndbd mycluster\G
get :mysqld mycluster\G
    
```

Assuming that you have created a cluster named `mycluster`, you can run this script in the client; the results vary according to how this cluster is actually configured, but should be similar to this:

```

mcm> \. ~/get-attributes.mcm
mcm> get :ndb_mgmd mycluster\G
***** 1. row *****
  Name: DataDir
  Value: /home/jon/bin/mcm/mcm_data/clusters/mycluster/49/data
Process1: ndb_mgmd
  NodeId1: 49
Process2:
  NodeId2:
  Level:
  Comment:
***** 2. row *****
  Name: HostName
  Value: flundra
Process1: ndb_mgmd
  NodeId1: 49
Process2:
  NodeId2:
  Level:
  Comment: Read only
***** 3. row *****
  Name: NodeId
  Value: 49
Process1: ndb_mgmd
  NodeId1: 49
Process2:
  NodeId2:
  Level:
  Comment: Read only
    
```

```

***** 4. row *****
  Name: PortNumber
  Value: 1186
Process1: ndb_mgmd
  NodeId1: 49
Process2:
  NodeId2:
  Level: Process
  Comment:
4 rows in set (0.09 sec)

mcm> get :ndbd mycluster\G
***** 1. row *****
  Name: DataDir
  Value: /home/jon/bin/mcm/mcm_data/clusters/mycluster/1/data
Process1: ndbd
  NodeId1: 1
Process2:
  NodeId2:
  Level:
  Comment:
***** 2. row *****
  Name: HostName
  Value: tonfisk
Process1: ndbd
  NodeId1: 1
Process2:
  NodeId2:
  Level:
  Comment: Read only
***** 3. row *****
  Name: NodeId
  Value: 1
Process1: ndbd
  NodeId1: 1
Process2:
  NodeId2:
  Level:
  Comment: Read only
***** 4. row *****
  Name: DataDir
  Value: /home/jon/bin/mcm/mcm_data/clusters/mycluster/2/data
Process1: ndbd
  NodeId1: 2
Process2:
  NodeId2:
  Level:
  Comment:
***** 5. row *****
  Name: HostName
  Value: grindval
Process1: ndbd
  NodeId1: 2
Process2:
  NodeId2:
  Level:
  Comment: Read only
***** 6. row *****
  Name: NodeId
  Value: 2
Process1: ndbd
  NodeId1: 2
Process2:
  NodeId2:
  Level:
  Comment: Read only
6 rows in set (0.10 sec)

mcm> get :mysqld mycluster\G
***** 1. row *****
  Name: datadir
  Value: /home/jon/bin/mcm/mcm_data/clusters/mycluster/50/data

```

```

Process1: mysqld
NodeId1: 50
Process2:
NodeId2:
Level:
Comment:
***** 2. row *****
    Name: HostName
    Value: haj
Process1: mysqld
NodeId1: 50
Process2:
NodeId2:
Level:
Comment: Read only
***** 3. row *****
    Name: log_error
    Value: /home/jon/bin/mcm/mcm_data/clusters/mycluster/50/data/mysqld_50_out.err
Process1: mysqld
NodeId1: 50
Process2:
NodeId2:
Level:
Comment:
***** 4. row *****
    Name: ndb_nodeid
    Value: 50
Process1: mysqld
NodeId1: 50
Process2:
NodeId2:
Level:
Comment: Read only
***** 5. row *****
    Name: ndbcluster
    Value:
Process1: mysqld
NodeId1: 50
Process2:
NodeId2:
Level:
Comment: Read only
***** 6. row *****
    Name: NodeId
    Value: 50
Process1: mysqld
NodeId1: 50
Process2:
NodeId2:
Level:
Comment: Read only
***** 7. row *****
    Name: port
    Value: 3306
Process1: mysqld
NodeId1: 50
Process2:
NodeId2:
Level:
Comment:
***** 8. row *****
    Name: socket
    Value: /tmp/mysql.mycluster.50.sock
Process1: mysqld
NodeId1: 50
Process2:
NodeId2:
Level:
Comment:
***** 9. row *****
    Name: tmpdir
    Value: /home/jon/bin/mcm/mcm_data/clusters/mycluster/50/data/tmp

```

```

Process1: mysqld
NodeId1: 50
Process2:
NodeId2:
Level:
Comment:
***** 10. row *****
    Name: datadir
    Value: /home/jon/bin/mcm/mcm_data/clusters/mycluster/51/data
Process1: mysqld
NodeId1: 51
Process2:
NodeId2:
Level:
Comment:
***** 11. row *****
    Name: HostName
    Value: torsk
Process1: mysqld
NodeId1: 51
Process2:
NodeId2:
Level:
Comment: Read only
***** 12. row *****
    Name: log_error
    Value: /home/jon/bin/mcm/mcm_data/clusters/mycluster/51/data/mysqld_51_out.err
Process1: mysqld
NodeId1: 51
Process2:
NodeId2:
Level:
Comment:
***** 13. row *****
    Name: ndb_nodeid
    Value: 51
Process1: mysqld
NodeId1: 51
Process2:
NodeId2:
Level:
Comment: Read only
***** 14. row *****
    Name: ndbcluster
    Value:
Process1: mysqld
NodeId1: 51
Process2:
NodeId2:
Level:
Comment: Read only
***** 15. row *****
    Name: NodeId
    Value: 51
Process1: mysqld
NodeId1: 51
Process2:
NodeId2:
Level:
Comment: Read only
***** 16. row *****
    Name: port
    Value: 3307
Process1: mysqld
NodeId1: 51
Process2:
NodeId2:
Level:
Comment:
***** 17. row *****
    Name: socket
    Value: /tmp/mysql.mycluster.51.sock
    
```

```

Process1: mysqld
NodeId1: 51
Process2:
NodeId2:
Level:
Comment:
***** 18. row *****
      Name: tmpdir
      Value: /home/jon/bin/mcm/mcm_data/clusters/mycluster/51/data/tmp
Process1: mysqld
NodeId1: 51
Process2:
NodeId2:
Level:
Comment:
18 rows in set (0.05 sec)

mcm>

```



Note

You are not returned to the client prompt until the script has finished executing.

Similarly, on Windows, you can create a batch file using Notepad or another text editor, copy the same `get` commands as shown previously into it, and save it as `get-attributes.bat` in a convenient location such as the Windows desktop.

You can view a list of available `mysql` client commands using the `help` command. For more information about these, view the `help` output or see [mysql Client Commands](#), in the *MySQL Manual*.

4.2 MySQL Cluster Manager Site and Agent Commands

In this section, we discuss commands used to work with MySQL Cluster Manager management sites. In addition, the `stop agents`, `show settings`, `version`, and `show warnings` commands, which relate to management agents, are also covered in this section.

A *site*, in terms of MySQL NDB Cluster and MySQL Cluster Manager, is a collection of one or more host computers where MySQL Cluster Manager agents are running. Each agent is identified by the combination of two pieces of information:

- The hostname or IP address of the machine where the agent is running
- The number of the port used by the agent for communications



Note

MySQL NDB Cluster makes extremely intensive use of network connections, and DNS lookups can contend with MySQL NDB Cluster and MySQL Cluster Manager for bandwidth, resulting in a negative impact on the performance of MySQL NDB Cluster and the applications using it. For this reason, we recommend that you use numeric IP addresses rather than hostnames for MySQL NDB Cluster and MySQL Cluster Manager host computers whenever feasible.

4.2.1 The `add hosts` Command

```

add hosts [--hosts=|-h ]host_list site_name

host_list:
  host[, host[, ...]]

```

This command adds one or more hosts to an existing management site. Agents using the same port as the management site must be running on any hosts added using this command. This command

takes two mandatory arguments: a list of hosts (using the `--hosts` option or its short form `-h`), and the name of the site to which the hosts are to be added.

The `--hosts` takes a comma-separated list of one or more hosts to be added to the site.

For example, the following command adds two hosts named `torsk` and `kolja` to management site `mysite`:

```
mcm> add hosts --hosts=torsk,kolja mysite;
+-----+
| Command result |
+-----+
| Hosts added successfully |
+-----+
1 row in set (0.48 sec)
```

None of the hosts added by this command may already be members of management site `site_name`. Do not attempt to add again a host that is already a member of the management site using its secondary (or alternate) IP address—the `mcmd` process on the host is already bound to the IP address that was supplied when the host was first added, and it cannot be bound again to another IP address.



Notes

- This command does not support the `--force` option.
- Do not use `localhost` in the host list, as MySQL Cluster Manager relies on the operating system for host name resolution, and `localhost` might be resolved differently on different systems. Use proper host names for the host list or, preferably, use the IP addresses for the hosts instead.
- When IPv6-enabled Windows systems are used as MySQL NDB Cluster hosts under MySQL Cluster Manager, you must reference these hosts using IPv4 addresses. Otherwise, MySQL Cluster Manager will be unable to connect to the agent processes on those hosts. See [Section 5.1, “MySQL Cluster Manager Usage and Design Limitations”](#).

4.2.2 The `remove hosts` Command

```
remove hosts {--hosts=|-h }host_list site_name

host_list:
  host[, host[, ...]]
```

This command removes one or more hosts from an existing management site. It takes as arguments a required option `--hosts` (or its short form `-h`), whose value is a comma-separated list of one or more hosts to be removed, and the name of the site from which the hosts are to be removed. A number of limitations apply:

- The name of the host to be removed must not be `localhost` or `127.0.0.1`.
- The host to be removed must not have any managed processes from any clusters assigned to them (remove those processes first with the `remove process` command); it can have unmanaged processes assigned to them though (typically, `ndbapi@hostname` or `mysqld*@hostname`).
- There should not be any packages defined with explicit paths pointing to the host to be removed.
- A quorum consists of a majority of hosts (i.e., half of the total number of hosts plus one) must exist for the site both before and after the host's removal, or it will not be possible to execute the `remove host` command.
- You cannot remove the last host from a site; use the `delete site` command instead.

The following command removes two hosts named `tonfisk` and `flundra` from the management site `mysite`:

```
mcm> remove hosts --hosts=tonfisk,flundra mysite;
+-----+
| Command result |
+-----+
| Hosts removed successfully |
+-----+
1 row in set (0.48 sec)
```

4.2.3 The `change log-level` Command

```
change log-level [{--hosts=|-h }host_list] log_level site_name

host_list:
  host[,host[,...]]
```

Set the management agent's cluster logging level. This has the same effect as using the `--log-level` option; however, unlike the option, this command can be used at run time and does not require a restart of `mcmd`. Issuing this command overrides any value for `--log-level` set on the command line or in the agent configuration file.

When used with the `log_level` alone without a `host_list` and a `site_name`, this command applies only to the agent the `mcm` client is connected to. In the following example, the logging level is set to `warning` only on the host managed directly by the agent to which the `mcm` client is connected:

```
mcm> change log-level warning;
+-----+
| Command result |
+-----+
| Log-level changed successfully |
+-----+
1 row in set (0.00 sec)
```

You can specify the name of a site to be affected by the command. For example, the following invocation of the command applies to the site named `mysite`:

```
mcm> change log-level debug mysite;
+-----+
| Command result |
+-----+
| Log-level changed successfully |
+-----+
1 row in set (0.05 sec)
```

You can also restrict the change to one or more hosts in a given site using the `--hosts` option (or its short form `-h`), with multiple host names separated by commas. The following command changes the logging level to `debug` on the hosts named `tonfisk` and `haj`, but not on any other hosts in `mysite`:

```
mcm> change log-level --hosts=tonfisk,haj debug mysite;
+-----+
| Command result |
+-----+
| Log-level changed successfully |
+-----+
1 row in set (0.09 sec)
```

You must specify a site when using the `--hosts` option; trying to use `--hosts` alone results in an error.

Accepted values for `log_level` are the same as for the `--log-level`: one of `debug`, `critical`, `error`, `info`, `message`, or `warning`. For more detailed information about the meanings and effects of these values, see [NDB Cluster Logging Management Commands](#).

4.2.4 The `rotate log` Command

```
rotate log [{--hosts=|-h }host_list] [site_name]
```

```
host_list:
  host[,host[,...]]
```

Rotate `mcmd` logs for the connected MySQL Cluster Manager agent, for agents running on certain hosts, or for agents on all hosts in a management site.

For example, to rotate logs for the agent to which the client session is connected:

```
mcm> rotate log;
+-----+
| Command result |
+-----+
| Log rotated successfully |
+-----+
1 row in set (0.03 sec)
```

A new log file, with an underscore and then a timestamp inserted into its file name before the file extension, is created as a result:

```
-rw-r----- 1 mcmd cluster 74265 Jul 15 22:45 mcmd.log
-rw-r----- 1 mcmd cluster 1197573 Jul 15 22:45 mcmd_2017-07-15T22-45-28.log
```



Note

For MySQL Cluster Manager 1.4.4 and earlier, the new file name is in the format of `old_filename.timestamp`, which will be `mcmd.log.2017-07-15T22-45-28` for the example above.

To rotate logs for agents on specific hosts like `nanna12` and `nanna13`, use the `--hosts` option (or its short form `-h`):

```
mcm> rotate log --hosts=nanna12,nanna13 mysite;
```

To rotate logs on all agents in the management site `mysite`:

```
mcm> rotate log mysite;
```

4.2.5 The `collect logs` Command

```
collect logs [cluster_name]
```

This command collects log files and other related files from all hosts. When the name of a cluster (`cluster_name`) is supplied with the command, it collects all the log files (`.log`) as well as the configuration files (`.ini`, `.cnf`), error files (`.err`), and trace files (`.trace.log.*`) used by all processes belonging to the cluster, and also all the agent log files.

When an `mcmd` agent receives the `collect logs` command from the `mcm` agent it is connected with, it sets up a TCP server socket using port 0 by default, and lets the operating system assign the actual port number. All agents in the site are then instructed to perform the copying, and each of them spawns a TCP client, which connects to the TCP server socket set up earlier to copy the files over.

For MySQL Cluster Manager 1.4.2 and later: To assign a specific port manually for file copying, use the `--copy-port` option when starting `mcmd`. Default value for the option is 0. The `collect logs` command times out if, in 30 seconds, no connections can be established by any of the clients or no incoming connections are detected by the TCP server.



Warning

If a firewall or other networking issues prohibit the TCP clients to connect to the TCP server socket, the `collect logs` command will never complete.

The collected files are put under the MySQL Cluster Manager data repository (`./mcm_data` (relative to the installation directory of MySQL Cluster Manager) by default, or specified by the option `--`

`manager-directory`) inside a folder named `collected_files`, under which the files are organized under a hierarchy of timestamps [for the file collections] and then host names; below the host name level, the file organization mimics that on the host. For example, the agent log for the host `tonfisk`, collected on 2014-07-31 at `07:44:05`, is found at:

```
/opt/mcm_data/collected-files/2014-07-31T07:44:05Z/tonfisk/opt/mysql/logs/mcmd-tonfisk-19001.log
```

If `cluster_name` is omitted, only the agent log files are collected.

4.2.6 The `create site` Command

```
create site [--hosts=|-h ]host_list site_name
```

```
host_list:
  host[,host[,...]]
```

The `create site` command is used to create a MySQL Cluster Manager management site; that is, a set of MySQL Cluster Manager management agents running on one or more host computers. The command requires a list of one or more hosts where management agents are running and a name for the site. The host list is passed as the value of the `--hosts` option (short form: `-h`).

This is an example of a `create site` command that creates a site named `mysite`, consisting of the hosts `tonfisk` and `flundra`:

```
mcm> create site --hosts=tonfisk,flundra mysite;
+-----+
| Command result |
+-----+
| Site created successfully |
+-----+
1 row in set (0.31 sec)
```



Tip

You can verify that the site was created as intended, using the `list sites` command, as shown here:

```
mcm> list sites;
+-----+-----+-----+-----+
| Site | Port | Local | Hosts |
+-----+-----+-----+-----+
| mysite | 1862 | Local | tonfisk,flundra |
+-----+-----+-----+-----+
1 row in set (0.06 sec)
```

(See [Section 4.2.8, “The `list sites` Command](#)”, for more information about this command.)

Agents must be running on all hosts specified in the `--hosts` option when `create site` is executed; otherwise, the command fails with the error `Agent on host host:port is unavailable`. The host where the agent used to issue the command is running must be one of the hosts listed. Otherwise, the command fails with the error `Host host_name is not a member of site site_name`.



Warning

Moreover, if the client and the agent it is connected to are on the same host, that host must be included in the host list using its host name or its own loopback address (which can be something other than `127.0.0.1` on some systems); otherwise, the cluster might become not restartable in the future.

A given agent may be a member of one site only; if one of the management agents specified in the `host_list` already belongs to a site, the command fails with the error `Host host is already a member of site site`.

**Notes**

- Using `localhost` as the argument for the `--hosts` option will result in the creation of a single-host site (consisting of the host on which the command is run) that cannot be scaled up later by the `add hosts` command. Also notice that you cannot mix `localhost` with other host names in the host list. Therefore, it is recommended that you use IP addresses (but not any addresses belonging to the `localhost` subnet 127.*.*) or proper host names in the list.
- When IPv6-enabled Windows systems are used as MySQL NDB Cluster hosts under MySQL Cluster Manager, you must reference these hosts using IPv4 addresses. Otherwise, MySQL Cluster Manager will be unable to connect to the agent processes on those hosts. See [Section 5.1, “MySQL Cluster Manager Usage and Design Limitations”](#).

4.2.7 The `delete site` Command

```
delete site site_name
```

The `delete site` command deletes an existing management site. The command does not stop or remove any agents making up the deleted site; instead, these agents continue to run, and remain available for use in other sites.

The command takes a single argument, the name of the site to be deleted. This example shows the deletion of a management site named `mysite`:

```
mcm> delete site mysite;
+-----+
| Command result |
+-----+
| Site deleted successfully |
+-----+
1 row in set (0.38 sec)
```

If the site to be deleted does not exist, the command fails with the error `Command requires a site to be defined`. If there are any packages referencing hosts belonging to the site, `delete site` fails with the error `Packages exist in site site_name`. The command also fails if there are defined any clusters that include hosts belonging to the site.

**Note**

The management client must be connected to a site in order to be able to delete it.

In addition, if you execute a `delete site` command with the `--force` option using one management agent while a different management agent is not running, you must remove the “missing” management agent's site files manually. For more information on site files, see [Section 2.4, “MySQL Cluster Manager Configuration File”](#).

4.2.8 The `list sites` Command

```
list sites
```

This command returns a list of the sites known to the management agent. It does not require any arguments. An example is shown here:

```
mcm> list sites;
+-----+-----+-----+-----+
| Site | Port | Local | Hosts |
+-----+-----+-----+-----+
| mysite | 1862 | Local | tonfisk,flundra |
+-----+-----+-----+-----+
```

```
1 row in set (0.06 sec)
```

The output of `list sites` contains the following columns:

- **Site.** The name of the site.
- **Port.** The TCP/IP port used for communications between clients and management agents.
- **Local.** Either one of `Local` or `Remote`.
- **Hosts.** A comma-separated list of the hosts making up the site.

4.2.9 The `list hosts` Command

```
list hosts [--verbose|-v] site_name
```

The `list hosts` command is used to obtain a list of the hosts comprising a given management site. The command requires a single argument, the name of the site to be examined. For each host listed, the information returned includes the hostname, status, and version of the management agent software in use, as shown in this example:

```
mcm> list hosts mysite;
+-----+-----+-----+
| Host      | Status   | Version |
+-----+-----+-----+
| tonfisk   | Available| 1.4.8   |
+-----+-----+-----+
| flundra   | Available| 1.4.8   |
+-----+-----+-----+
2 rows in set (0.16 sec)
```

`Status` can be one of :

- **Available:** Agent on the host is active
- **Recovery:** Agent on the host is in the process of recovering itself (for MySQL Cluster Manager 1.4.7 and later)
- **Unresponsive:** Agent on the host rejected an attempt to connect (for MySQL Cluster Manager 1.4.7 and later).
- **Unavailable:** Agent on the host is unreachable

If an agent is reported persistently as `Unresponsive` or `Unavailable`, you may have to restart it.

If you omit the `site_name` argument, the command fails with an error, as shown here:

```
mcm> list hosts;
ERROR 6 (00MGR): Illegal number of operands
```

Using the `--verbose` option (short form : `-v`) causes the command to print additional information on the hosts:

```
mcm> list hosts --verbose mysite;
+-----+-----+-----+-----+-----+-----+
| Host      | Status   | Version | Cores | Memory | OS                               |
+-----+-----+-----+-----+-----+-----+
| tonfisk   | Available| 1.4.8   | 1     | 1819 Mb | Linux 3.13.11-100.fc19.x86_64 |
| flundra   | Available| 1.4.8   | 1     | 1819 Mb | Linux 3.13.11-100.fc19.x86_64 |
+-----+-----+-----+-----+-----+-----+
2 rows in set (0.07 sec)
```

4.2.10 The `show settings` Command

```
show settings [--hostinfo]
```

This command lists the current values of a number of `mcmd` options:

```
mcm> show settings;
```

```
+-----+-----+
| Setting      | Value      |
+-----+-----+
| log-file     | /opt/mcm_data/mcmd.log |
| log-level    | message    |
| log-use-syslog | FALSE      |
| manager-directory | /opt/mcm_data |
| manager-username | mcmd      |
| manager-password | ****      |
| manager-port  | 1862      |
| xcom-port    | 18620     |
+-----+-----+
8 rows in set (0.00 sec)
```

Using the `--hostinfo` option makes the command print out information on the host that the `mcm` client is connected to:

```
mcm> show settings --hostinfo;
+-----+-----+
| Property     | Value      |
+-----+-----+
| Hostname     | localhost.localdomain |
| Platform     | Linux 3.13.11-100.fc19.x86_64 |
| Processor cores | 1          |
| Total memory  | 1819 Mb   |
+-----+-----+
4 rows in set (0.00 sec)
```

4.2.11 The `stop agents` Command

```
stop agents [[--hosts=host_list] site_name]
```

This command stops one or more MySQL Cluster Manager agents on one or more hosts.

When used without any arguments, `stop agents` stops the agent to which the client is currently connected.

When used with the name of a management site, the command stops all agents running on hosts making up the site. The following stops all MySQL Cluster Manager agents running on hosts in `mysite`:

```
mcm> stop agents mysite;
```

You can also stop a subset of the agents in a given management site by listing the hosts where they are running with the `--hosts` option (short form: `-h`), along with the name of the site to which they belong. The result of the following command is to stop MySQL Cluster Manager agents running on hosts `kolja` and `torsk`, both of which are members of the management site `mysite`:

```
mcm> stop agents --hosts=kolja,torsk mysite;
```

Multiple host names following the `--hosts` option should be separated by commas, with no intervening spaces. Invoking `stop agents` with this option without supplying a `site_name` causes a syntax error. Using an undefined `site_name` or names of hosts not belonging to the site with this command also results in an error.



Note

When IPv6-enabled Windows systems are used as MySQL NDB Cluster hosts under MySQL Cluster Manager, you must reference these hosts using IPv4 addresses. Otherwise, MySQL Cluster Manager will be unable to connect to the agent processes on those hosts. See [Section 5.1, “MySQL Cluster Manager Usage and Design Limitations”](#).

4.2.12 The `version` Command

```
version
```

This command displays the version of the MySQL Cluster Manager software in use by the MySQL Cluster Manager agent to which this client is connected, as shown here:

```
mcm> version;
+-----+
| Version |
+-----+
| MySQL Cluster Manager 1.4.8 (64bit) |
+-----+
1 row in set (0.00 sec)
```

The `version` command takes no arguments.

4.2.13 The `show warnings` Command

Using the `show warnings` command, you can check the warnings (up to the last five) issued to the agent log (`mcmd.log`). For example:

```
mcm> set delayed_insert_timeout:mysqld=400 mycluster;
+-----+
| Command result |
+-----+
| Cluster reconfigured successfully |
+-----+

mcm> show warnings;
+-----+-----+-----+
| Level | Code | Message |
+-----+-----+-----+
| Warning | -1 | Config variable delayed_insert_timeout was deprecated in mysql 5.6.7 |
+-----+-----+-----+
```

4.3 MySQL Cluster Manager Package Commands

This section contains information about MySQL Cluster Manager client commands used to register, extend, unregister, and obtain information about the software packages making up instances of MySQL NDB Cluster that are to be managed using the MySQL Cluster Manager.

4.3.1 The `add package` Command

```
add package [--basedir=-b ]path
            [--hosts=-h ]host_list package_name

host_list:
    host[,host[,...]]
```

This command creates a new package, or, if the package named `package_name` already exists, this command extends the package definition.

The `--basedir` option (short form: `-b`) indicates the location of the MySQL NDB Cluster installation directory on the listed hosts, and is required. This must be the path to the top-level directory where the MySQL NDB Cluster software is located (for example, `/usr/local/mysql`), and should *not* include the MySQL NDB Cluster `bin`, `libexec`, or other subdirectory within the installation directory.

Hosts may be specified as a comma-separated list, using the `--hosts` option (short form: `-h`); however, this option is not required. If `--hosts` is omitted, the `path` is assumed to be valid for all hosts in the cluster that is created using this package (see [Section 4.4.1, “The create cluster Command”](#)).



Important

- You cannot perform `add package` if you have not yet defined any sites (each host referenced in an `add package` command must be associated with a site). See [Section 4.2.6, “The create site Command”](#), for more information about defining sites.

- When a package is first added for a site with the `add package` command, whenever the `--hosts` option is used, the host list must contain the host for the `mcmd` agent to which the `mcm` client is currently connected, in order to allow the MySQL Cluster Manager to access the version information of the package.

Suppose we have two Linux hosts named `tonfisk` and `flundra`, and the MySQL NDB Cluster software is installed in `/usr/local/mysql` on both hosts. In this case, you can create a package named `mypackage` that accounts for both hosts as shown here:

```
mcm> add package --basedir=/usr/local/mysql mypackage;
+-----+
| Command result |
+-----+
| Package added successfully |
+-----+
1 row in set (0.71 sec)
```

When this package is used to create a cluster, the MySQL Cluster Manager knows that it should find the MySQL NDB Cluster software in the `/usr/local/mysql` directory on each of the hosts.

For options to MySQL Cluster Manager client command options having Windows paths as values, you must use forward slashes (/) in place of backslashes (\), so if `tonfisk` and `flundra` are Windows hosts where MySQL NDB Cluster has been installed to the directory `C:\mysql`, the corresponding `add package` command would look like this (with the `--basedir` option highlighted):

```
mcm> add package --basedir=c:/mysql mypackage;
+-----+
| Command result |
+-----+
| Package added successfully |
+-----+
1 row in set (0.71 sec)
```

In the example just given, we could also have issued the command as `add package --basedir=/usr/local/mysql --hosts=tonfisk,flundra mypackage` (or `add package --basedir=c:/mysql --hosts=tonfisk,flundra mypackage` on Windows) with the same result, but the `--hosts` option was not required, since the MySQL NDB Cluster software's location is the same on each host. Let us suppose, however, that the software is installed in `/usr/local/ndb-host-10` on host `tonfisk` and in `/usr/local/ndb-host-20` on host `flundra`. In this case, we must issue 2 separate commands, specifying the host as well as the base directory in each case, as shown here:

```
mcm> add package --basedir=/usr/local/ndb-host-10
> --hosts=tonfisk yourpackage;
+-----+
| Command result |
+-----+
| Package added successfully |
+-----+
1 row in set (0.68 sec)

mcm> add package --basedir=/usr/local/ndb-host-20
> --hosts=flundra yourpackage;
+-----+
| Command result |
+-----+
| Package added successfully |
+-----+
1 row in set (0.81 sec)
```

Assuming that both hosts belong to a site called `mysite`, you can verify that these packages have been created as desired using the `list packages` command, as shown here:

```
mcm> list packages mysite;
+-----+-----+-----+
| Package | Path | Hosts |
+-----+-----+-----+
```

```

+-----+-----+-----+
| yourpackage | /usr/local/ndb-host-10 | tonfisk |
| mypackage   | /usr/local/ndb-host-20 | flundra |
|             | /usr/local/mysql      | tonfisk,flundra |
+-----+-----+-----+
3 rows in set (1.07 sec)

```

(For more information about this command, see [Section 4.3.3, “The `list packages` Command”](#).)

It is possible to assign the same base directory (or directories) on the same host (or hosts) to multiple packages, as shown in this example, in which we assume that hosts `tonfisk` and `flundra` have previously been assigned to a site named `mysite`:

```

mcm> add package -b /usr/local/mysql-cluster mypackage;
+-----+-----+
| Command result |
+-----+-----+
| Package added successfully |
+-----+-----+
1 row in set (1.41 sec)

mcm> add package -b /usr/local/mysql-cluster yourpackage;
+-----+-----+
| Command result |
+-----+-----+
| Package added successfully |
+-----+-----+
1 row in set (1.58 sec)

mcm> list packages mysite;
+-----+-----+-----+
| Package | Path | Hosts |
+-----+-----+-----+
| mypackage | /usr/local/mysql-cluster | tonfisk,flundra |
| yourpackage | /usr/local/mysql-cluster | tonfisk,flundra |
+-----+-----+-----+
2 rows in set (0.50 sec)

```



Note

When IPv6-enabled Windows systems are used as MySQL NDB Cluster hosts under MySQL Cluster Manager, you must reference these hosts using IPv4 addresses. Otherwise, MySQL Cluster Manager will be unable to connect to the agent processes on those hosts. See [Section 5.1, “MySQL Cluster Manager Usage and Design Limitations”](#).

4.3.2 The `delete package` Command

```

delete package [--hosts=-h ]host_list package_name

host_list:
    host[,host[,...]]

```

This command is used to unregister a package. More specifically, it removes any references to MySQL NDB Cluster software installations added to the agent's repository when the package was created. `delete package` does *not* remove any MySQL NDB Cluster installations; the command removes only references to the installations. Once a package has been unregistered, it can no longer be used for a `create cluster` command. The MySQL NDB Cluster binaries remain, but cannot be used in a MySQL NDB Cluster administered using the MySQL Cluster Manager unless and until the base directory containing them has been registered with another package. (It is possible to register a base directory with multiple packages; see [Section 4.3.1, “The `add package` Command”](#), for more information and an example.)

If the `--hosts` option (short form: `-h`) is used with this command, the base directory settings for the host or hosts named by the option are removed as well. All hosts given in the `host_list` must be members of the site to which the package is registered. Otherwise, the command fails.

A package that is in use by a cluster cannot be unregistered; the cluster must first be deleted (see [Section 4.4.2, “The `delete cluster` Command”](#)).

Here is an example that demonstrates how to unregister a package named `mypackage`:

```
mcm> delete package mypackage;
+-----+
| Command result |
+-----+
| Package deleted successfully |
+-----+
1 row in set (1.23 sec)
```

You can also verify that the package was unregistered using the `list packages` command; the package name should no longer appear in the output of this command. If you attempt to use the unregistered package in a `create cluster` command, the command fails, as shown here:

```
mcm> create cluster --package=mypackage
> --processhosts=ndb_mgmd@tonfisk,ndbd@grindval,ndbd@flundra,mysqld@tonfisk mycluster;
ERROR 4001 (00MGR): Package mypackage not defined
```

An `upgrade cluster` command that references an unregistered package also fails.



Note

When IPv6-enabled Windows systems are used as MySQL NDB Cluster hosts under MySQL Cluster Manager, you must reference these hosts using IPv4 addresses. Otherwise, MySQL Cluster Manager will be unable to connect to the agent processes on those hosts. See [Section 5.1, “MySQL Cluster Manager Usage and Design Limitations”](#).

4.3.3 The `list packages` Command

```
list packages [package_name] site_name
```

This command lists registered packages. It requires a single argument, that being the name of the site with which the packages are registered, as shown in this example:

```
mcm> list packages mysite;
+-----+-----+-----+
| Package | Path | Hosts |
+-----+-----+-----+
| yourpackage | /usr/local/ndb-host-10 | tonfisk |
| | /usr/local/ndb-host-20 | flundra |
| mypackage | /usr/local/mysql | tonfisk,flundra |
+-----+-----+-----+
3 rows in set (1.07 sec)
```

If `tonfisk` and `flundra` are Windows hosts, the list of packages might look something like this:

```
mcm> list packages mysite;
+-----+-----+-----+
| Package | Path | Hosts |
+-----+-----+-----+
| yourpackage | c:/cluster/ndb-host-10 | tonfisk |
| | c:/cluster/ndb-host-20 | flundra |
| mypackage | c:/mysql | tonfisk,flundra |
+-----+-----+-----+
3 rows in set (1.07 sec)
```

In the example just shown, `yourpackage` uses the MySQL NDB Cluster binaries installed at `C:\cluster\ndb-host-10` on host `tonfisk`, and at `C:\cluster\ndb-host-20` on `flundra`; `mypackage` uses MySQL NDB Cluster binaries installed at `C:\mysql` on both hosts.

The output contains three columns; these are described in the following list:

- **Package.** The name of the package. This can sometimes be empty when a package includes MySQL NDB Cluster installations that are in different locations on different hosts (see next example).
- **Path.** The path to the MySQL NDB Cluster installation directory (base directory) on the indicated host or hosts. This is the same as the value given for the `--basedir` option in the `add package` command that was used to create or augment the package.

On Windows, paths shown in this column have any backslash characters converted to forward slashes, just as must be done for the `--basedir` option (see the earlier example in this section).

- **Hosts.** The host or hosts where the MySQL NDB Cluster installation or installations are located.

You can filter the results so that information relating to only a single package is displayed by supplying the package name before the site name, as shown here:

```
mcm> list packages yourpackage mysite;
+-----+-----+-----+
| Package | Path | Hosts |
+-----+-----+-----+
| yourpackage | /usr/local/ndb-host-10 | tonfisk |
| | /usr/local/ndb-host-20 | flundra |
+-----+-----+-----+
2 rows in set (0.55 sec)
```

(See [Section 4.3.1, “The add package Command”](#), for the `add package` commands that were used to create `yourpackage`.)

When a package contains MySQL NDB Cluster installations using different base directories on different hosts, each unique combination of path and host is shown in its own row. However, the name of the package is displayed in the first row only; all rows that immediately follow this row and that do not contain the package name also relate to the same package whose name is shown in the first preceding row to display a package name. For example, consider the `list packages` command and output shown here:

```
mcm> list packages mysite;
+-----+-----+-----+
| Package | Path | Hosts |
+-----+-----+-----+
| yourpackage | /usr/local/ndb-host-10 | tonfisk |
| | /usr/local/ndb-host-20 | flundra |
| mypackage | /usr/local/mysql | tonfisk |
| | /usr/local/bin/mysql | flundra |
+-----+-----+-----+
3 rows in set (1.07 sec)
```

This output shows that there are two packages defined for the site named `mysite`; these packages are named `yourpackage` and `mypackage`. The package `yourpackage` consists of the MySQL NDB Cluster binaries in the directory `/usr/local/ndb-host-10` on host `tonfisk`, and in the directory `/usr/local/ndb-host-20` on host `flundra`. The package named `mypackage` consists of the MySQL NDB Cluster binaries in the directory `/usr/local/mysql` on host `tonfisk`, and in the directory `/usr/local/bin/mysql` on host `flundra`.

If you omit the `site_name` argument, the command fails with an error, as shown here:

```
mcm> list packages;
ERROR 6 (00MGR): Illegal number of operands
```

4.4 MySQL Cluster Manager Cluster Commands

This section contains descriptions of MySQL Cluster Manager commands used to perform operations on clusters. These include creating and deleting a cluster; starting, stopping, and restarting a cluster; upgrading a cluster (that is, upgrading the MySQL NDB Cluster software used by a given cluster); and listing clusters known to MySQL Cluster Manager.

4.4.1 The `create cluster` Command

```
create cluster [--package=-P ]package_name
  [--processhosts=-R ]process_host_list cluster_name
  [--import|-m ]cluster_name [--verbose | -v]

process_host_list:
  process_name[ :node_id]@host[ ,process_name@host[ ,... ] ]

process_name:
  {ndb_mgmd | ndbd | ndbmtd | mysqld | ndbapi}
```

This command creates a cluster to be managed by the MySQL Cluster Manager. However, it does not start the cluster (see [Section 4.4.7, “The `start cluster` Command”](#)).

This command can also be used to create a cluster earmarked specifically as a target for importing another cluster that is not already under MySQL Cluster Manager control, as described later in this section, by employing the `--import` option. See also [Section 3.5, “Importing MySQL NDB Clusters into MySQL Cluster Manager”](#).

`create cluster` requires the following arguments:

- A *package_name*, supplied as the value of the `--package` option (short form: `-P`). This must be the name of a package previously registered using `add package`.
- A list (*process_host_list*) of MySQL NDB Cluster processes, the hosts on which they are to run, and—optionally—their node IDs, supplied as the value of the `--processhosts` option (short form: `-R`), with list items separated by commas. As with other lists passed as option values in MySQL Cluster Manager commands, you must not use spaces before or after the commas.

Each item in the *process_host_list* consists of the name of a MySQL NDB Cluster process—possibly suffixed with a colon (:) character followed by the process node ID—joined with the name of the host on which it is located using an amphora (@) sign (also sometimes known as the “at” sign). Permitted values for processes are `ndb_mgmd`, `ndbd`, and `mysqld`. When the cluster employs MySQL NDB Cluster 7.0 or later, you can also use `ndbmt`*d* as process name; in other words, a valid process name is the name of a MySQL NDB Cluster process daemon binary. If node IDs are specified, they must be within the allowed range for the type of node defined.

To support running your own NDB API applications with a cluster under MySQL Cluster Manager, it is also possible to use `ndbapi` as a process type. Such applications can be connected to a managed cluster. Currently, MySQL Cluster Manager recognises only that an NDB API application is connected to the cluster; the NDB API application itself must be started, stopped, and configured manually.

It is also possible to specify one or more “free” `mysqld` and `ndbapi` processes without any hosts. To do this, simply use the wildcard * (asterisk character) in place of the hostname or IP address, as shown below:

- “Free” `mysqld` process: `mysqld@*`
- “Free” `ndbapi` process: `ndbapi@*`

It is also possible to specify a node ID for a “free” process. (If this is not specified, MySQL Cluster Manager assigns a suitable node ID automatically.)

A `mysqld` process or `ndbapi` process that is specified without a host in this fashion is permitted to connect to the cluster from any host that can access the cluster over the network. Otherwise, the process may connect to the cluster only from the specified host.

By convention, items in the *process_host_list* are listed according to the process type, in the following order:

1. Management node processes (`ndb_mgmd`)

2. Data node processes (*MySQL NDB Cluster 6.3*: `ndbd`; *MySQL NDB Cluster 7.0 and later*: `ndbd`, `ndbmtid`)
3. SQL node processes (`mysqld`)
4. Custom NDB API applications (`ndbapi`)

For information about writing your own NDB API applications, see [The NDB API](#), in the *MySQL NDB Cluster API Developer Guide*.

While the order in which the items are listed does not affect whether the `create cluster` command succeeds, we suggest that you follow this convention for readability, as well as compatibility with other MySQL NDB Cluster management tools such as `ndb_mgm`.

`create cluster` causes cluster node IDs to be assigned consecutively, in the order that the nodes are specified in the `process_host_list`, with node IDs for data node processes starting with 1, and node IDs for processes other than data node processes starting with 49. For MySQL Cluster Manager 1.3.3 and earlier, trying to manually assign node IDs less than 49 for `ndb_mgmd`, `mysqld`, or `ndbapi` fails with an error; the restriction, however, has been totally lifted for MySQL Cluster Manager 1.3.4 and later. Nevertheless, you are still recommended to follow the best practice of reserving node ID 1 to 48 for data nodes.



Note

NDB 8.0 supports up to 144 data nodes (for release 8.0.18 and later). Therefore, `create cluster` assigns node IDs for data node processes starting with 1, and node IDs for processes other than data node processes starting with 145. When manually assigning node IDs for an NDB 8.0 cluster, you are recommended to follow the best practice of reserving node ID 1 to 144 for data nodes.

Each host referenced in the list must be part of the site for which the package used in `create cluster` is defined.

For processes of types `mysqld` and `ndbapi`, the hostname is required, but not enforced in the running cluster. In other words, an `[api]` section is created in the cluster `config.ini` file, but no `HostName` parameter is specified; thus, the `mysqld` or `ndbapi` can connect from any host. (Currently, there is no way using MySQL Cluster Manager to specify that a `mysqld` or `ndbapi` process is restricted to connecting from a single host.)

- A name for the cluster. Once the cluster has been created, this name is used to refer to it in other cluster management commands such as `delete cluster`, `start cluster`, and `stop cluster`. Like other object names used with MySQL Cluster Manager, the `cluster_name` must be valid according to the rules given elsewhere in this document for identifiers (see [Chapter 4, MySQL Cluster Manager Client Commands](#)).

An additional `--verbose` option for this command causes `create cluster` to output extra information as it is executed, as shown later in this section.

The `--import` option flags the cluster as being created as a target for importing a cluster created outside MySQL Cluster Manager. This option causes the cluster's status to appear as `import` in the output of `show status`, as shown here:

```
mcm> show status --process newcluster;
```

NodeId	Process	Host	Status	Nodegroup	Package
1	ndb_mgmd	alpha	import		newpackage
5	ndbd	beta	import	n/a	newpackage
6	ndbd	gamma	import	n/a	newpackage

```

| 10 | mysql | delta | import | | newpackage |
| 11 | ndbapi | * | import | | |
+-----+-----+-----+-----+-----+
6 rows in set (0.04 sec)

```

Having the `import` status causes any of the commands `start cluster`, `restart cluster`, `start process`, and `stop process` to fail if they are executed before an `import cluster` command has been executed against this cluster. It is also not possible to execute `upgrade cluster` on a cluster having processes with `import` status. Other operations on this cluster continue to be performed normally.



Caution

While it is possible to import into a cluster that was created without this option, it is not advisable, since the cluster is not protected against accidentally performing any of the operations listed previously, which may result in confusing or misleading errors, and possibly other problems. For this reason, it is strongly recommended that you always use the `--import` option for creating the cluster in such cases.

For more information about importing clusters into MySQL Cluster Manager, including examples, see [Section 3.5, “Importing MySQL NDB Clusters into MySQL Cluster Manager”](#).

Example

Consider the following command issued in the MySQL Cluster Manager client, which creates a cluster named `mycluster`:

```

mcm> create cluster --package=myspackage
-> --processhosts=ndb_mgmd@flundra,ndbd@tonfisk,ndbd@grindval,mysql@flundra
-> mycluster;
+-----+-----+
| Command result |
+-----+-----+
| Cluster created successfully |
+-----+-----+
1 row in set (7.71 sec)

```

As defined by the command just shown, `mycluster` consists of four nodes: a management node on host `flundra`; two data nodes—one on each of the hosts `tonfisk` and `grindval`; and one SQL node, also on host `flundra`.

Using the `--verbose` option causes the command to print output similar to that produced by the `list processes` command, as shown here:

```

mcm> create cluster --verbose --package=myspackage
-> --processhosts=ndb_mgmd@flundra,ndbd@tonfisk,ndbd@grindval,mysql@flundra
-> mycluster;
+-----+-----+-----+
| NodeId | Name      | Host      |
+-----+-----+-----+
| 49     | ndb_mgmd | flundra   |
| 1      | ndbd     | tonfisk   |
| 2      | ndbd     | grindval  |
| 50     | mysql    | flundra   |
+-----+-----+-----+
4 rows in set (0.32 sec)

```

You can also create this cluster in such a way that the `mysql` process is permitted to connect to the cluster from any host able to reach the other cluster hosts over the network as shown here:

```

mcm> create cluster --package=myspackage
-> --processhosts=ndb_mgmd@flundra,ndbd@tonfisk,ndbd@grindval,mysql@*
-> mycluster;
+-----+-----+
| Command result |
+-----+-----+

```

```

+-----+
| Cluster created successfully |
+-----+
1 row in set (7.71 sec)

```

**Note**

In the case of a “free” `ndbapi` process, it is not necessary to have the MySQL Cluster Manager software installed on the host where the `ndbapi` process is running.

Configuration changes to the newly-created cluster can be made using the `set` command prior to starting the cluster. This is often preferable to doing after the cluster has been started, since `set` commands used to make configuration changes in a running cluster can require a rolling restart, and rolling restarts of clusters having many nodes or large quantities of data (or both) may take a great deal of time to complete.

**Note**

When creating a cluster having more than one `mysqld` process on the same host machine, MySQL Cluster Manager assigns the MySQL default port (3306) to each of them. Therefore, you must assign a unique port for each `mysqld` process in the cluster.

4.4.2 The `delete cluster` Command

```
delete cluster [--removedirs] cluster_name
```

This command deletes the cluster named `cluster_name`, removing it from the list of clusters managed by MySQL Cluster Manager.

`delete cluster` does *not* remove any MySQL NDB Cluster binaries from hosts. However, it *does* remove the cluster configuration, data, and log files that reside in the MySQL Cluster Manager data repository.

This example demonstrates how to delete a cluster named `mycluster`:

```

mcm> delete cluster mycluster;
+-----+
| Command result          |
+-----+
| Cluster deleted successfully |
+-----+
1 row in set (1.22 sec)

```

A look at the MySQL Cluster Manager data repository (at `/opt/mcm_data/` in this case) shows that the folder that used to host the configuration, data, and log files for `mycluster` (`/opt/mcm_data/clusters/mycluster`) no longer exists:

```

$> ls -l /opt/mcm_data/clusters
total 0

```

To remove the configuration and data files outside of the MySQL Cluster Manager data repository, `delete cluster` must be invoked with the `--removedirs` option, like this:

```

mcm> delete cluster --removedirs mycluster;
+-----+
| Command result          |
+-----+
| Cluster deleted successfully |
+-----+
1 row in set (1.22 sec)

```

For example, if one of the data node on `mycluster` has its data directory outside of the MySQL Cluster Manager data repository:

```
mcm> get Datadir mycluster;
+-----+-----+-----+-----+-----+-----+-----+
| Name   | Value                               | Process1 | NodeId1 | Process2 | NodeId2 | Level  | Comment |
+-----+-----+-----+-----+-----+-----+-----+
| DataDir | /home/dso/mycluster/cdata          | ndbd     | 1       |           |          | Process |         |
...

```

Deleting mycluster without using `--removedirs` does *not* remove the data directory for node 1:

```
$> ls -l /home/dso/mycluster
total 4 drwxr-xr-x. 3 dso dso 4096 Sep 10 18:00 cdata
```

However, if the `--removedirs` option is used, the data directory for node 1 also gets removed:

```
$> ls -l /home/dso/mycluster
total 0
```

`delete cluster` fails if the cluster to be deleted is running, as shown here:

```
mcm> delete cluster mycluster;
ERROR 5010 (OOMGR): All processes must be stopped to delete cluster mycluster
```

You must shut down the cluster first, using `stop cluster`.

4.4.3 The `list clusters` Command

```
list clusters site_name
```

This command lists all clusters defined for a given management site named `site_name`, together with the package used by each cluster. For example, the command shown here displays a list of all clusters defined for the site named `mysite`:

```
mcm> list clusters mysite;
+-----+-----+
| Cluster      | Package |
+-----+-----+
| mycluster    | m-7.1.26 |
| yourcluster  | y-7.1.26 |
| someothercluster | s-7.2.9 |
+-----+-----+
3 rows in set (2.07 sec)
```

If `site_name` is omitted, the command fails with an error, as shown here:

```
mcm> list clusters;
ERROR 6 (OOMGR): Illegal number of operands
```

4.4.4 The `list nextnodeids` Command

```
list nextnodeids cluster_name
```

MySQL Cluster Manager normally assigns IDs to new node processes automatically (although this can be overridden when issuing the `create cluster` or `add process` command). The `list nextnodeids` command can be used to see the next node ID that MySQL Cluster Manager has reserved for the next new process (of each possible process type) to be added to the cluster named `cluster_name`, which is required.

```
mcm> list nextnodeids mycluster;
+-----+-----+-----+-----+
| Category  | NodeId Range | Next NodeId | Processes |
+-----+-----+-----+-----+
| Datanodes | 1 - 48       | 5           | ndbd, ndbmt |
| Others    | 49 - 255    | 52         | ndb_mgmd, mysql |
+-----+-----+-----+-----+
2 rows in set (0.07 sec)
```

4.4.5 The `restart cluster` Command

```
restart cluster [--sequential-restart] cluster_name
```

This command performs a rolling restart (see [Performing a Rolling Restart of an NDB Cluster](#)) of the cluster named *cluster_name*. The cluster must already be running in order for this command to succeed. (For information about how to determine the operation state of the cluster, see [Section 4.4.6, “The show status Command”](#).)

For example, the command shown here performs a rolling restart of the cluster named `mycluster`:

```
mcm> restart cluster mycluster;
+-----+
| Command result |
+-----+
| Cluster restarted successfully |
+-----+
1 row in set (1 min 22.53 sec)
```

If the cluster is not already running, `restart cluster` fails with an error, as shown here:

```
mcm> show status --cluster mycluster;
+-----+-----+-----+
| Cluster | Status | Comment |
+-----+-----+-----+
| mycluster | stopped | |
+-----+-----+-----+
1 row in set (1.49 sec)

mcm> restart cluster mycluster;
ERROR 5009 (OOMGR): Restart can not be performed as processes are
stopped in cluster mycluster
```

For MySQL Cluster Manager 1.4.8 and later: By default, a rolling restart is performed on the nodes in a *parallel* manner (that is, half of the nodes are stopped and restarted together, followed by the second half of the nodes). In some situations, you might want to have a rolling restart performed in a *sequential* manner by adding the `--sequential-restart` option, in which case nodes are stopped and restarted one after another.

For MySQL Cluster Manager 1.4.7 and earlier: A rolling restart is performed on the nodes in a *parallel* manner (i.e., half of the nodes are stopped and restarted together, followed by the second half of the nodes).



Note

Depending on the number of nodes and the amount of data stored in the cluster, a rolling restart can take a considerable amount of time, up to several hours for a cluster with a great many data nodes and a large amount of data. For this reason, you may want to execute this command with the `--background` option (short form `-B`) to allow it to run in the background, freeing the MySQL Cluster Manager client for other tasks.

Currently, there is no mechanism in MySQL Cluster Manager for performing system initial restarts of a cluster. This means that attributes that require an initial restart to be changed must be set before the cluster is started for the first time.

4.4.6 The `show status` Command

```
show status --cluster|-c cluster_name
show status --operation|-o cluster_name
show status --backup|-b cluster_name
show status --process|-r cluster_name
show status --progress cluster_name
show status --progressbar cluster_name
```

This command is used to check the status of clusters, cluster processes, backups, and commands issued in the MySQL Cluster Manager client. The type of status returned depends on which of the four

options of `--cluster` (short form: `-c`), `--operation` (short form: `-o`), `--backup` (short form: `-b`), or `--process` (short form `-r`) is used with the command. (If none of these is used, `--cluster` is assumed.) These options are described in more detail in the next few paragraphs.

`--cluster` option

When this option is used, `show status` reports on the status of the cluster named `cluster_name`, as shown in this example:

```
mcm> show status --cluster mycluster;
+-----+-----+-----+
| Cluster | Status           | Comment |
+-----+-----+-----+
| mycluster | fully operational |         |
+-----+-----+-----+
1 row in set (0.01 sec)
```

When used with the `--cluster` option (short form: `-c`), the output of this command consist of two columns. The `Cluster` column contains the name of the cluster. The `Status` column contains a description of the cluster's status; possible values and their meanings are shown in the following table:

Table 4.1 Status values shown by `show status --cluster`

Status Value	Meaning
<code>fully operational</code>	All cluster processes are running.
<code>operational</code>	All node groups are up and running, but at least one data node process (<code>ndbd</code> or <code>ndbmtid</code>) is not running. The cluster is online, but you should determine why any “missing” data nodes are not running and correct the problem as soon as possible.
<code>non-operational</code>	The cluster is not operational, because at least one node group is offline. You must investigate and fix the problem or problems, then restart the cluster, before the cluster can be used for data storage and retrieval operations.
<code>stopped</code>	The cluster is not running, because it has been stopped by the user. This normally does not indicate any problem as such, but you must restart the cluster before it can be used by any applications.
<code>created</code>	The cluster has been created successfully using the <code>create cluster</code> command, but has never been started. You must start the cluster using the <code>start cluster</code> command before you can make use of it.
<code>unknown</code>	The MySQL Cluster Manager was unable to determine the cluster's status. This may or may not indicate a problem with the cluster; it is possible that the problem lies with one or more MySQL Cluster Manager agents or the MySQL Cluster Manager client. You should attempt to determine the status of the cluster by other means, such as using <code>show status --process</code> in the MySQL Cluster Manager client (described later in this section), or employing one of the commands available in the <code>ndb_mgm</code> client (see ndb_mgm — The NDB Cluster Management Client) such as <code>SHOW</code> or <code>ALL STATUS</code> .

`--operation` option

When the `--operation` option (short form: `-o`) is used, it causes `SHOW STATUS` to display the status of the latest command to be executed. This includes commands that were issued using the `--background` option (short form `-B`). An example of this command is shown here:

```
mcm> show status --operation mycluster;
+-----+-----+-----+
| Command          | Status   | Description |
+-----+-----+-----+
| restart cluster | executing| <no message>|
+-----+-----+-----+
1 row in set (1.60 sec)
```

The output contains 3 columns, described in the following list:

- **Command.** The text of the command last issued (previous to the `show status --operation` command), less any options or arguments.
- **Status.** The current state of the command. Possible values and their meanings are listed later in this section.
- **Description.** In some cases, depending on the command and its status, this column may contain additional information. Otherwise, `<no message>` is displayed here.

Possible values for the `Status` column, together with descriptions of these values, are shown in the following table:

Table 4.2 Status values shown by `show status --operation`

Status Value	Description
<code>executing</code>	MySQL Cluster Manager is executing the command, but has not yet completed doing so.
<code>finished</code>	The command has executed (and completed) successfully.
<code>failed</code>	The command failed to execute. The <code>Description</code> column may contain information about the reason for the failure.
<code>unknown</code>	MySQL Cluster Manager was unable to determine the status of this command.

`--backup` option

When this option is used, `show status` reports on the status of the backup process for the cluster named `cluster_name`, as shown in the following examples:

```
mcm> show status --backup mycluster;
+-----+-----+
| Command result |
+-----+-----+
| No backup currently active in mycluster |
+-----+-----+
1 row in set (0.05 sec)
```

```
mcm> show status --backup mycluster;
+-----+-----+
| Command result |
+-----+-----+
| BackupId 5 currently active in mycluster|
+-----+-----+
1 row in set (0.09 sec)
```

`--process` option

When run with this option, `show status` returns information about each process in the cluster named `cluster_name`, as shown in this example:

```
mcm> show status --process mycluster;
+-----+-----+-----+-----+-----+
| Id   | Process | Host   | Status | Nodegroup |
+-----+-----+-----+-----+-----+
| 1   | ndb_mgmd | tonfisk | running |           |
| 2   | ndbd     | flundra | running | 0         |
| 3   | ndbd     | grindval | running | 0         |
| 4   | mysqld   | lax    | running |           |
+-----+-----+-----+-----+-----+
4 rows in set (1.67 sec)
```

When the `--process` option (short form: `-r`) is used with `show status`, the output contains 5 columns, described in the following list:

- **Id.** This is the node ID of the process as a node in cluster `cluster_name`.
- **Process.** The type of process, that is, the name of the corresponding MySQL NDB Cluster executable. Allowed values are `ndb_mgmd`, `ndbd`, `ndbmtbd`, and `mysqld`.
- **Host.** The hostname or IP address of the computer where the process is running.
- **Status.** The state or condition of this process. Possible values for this column are given later in this section.
- **Nodegroup.** If the `Process` is `ndbd` or `ndbmtbd`—that is, if the process is a data node process—then this column shows the ID of the node group to which the process belongs. For any other value of `Process`, this column is empty.

Possible values for the `Status` column are shown in the following table, together with a description of what this value represents:

Table 4.3 Status values shown by `show status --process`

Status Value	Meaning
<code>running</code>	The process is running normally.
<code>stopped</code>	The process has been stopped by the user.
<code>added</code>	The process has been added to the cluster, but not yet started.
<code>connected</code>	The <code>ndbapi</code> or <code>mysqld</code> process is connected to the cluster.
<code>starting</code>	The process has been started, but is not yet fully running. (For data nodes, you can determine which start phase the node is currently in by using the <code>status</code> command in the <code>ndb_mgm</code> client.)
<code>stopping</code>	The process has received a command to stop, and is now shutting down.
<code>failed</code>	The process has shut down unexpectedly (likely to have crashed). You should determine the cause for this unplanned shutdown, fix the problem, and restart the process as soon as possible.
<code>import</code>	The process is part of a cluster that was created for import, but the actual migration of processes and data from the original cluster has not yet taken place. <code>start process</code> and <code>stop</code>

Status Value	Meaning
	<code>process</code> commands fail for this process until this migration has occurred.
unknown	MySQL Cluster Manager is unable to establish the current status of this process. You should try to determine its status using other means.

`--progress` option

For MySQL Cluster Manager 1.4.2 and later: When run with this option, `show status` returns, when available, progress on the current action of `mcmd` on the cluster named `cluster_name`, in terms of the percentage of the total number of steps completed:

```
mcm> show status --progress mycluster;
+-----+-----+-----+
| Command      | Status   | Progress |
+-----+-----+-----+
| restore cluster | executing | 47%      |
+-----+-----+-----+
1 row in set (0.02 sec)
```

`--progressbar` option

For MySQL Cluster Manager 1.4.2 and later: The option provides the same function as the `--progress` option, but also adds an ASCII-art progress bar:

```
mcm> show status --progressbar mycluster;
+-----+-----+-----+
| Command      | Status   | Progress |
+-----+-----+-----+
| restore cluster | executing | 47% [#####          ] |
+-----+-----+-----+
1 row in set (0.02 sec)
```

You must supply the name of an existing cluster with this command, or else `show status` fails with an error, as shown here:

```
mcm> show status;
ERROR 6 (00MGR): Illegal number of operands

mcm> show status -c nosuchcluster;
ERROR 5001 (00MGR): Cluster nosuchcluster not defined
```



Important

Do not confuse this command with the MySQL `SHOW STATUS` statement, which has a different syntax and can be used only in the standard `mysql` client. The MySQL Cluster Manager client command accepts only those options shown at the beginning of this section, and does not accept a `LIKE` or `WHERE` clause.

4.4.7 The `start cluster` Command

```
start cluster [--initial|-i] [--skip-init=process_id_list] cluster_name
process_id_list:
    process_id[, process_id[, ...]]
```

This command starts the cluster named `cluster_name`, as shown in this example:

```
mcm> start cluster mycluster;
+-----+-----+
| Command result          |
+-----+-----+
| Cluster started successfully |
+-----+-----+
1 row in set (45.37 sec)
```

In order for the command to succeed, the cluster named in the command must already exist; otherwise the command fails with the error `Cluster cluster_name not defined`, as shown here:

```
mcm> list sites;
+-----+-----+-----+-----+
| Site   | Port | Local | Hosts                               |
+-----+-----+-----+-----+
| mysite | 1862 | Local | tonfisk,flundra,grindval,haj      |
+-----+-----+-----+-----+
1 row in set (1.72 sec)

mcm> list clusters mysite;
+-----+-----+
| Cluster | Package |
+-----+-----+
| mycluster | mypackage |
+-----+-----+
1 row in set (1.70 sec)

mcm> start cluster yourcluster;
ERROR 5001 (00MGR): Cluster yourcluster not defined
```

In addition, the cluster must not already be running, as shown here:

```
mcm> show status --cluster mycluster;
+-----+-----+-----+
| Cluster | Status           | Comment |
+-----+-----+-----+
| mycluster | fully operational |         |
+-----+-----+-----+
1 row in set (0.01 sec)

mcm> start cluster mycluster;
ERROR 5005 (00MGR): Cluster mycluster is running
```

A cluster created for import cannot be started until the import has been completed. See [Section 4.4.1, “The create cluster Command”](#), and [Section 3.5, “Importing MySQL NDB Clusters into MySQL Cluster Manager”](#), for more information.

`--initial` option

The `--initial` option (short form: `-i`) causes the following to happen:

- All cluster data nodes are started as if `start process --initial` had been used on them, which means that all data nodes wipe their data and start with clean data node file systems. NDB tables that were previously stored in the cluster are lost.
- *For MySQL Cluster Manager 1.4.3 and later:* All cluster SQL nodes are started as if `start process --initial` had been used on them, which means MySQL Cluster Manager rebuilds the `mysqld` data directory with the `mysqld --initialize-insecure` command for MySQL 5.7 and with the `mysql_install_db` command for MySQL 5.6. However, the node's data directory must be empty, or the reinitialization will not be attempted.

To skip reinitialization for any SQL nodes, list their process IDs (separated by commas if there are more than one) using the `--skip-init=process_id_list` option, for example:

```
mcm> start cluster --initial --skip-init=50,51 mycluster;
```

The `--skip-init` option only accepts SQL node IDs as its argument; it cannot be used to skip the initialization of data nodes.

Under normal circumstances, you should use this option to start a cluster only when either you do not wish to preserve any of its data (and want to make a clean start), or you intend to restore the cluster from backup to a known good state (see [Section 4.7.4, “The restore cluster Command”](#)). You should also be aware that no special warnings are printed by the `mcm` client when `--initial` is used with `start cluster`; the command is immediately executed.

For information about creating cluster backups, see [Section 4.7.2, “The `backup cluster` Command”](#). If you need to know which backups are available (if any), use `list backups`.

4.4.8 The `stop cluster` Command

```
stop cluster cluster_name
```

This command stops the cluster named `cluster_name`, if it is running, as shown in this example:

```
mcm> stop cluster mycluster;
+-----+
| Command result |
+-----+
| Cluster stopped successfully |
+-----+
1 row in set (21.31 sec)
```

`stop cluster` fails if the cluster is not in an operational state (see [Section 4.4.6, “The `show status` Command”](#)), for information about obtaining the cluster's status):

```
mcm> show status --cluster mycluster;
+-----+-----+-----+
| Cluster | Status | Comment |
+-----+-----+-----+
| mycluster | stopped | |
+-----+-----+-----+
1 row in set (0.01 sec)

mcm> stop cluster mycluster;
ERROR 5006 (00MGR): Cluster mycluster is stopped
```

`stop cluster` cannot be used on a cluster created for import until the import has been completed. See [Section 4.4.1, “The `create cluster` Command”](#), and [Section 3.5, “Importing MySQL NDB Clusters into MySQL Cluster Manager”](#), for more information.

4.4.9 The `autotune` Command

```
autotune [--dryrun] [--sequential-restart] [--writeload=writeload] template cluster_name

writeload:
    {low|medium|high}

template:
    {web|realtime|test}
```

The command `autotune` a number of parameters for the cluster according to the specified values for the `template` argument and the [optional] `writeload` option, in order to optimize the cluster's performance.

The valid values for `template` are:

- `web`: Maximize performance for the given hardware.
- `realtime`: Maximize performance while maximizing sensitivity to timeouts in order to minimize the time needed to detect failed cluster processes.
- `test`: Minimal resource usage for small-scale testing. Not intended for production environments.

The valid values for `--writeload` are:

- `low`: The expected load includes fewer than 100 write transactions for second.
- `medium`: The expected load includes 100 to 1000 write transactions per second. This is the default value used when `--writeload` is not specified.
- `high`: The expected load includes more than 1000 write transactions per second.

The cluster must be in the `created` or `fully operational` status for this command to work, or an error will result. The command tunes the cluster by issuing a number of `set` commands to adjust different parameters, and then performs a rolling restart for the cluster. For *MySQL Cluster Manager 1.4.8 and later*, use the `--sequential-restart` option to make the rolling restart a `sequential` one.

When the `--dryrun` option is used, the command does not make any actual changes to the cluster, but writes the `set` commands that it will issue for tuning into the file `/path-to-mcm-data-repository/clusters/clustername/tmp/autotune.message_id.mcm`.

```
mcm> autotune --dryrun --writeload=high realtime mycluster;
+-----+
| Command result |
+-----+
| Autotuning calculation complete. Please check /opt/mcm_data/clusters/mycluster/tmp/autotune.30fcce24_2184_0.mcm |
+-----+
1 row in set (0.62 sec)
```

```
$> cat /opt/mcm_data/clusters/mycluster/tmp/autotune.30fcce24_2184_0.mcm
# The following will be applied to the current cluster config:
set HeartbeatIntervalDbDb:ndbmttd=1500 mycluster;
set HeartbeatIntervalDbApi:ndbmttd=1500 mycluster;
set RedoBuffer:ndbmttd=64M mycluster;
set SharedGlobalMemory:ndbmttd=20M mycluster;
set DataMemory:ndbmttd=83886080 mycluster;
set IndexMemory:ndbmttd=18874368 mycluster;
set MaxNoOfExecutionThreads:ndbmttd=2 mycluster;
set FragmentLogFileSize:ndbmttd=256M mycluster;
set NoOfFragmentLogFiles:ndbmttd=3 mycluster;
```

After checking out those changes in the `.mcm` file, if you do not want to apply all of them to your cluster, you can edit the `.mcm` file as desired, and then execute it at the `mcm` client (see [Section 3.5.2.3, “Creating and Configuring the Target Cluster”](#) for how to do that). If you are happy with all the changes described in the file, issue the `autotune` command again without the `--dryrun` option, to perform the tuning:

```
mcm> autotune --writeload=high realtime mycluster;
+-----+
| Command result |
+-----+
| Cluster successfully autotuned to template realtime |
+-----+
1 row in set (2 min 58.09 sec)
```



Note

The command is supported only for MySQL NDB Cluster 7.4 and later.

4.4.10 The `upgrade cluster` Command

```
upgrade cluster [--package=|-P ]package_name
  [{--nodeid|-n }node_id_list] [--force|-f]
  [--retry|-L] [--set=attribute_assignment_list] cluster_name

node_id_list:
  node_id[, node_id[, ...]]

attribute_assignment_list:
  attribute_assignment[,attribute_assignment][,...]

attribute_assignment:
  attribute_name:process_name[=value]
```

This command upgrades the cluster named `cluster_name` to the software package `package_name` specified with the `--package`. It finishes an upgrade by performing a rolling restart for the cluster, in which data nodes are restarted with the `--initial` option to have their data file systems rebuilt.

The new package must be registered using `add package` before you can use it for an upgrade; otherwise, `upgrade cluster` fails with an error.

To use the command to perform an upgrade on a cluster, unless some special options are used (see explanations on the `--force`, `--retry`, and `--nodeid` options below), the cluster must be in the `fully operational` status (you can check that using the command `show status --cluster cluster_name`). A cluster created for import cannot be upgraded until the import has been completed. See [Section 4.4.1, “The `create cluster` Command”](#), and [Section 3.5, “Importing MySQL NDB Clusters into MySQL Cluster Manager”](#), for more information.

Suppose `mycluster` is using MySQL NDB Cluster 7.4.8, and the binaries are registered with a package named `7.4.8`, as shown by this `list clusters` command:

```
mcm> list clusters mysite;
+-----+-----+
| Cluster | Package |
+-----+-----+
| mycluster | 7.4.8 |
+-----+-----+
1 row in set (1.80 sec)
```

Now you wish to upgrade `mycluster` to MySQL NDB Cluster. Assuming that you have placed the NDB 7.6.13 binaries in the same directory on each host, the `add package` command to create a new package named `7.6.13` that contains these binaries might look something like this:

```
mcm> add package --basedir=/usr/local/ndb-7.6.13 7.6.13;
+-----+-----+
| Command result |
+-----+-----+
| Package added successfully |
+-----+-----+
1 row in set (0.88 sec)
```



Note

On Windows, you must replace any backslash (`\`) characters in the path used for the `add package` command's `--basedir` option with forward slashes (`/`). See [Section 4.3.1, “The `add package` Command”](#), for additional information and examples.

Both packages should now be listed in the output of the command `list packages mysite`. To perform the upgrade to the `7.6.13` package, use the `upgrade cluster` command as shown here:

```
mcm> upgrade cluster --package=7.6.13 mycluster;
+-----+-----+
| Command result |
+-----+-----+
| Cluster upgraded successfully |
+-----+-----+
1 row in set (3 min 17.00 sec)
```

Once the `upgrade cluster` command has been successfully executed, you can verify that `mycluster` is now using the `7.6.13` package from the output of the appropriate `list clusters` command:

```
mcm> list clusters mysite;
+-----+-----+
| Cluster | Package |
+-----+-----+
| mycluster | 7.6.13 |
+-----+-----+
1 row in set (1.80 sec)
```

The command can perform major as well as minor series upgrades. Despite the name of this command, `upgrade cluster` can also be used to perform MySQL NDB Cluster downgrades.

Not all upgrades and downgrades between different versions of MySQL NDB Cluster are supported by the command. The following criteria must also be met:

- The upgrade or downgrade must be supported by the MySQL NDB Cluster versions involved. See [Upgrading and Downgrading NDB Cluster](#).
- Both the versions you upgrade or downgrade to and from must be supported by the version of MySQL Cluster Manager you are using (for example, upgrade from MySQL NDB Cluster 6.3 to 7.5 must be performed manually, because MySQL Cluster Manager no longer supports MySQL NDB Cluster 6.3).

When using the `upgrade cluster` command, you can use the `--set` option to reconfigure your MySQL NDB Cluster at the same time. This is particularly helpful when the upgrade requires configuration changes to your cluster. This option takes as its argument an attribute assignment list similar in format to that used with the `get` and `set` commands; see description of the `set` command on the proper way to formulate an attribute assignment list. For example: if you want to change the memory assigned to each data node for storing database records to 750M, specify that with the `--set` option in your `upgrade cluster` command:

```
mcm> upgrade cluster --package=7.6.13 --set=DataMemory:ndbd=750M mycluster;
+-----+
| Command result |
+-----+
| Cluster upgraded successfully |
+-----+
1 row in set (3 min 17.04 sec)
```



Note

Unlike the way you use the `set` command, an equal sign (=) immediately following the `--set` option is required.

Options for dealing with failed upgrades

The `--force` option (`-f` for short) should be used when you want to run the `upgrade cluster` command again after a failed upgrade attempt that ends up with any failed management or data nodes. Without the `--force` option, the `upgrade cluster` command only runs when the cluster is in the `fully operational` status.

The `--retry` option (`-L` for short) should be used when you want to retry the `upgrade cluster` command after a failed attempt that ends up with some nodes being upgraded, and some not. Without the `--retry` option, the `upgrade cluster` command cannot be run on the same cluster twice using the same package.

In the case of a failed or incomplete upgrade, instead of using the `--force` and `--retry` option, you can also choose to retry the upgrade only on the failed nodes by specifying them using the `--nodeid` option (`-n` in short). Check for any failed nodes after a failed upgrade:

```
mcm> upgrade cluster -P next mycluster;
ERROR 7006 (OOMGR): Process error: <reason of failure>
mcm> show status --process mycluster;
+-----+-----+-----+-----+-----+-----+
| NodeId | Process | Host | Status | Nodegroup | Package |
+-----+-----+-----+-----+-----+-----+
| 49 | ndb_mgmd | thinkpad | running | | next |
| 1 | ndbmttd | thinkpad | running | 0 | next |
| 2 | ndbmttd | thinkpad | running | 0 | next |
| 50 | mysqld | thinkpad | running | | next |
| 51 | mysqld | thinkpad | failed | | next |
| 52 | ndbapi | * | added | | |
+-----+-----+-----+-----+-----+-----+
6 rows in set (0.03 sec)
```

Then, issue the command again, specifying the failed node with the `--nodeid` option:

```
mcm> upgrade cluster --nodeid=51 -P next mycluster;
+-----+
| Command result |
+-----+
| Cluster upgraded successfully |
+-----+
1 row in set (26.03 sec)
```



Warning

Using the `--nodeid` option inappropriately with the `upgrade cluster` command may result in a partial upgrade. Use it only when a previous attempt to upgrade failed, and only with guidance from the proper support personnel.

4.5 MySQL Cluster Manager Configuration Commands

This section covers the commands used in the MySQL Cluster Manager for getting and setting values of various types used in MySQL NDB Cluster configuration. We begin with a discussion of what we mean by the term “configuration attribute”, and how this relates to the manual configuration of a MySQL NDB Cluster using MySQL NDB Cluster configuration parameters and MySQL Server options and variables that you may already be familiar with.

Configuration attributes.

Traditionally, when administering MySQL NDB Cluster, it has been necessary to distinguish between 3 types of configuration data:

- *Configuration parameters* set in the MySQL NDB Cluster global configuration file read by the management server (or servers), by convention named `config.ini`
- *Configuration variables* set in a running MySQL server (SQL node) by using the SQL `SET` statement in the `mysql` command-line client (or in another MySQL client application)
- *Configuration options* passed to MySQL NDB Cluster executable programs when invoking them



Note

Configuration options passed to `mysqld` often have the effect of setting values for configuration variables, many—but not all—of which can be overridden in a running MySQL server using a MySQL client application such as `mysql`.

MySQL Cluster Manager simplifies this configuration scheme by treating all 3 types of configuration data as *attributes*, where the term “attribute” refers to a MySQL NDB Cluster configuration parameter, a MySQL Server variable, or a command-line option used with one or more MySQL NDB Cluster binary programs. It does this transparently, handling all necessary changes in a unified interface.

Suppose that you wish to know how much data memory is allocated to the data nodes in a given MySQL NDB Cluster. Rather than having to determine that this is controlled using the `DataMemory` configuration parameter that is written in the `config.ini` file and then reading that file to find the value, you merely invoke the MySQL Cluster Manager `get` command, and MySQL Cluster Manager handles reading from the file for you, and displays the value without the necessity of opening the file in a separate application such as `more` or `less`. If you wish to change the amount of data memory allocated to the data nodes, you can issue a MySQL Cluster Manager `set` (or `reset`) command; MySQL Cluster Manager then writes the desired value to `config.ini`. If—as is the case with `DataMemory`—updating a configuration value in a running MySQL NDB Cluster requires a rolling restart to be performed, MySQL Cluster Manager can perform this operation automatically so that the configuration change takes effect without further intervention required on the part of the operator.

Configuration attribute levels.

A configuration attribute value applies at one of the three levels, described here:

- *Default*: This value is always used by any MySQL NDB Cluster process of the type or types (such as `ndbd` or `mysqld`) to which the attribute applies, unless this value is overridden by the user.
- *Process*: This value is used for all instances of a given type of MySQL NDB Cluster process.
- *Instance*: This value is used for a specific instance of a MySQL NDB Cluster process, the instance being identified by its MySQL NDB Cluster node ID.

Default values are hard-coded into MySQL NDB Cluster; you can override a default value for a given configuration attribute (using the `set` command) or reset a given attribute value to its default (using the `reset` command), but you cannot change a default value itself. You can set or reset an configuration attribute's value on either the process level or the instance level using a single `set` or `reset` command. Once you have set or reset the value of a configuration attribute, this value persists until it is changed by executing another `set` or `reset` command.



Note

When setting or resetting a configuration attribute value, you must specify the level at which the setting applies.

MySQL Cluster Manager determines what value to use for a configuration attribute relating to a given process by following these steps for each MySQL NDB Cluster process:

(For each configuration attribute:)

1. Is an attribute value defined for the node ID of this process?
 - Yes: Use the value that was defined for this node ID, and exit.
 - No: Proceed to the next step.
2. Is an attribute value specified on the process level, that is, for all processes of this type?
 - Yes: Use the value that was specified for all processes of this type, and exit.
 - No: Use the default value that applies to processes of this type, and exit.

(In the steps just shown, “exit” can be taken to mean “If there are more configuration attributes applicable to this process that have not yet been set, proceed to the next attribute until there are no more attributes to be set for this process”.)



Note

The most recently specified value takes precedence. This means that if you set a configuration attribute for a specific process, then later specify a process-level value for this attribute, the process-level value is used for all processes of that type, including the instance for which you earlier set an instance-specific value.

Mandatory attributes.

Some attributes must be defined in the MySQL Cluster Manager at the process type or instance level for all processes of the applicable type or types for the cluster configuration to be valid. Such *mandatory attributes* may be changed, but not reset; in other words, the definition can be changed, but the definition itself cannot be removed entirely. Another way of stating this is that a mandatory attribute has no default value.

An example of a mandatory attribute is `NodeId`. If you try to reset a mandatory attribute, the attempt fails with an error, as shown here:

```
mcm> reset NodeId:ndb_mgmd:1 mycluster;
ERROR 6007 (00MGR): Config attribute NodeId is mandatory and cannot be reset
mcm> reset NodeId:ndbd:2 mycluster;
ERROR 6007 (00MGR): Config attribute NodeId is mandatory and cannot be reset
mcm> reset NodeId:mysqld:4 mycluster;
```

```
ERROR 6007 (OOMGR): Config attribute NodeId is mandatory and cannot be reset
```

Read-only attributes.

A *read-only attribute* is an attribute that must be defined by the MySQL Cluster Manager when a cluster is created. A read-only attribute can be neither changed nor reset by the user. This means that a read-only attribute is always a mandatory attribute.

One such attribute is `HostName`, which is read only for any type of MySQL NDB Cluster process. Any attempt to change or reset a read-only attribute fails, as shown here:

```
mcm> reset HostName:ndb_mgmd mycluster;
ERROR 6008 (OOMGR): Config attribute HostName is readonly and cannot be changed
mcm> reset HostName:ndbd mycluster;
ERROR 6008 (OOMGR): Config attribute HostName is readonly and cannot be changed
mcm> reset HostName:mysqlqld mycluster;
ERROR 6008 (OOMGR): Config attribute HostName is readonly and cannot be changed

mcm> set HostName:ndb_mgmd mycluster;
ERROR 6008 (OOMGR): Config attribute HostName is readonly and cannot be changed
mcm> set HostName:ndbd mycluster;
ERROR 6008 (OOMGR): Config attribute HostName is readonly and cannot be changed
mcm> set HostName:mysqlqld mycluster;
ERROR 6008 (OOMGR): Config attribute HostName is readonly and cannot be changed
```

An attribute that is mandatory or read only is set when a cluster is created. Neither a mandatory attribute nor a read-only attribute can be reset. (Neither type of attribute has a default value other than what is set for it when the cluster is created.) A mandatory attribute can be changed at any time by the user; a read-only attribute cannot be changed once the cluster has been created. You can obtain a listing of mandatory and read-only attributes using the `get` command.

A listing of attribute properties also can be found in the output of `ndb_config --configinfo --xml` (see [ndb_config — Extract NDB Cluster Configuration Information](#)); for more complete information, see [Configuration of NDB Cluster](#).

MySQL Cluster Manager determines internally which attributes are considered read-only for reasons of cluster stability and performance. You can use the `get` command to see which attributes are read only.

Command-line-only attributes.

Command-line-only attributes are attributes that, when outside of MySQL Cluster Manager, must be specified as command-line options instead of parameters in a configuration file (for example, `config.ini` or `my.cnf`). These include all the command-line options of the `ndb_mgmd`, `ndbd`, and `ndbmttd` nodes, as well as `mysqlqld` options listed in [Server Option, System Variable, and Status Variable Reference](#) as not valid in option files. A small number of these command-line-only attributes can, however, be configured with MySQL Cluster Manager using the `set` and `reset` commands, and their values can be checked with the `get` command; they include:

- For `ndb_mgmd`: `--core-file`, `--log-name`, `--verbose`
- For `ndbd` and `ndbmttd`: `--core-file`, `--verbose`
- The `mysqlqld --core-file` option.

These command-line-only attributes supported by the `get`, `set`, and `reset` commands are marked with `Command Line` in the `Comment` column of the `get` command's output:

```
mcm> set core-file:ndb_mgmd mycluster;
+-----+
| Command result |
+-----+
| Cluster reconfigured successfully |
+-----+
1 row in set (32.43 sec)

mcm> get -d core-file:ndb_mgmd mycluster;
+-----+-----+-----+-----+-----+-----+-----+-----+
| Name | Value | Process1 | NodeId1 | Process2 | NodeId2 | Level | Comment |
+-----+-----+-----+-----+-----+-----+-----+-----+
```

```

+-----+-----+-----+-----+-----+-----+-----+-----+
| core-file |      | ndb_mgmd | 49      |      |      | Process | Command Line |
+-----+-----+-----+-----+-----+-----+-----+-----+
1 row in set (0.07 sec)

mcm> reset core-file:ndb_mgmd mycluster;
+-----+-----+
| Command result |
+-----+-----+
| Cluster reconfigured successfully |
+-----+-----+
1 row in set (9.57 sec)

mcm> get -d core-file:ndb_mgmd mycluster;
+-----+-----+-----+-----+-----+-----+-----+-----+
| Name      | Value | Process1 | NodeId1 | Process2 | NodeId2 | Level  | Comment      |
+-----+-----+-----+-----+-----+-----+-----+-----+
| core-file | false | ndb_mgmd | 49      |      |      | Default | Command Line |
+-----+-----+-----+-----+-----+-----+-----+-----+
1 row in set (0.05 sec)

```

4.5.1 The `get` Command

```

get [--include-defaults|-d] [filter_specification_list] cluster_name

filter_specification_list:
    filter_specification[,filter_specification][,...]

filter_specification:
    [attribute_name][:process_specification][+process_specification]

process_specification:
    [process_name][:process_id]

process_name:
    {ndb_mgmd|ndbd|ndbmt|mysqld|ndbapi}

```

This command is used in the MySQL Cluster Manager client to obtain configuration attribute values from a MySQL NDB Cluster. (See [Section 4.5, “MySQL Cluster Manager Configuration Commands”](#), for a definition of the term “attribute” as it applies in the MySQL Cluster Manager.) The output includes the following columns:

- **Name:** This column contains the name of the configuration attribute.
- **Value:** This column shows the attribute's current value.
- **Process1:** This column holds the process type to which the attribute applies. This is one of `ndb_mgmd`, `ndbd`, `ndbmt` (MySQL NDB Cluster 7.0 and later), or `mysqld`.
- **Id1:** This is the process ID of the process to which the attribute applies.
- **Process2:** For attributes that require specifying two nodes, such as those relating to TCP/IP connections, this column shows the process type of the second node.
- **Id2:** For attributes that require specifying two nodes, this column shows the process ID for the second node.
- **Level:** This is the attribute process level. This value in this column can be `Default`, `Process`, or empty; if this column is empty, it means that the attribute applies on the instance level.
- **Comment:** This column is used to show whether the attribute is `Mandatory`, `Read only`, `Default` attribute, or user defined (in which case the `Comment` column is empty).

By default, `get` returns only those attributes that have been set explicitly, either by the MySQL Cluster Manager itself, or by the user. In other words, it shows only attributes that are mandatory (including read-only attributes), or that have been set by the user after the cluster was created. Hereafter in this discussion, we refer to these as “non-default attributes”.

Thus, prior to setting any configuration attributes, you can obtain a list of all mandatory and read-only attributes by running the simplest possible form of this command, as shown here:

```
mcm> get mycluster\G
***** 1. row *****
  Name: Name
  Value: mycluster
Process1:
  NodeId1:
Process2:
  NodeId2:
  Level:
  Comment: Read only
***** 2. row *****
  Name: DataDir
  Value: /opt/mcm_data/clusters/mycluster/49/data
Process1: ndb_mgmd
  NodeId1: 49
Process2:
  NodeId2:
  Level:
  Comment:
***** 3. row *****
  Name: HostName
  Value: torsk
Process1: ndb_mgmd
  NodeId1: 49
Process2:
  NodeId2:
  Level:
  Comment: Read only
***** 4. row *****
  Name: NodeId
  Value: 49
Process1: ndb_mgmd
  NodeId1: 49
Process2:
  NodeId2:
  Level:
  Comment: Read only
***** 5. row *****
  Name: PortNumber
  Value: 1186
Process1: ndb_mgmd
  NodeId1: 49
Process2:
  NodeId2:
  Level: Process
  Comment:
***** 6. row *****
  Name: DataDir
  Value: /opt/mcm_data/clusters/mycluster/1/data
Process1: ndbmttd
  NodeId1: 1
Process2:
  NodeId2:
  Level:
  Comment:
***** 7. row *****
  Name: HostName
  Value: torsk
Process1: ndbmttd
  NodeId1: 1
Process2:
  NodeId2:
  Level:
  Comment: Read only
***** 8. row *****
  Name: NodeId
  Value: 1
Process1: ndbmttd
  NodeId1: 1
```

```

Process2:
NodeId2:
Level:
Comment: Read only
***** 9. row *****
Name: DataDir
Value: /opt/mcm_data/clusters/mycluster/2/data
Process1: ndbmttd
NodeId1: 2
Process2:
NodeId2:
Level:
Comment:
***** 10. row *****
Name: HostName
Value: torsk
Process1: ndbmttd
NodeId1: 2
Process2:
NodeId2:
Level:
Comment: Read only
***** 11. row *****
Name: NodeId
Value: 2
Process1: ndbmttd
NodeId1: 2
Process2:
NodeId2:
Level:
Comment: Read only
***** 12. row *****
Name: datadir
Value: /opt/mcm_data/clusters/mycluster/50/data
Process1: mysqld
NodeId1: 50
Process2:
NodeId2:
Level:
Comment:
***** 13. row *****
Name: default_storage_engine
Value: ndbcluster
Process1: mysqld
NodeId1: 50
Process2:
NodeId2:
Level: Process
Comment:
***** 14. row *****
Name: HostName
Value: torsk
Process1: mysqld
NodeId1: 50
Process2:
NodeId2:
Level:
Comment: Read only
***** 15. row *****
Name: ndb_nodeid
Value: 50
Process1: mysqld
NodeId1: 50
Process2:
NodeId2:
Level:
Comment: Read only
***** 16. row *****
Name: ndbcluster
Value: on
Process1: mysqld
NodeId1: 50

```

```

Process2:
  NodeId2:
  Level:
  Comment: Read only
***** 17. row *****
  Name: NodeId
  Value: 50
Process1: mysqld
  NodeId1: 50
Process2:
  NodeId2:
  Level:
  Comment: Read only
***** 18. row *****
  Name: port
  Value: 3306
Process1: mysqld
  NodeId1: 50
Process2:
  NodeId2:
  Level:
  Comment:
***** 19. row *****
  Name: socket
  Value: /tmp/mysql.mycluster.50.sock
Process1: mysqld
  NodeId1: 50
Process2:
  NodeId2:
  Level:
  Comment:
***** 20. row *****
  Name: tmpdir
  Value: /opt/mcm_data/clusters/mycluster/50/tmp
Process1: mysqld
  NodeId1: 50
Process2:
  NodeId2:
  Level:
  Comment:
***** 21. row *****
  Name: datadir
  Value: /opt/mcm_data/clusters/mycluster/51/data
Process1: mysqld
  NodeId1: 51
Process2:
  NodeId2:
  Level:
  Comment:
***** 22. row *****
  Name: default_storage_engine
  Value: ndbcluster
Process1: mysqld
  NodeId1: 51
Process2:
  NodeId2:
  Level: Process
  Comment:
***** 23. row *****
  Name: HostName
  Value: torsk
Process1: mysqld
  NodeId1: 51
Process2:
  NodeId2:
  Level:
  Comment: Read only
***** 24. row *****
  Name: ndb_nodeid
  Value: 51
Process1: mysqld
  NodeId1: 51

```

```

Process2:
NodeId2:
Level:
Comment: Read only
***** 25. row *****
Name: ndbcluster
Value: on
Process1: mysqld
NodeId1: 51
Process2:
NodeId2:
Level:
Comment: Read only
***** 26. row *****
Name: NodeId
Value: 51
Process1: mysqld
NodeId1: 51
Process2:
NodeId2:
Level:
Comment: Read only
***** 27. row *****
Name: port
Value: 3307
Process1: mysqld
NodeId1: 51
Process2:
NodeId2:
Level:
Comment:
***** 28. row *****
Name: socket
Value: /tmp/mysql.mycluster.51.sock
Process1: mysqld
NodeId1: 51
Process2:
NodeId2:
Level:
Comment:
***** 29. row *****
Name: tmpdir
Value: /opt/mcm_data/clusters/mycluster/51/tmp
Process1: mysqld
NodeId1: 51
Process2:
NodeId2:
Level:
Comment:
***** 30. row *****
Name: NodeId
Value: 52
Process1: ndbapi
NodeId1: 52
Process2:
NodeId2:
Level:
Comment: Read only
30 rows in set (0.07 sec)

```

On Windows, no substitutions for backslashes or other characters used in values of paths reported by the `get` command is performed. However, it is possible to see forward slashes used in such paths if the values were set using the `set` command. See [Setting Attributes Containing Paths on Windows \[129\]](#), for more information.

Although a `socket` attribute is shown for `mysqld` nodes in the `get` output from the previous example and is not marked `Read only`, MySQL Cluster Manager does not support socket files on Windows. For this reason; you should not attempt to set `socket` attributes for Windows `mysqld` processes using MySQL Cluster Manager.

To include default values for attributes that have not (or not yet) been set explicitly, you can invoke this command with the `--include-defaults` option (short form: `-d`), as shown here (in part):

```
mcm> get --include-defaults mycluster\G
***** 1. row *****
  Name: Name
  Value: mycluster
Process1:
  NodeId1:
Process2:
  NodeId2:
  Level:
  Comment: Read only
***** 2. row *****
  Name: Checksum
  Value: false
Process1: ndb_mgmd
  NodeId1: 49
Process2: ndbmttd
  NodeId2: 1
  Level: Default
  Comment:
***** 3. row *****
  Name: Group
  Value: 55
Process1: ndb_mgmd
  NodeId1: 49
Process2: ndbmttd
  NodeId2: 1
  Level: Default
  Comment:
***** 4. row *****
  Name: HostName1
  Value: NULL
Process1: ndb_mgmd
  NodeId1: 49
Process2: ndbmttd
  NodeId2: 1
  Level: Default
  Comment:
***** 5. row *****
  Name: HostName2
  Value: NULL
Process1: ndb_mgmd
  NodeId1: 49
Process2: ndbmttd
  NodeId2: 1
  Level: Default
  Comment:
***** 6. row *****
  Name: NodeId1
  Value: NULL
Process1: ndb_mgmd
  NodeId1: 49
Process2: ndbmttd
  NodeId2: 1
  Level: Default
  Comment: Mandatory
***** 7. row *****
  Name: NodeId2
  Value: NULL
Process1: ndb_mgmd
  NodeId1: 49
Process2: ndbmttd
  NodeId2: 1
  Level: Default
  Comment: Mandatory
***** 8. row *****
  Name: NodeIdServer
  Value: NULL
Process1: ndb_mgmd
  NodeId1: 49
```

```

Process2: ndbmttd
NodeId2: 1
Level: Default
Comment: Mandatory
***** 9. row *****
Name: OverloadLimit
Value: 0
Process1: ndb_mgmd
NodeId1: 49
Process2: ndbmttd
NodeId2: 1
Level: Default
Comment:
***** 10. row *****
Name: Proxy
Value: NULL
Process1: ndb_mgmd
NodeId1: 49
Process2: ndbmttd
NodeId2: 1
Level: Default
Comment:
***** 11. row *****
Name: ReceiveBufferMemory
Value: 2097152
Process1: ndb_mgmd
NodeId1: 49
Process2: ndbmttd
NodeId2: 1
Level: Default
Comment:
***** 12. row *****
Name: SendBufferMemory
Value: 2097152
Process1: ndb_mgmd
NodeId1: 49
Process2: ndbmttd
NodeId2: 1
Level: Default
Comment:
***** 13. row *****
Name: SendSignalId
Value: true
Process1: ndb_mgmd
NodeId1: 49
Process2: ndbmttd
NodeId2: 1
Level: Default
Comment:
***** 14. row *****
Name: TCP_MAXSEG_SIZE
Value: 0
Process1: ndb_mgmd
NodeId1: 49
Process2: ndbmttd
NodeId2: 1
Level: Default
Comment:
...
***** 1901. row *****
Name: StartConnectBackoffMaxTime
Value: 0
Process1: ndbapi
NodeId1: 52
Process2:
NodeId2:
Level: Default
Comment:
***** 1902. row *****
Name: TotalSendBufferMemory

```

```

Value: 0
Process1: ndbapi
NodeId1: 52
Process2:
NodeId2:
Level: Default
Comment:
***** 1903. row *****
Name: wan
Value: false
Process1: ndbapi
NodeId1: 52
Process2:
NodeId2:
Level: Default
Comment:
1903 rows in set (0.11 sec)

```

As you can see, the output from this `get` command is quite long (and the number of rows generated increases with the number of nodes in the cluster.) However, it is possible to filter the output so that you can view only the attribute or attributes in which you are interested. This can be done by using a comma-separated list of one or more filter specifications. A filter specification is defined as shown here (condensed from that given at the beginning of this section, but effectively the same):

```
[attribute_name][:[process_name][:process_id]]
```

Filtering can be applied per attribute, per process type, and per process instance. We now provide some examples illustrating the use of such filters.

To obtain the value of a given attribute for all processes to which it applies in the cluster, you need only use the name of the attribute as a filter. For example, to obtain the `HostName` of all processes in the cluster named `mycluster`, you can execute the command shown here:

```
mcm> get HostName mycluster;
```

Name	Value	Process1	NodeId1	Process2	NodeId2	Level	Comment
HostName	flundra	ndbd	1				Read only
HostName	tonfisk	ndbd	2				Read only
HostName	grindval	ndb_mgmd	49				Read only
HostName	haj	mysqld	50				Read only
HostName	torsk	mysqld	51				Read only

```

5 rows in set (0.04 sec)

```

The wildcard `*` (asterisk character) can be used to match a single or multiple attribute names; for example:

```
mcm> get Host* mycluster;
```

Name	Value	Process1	NodeId1	Process2	NodeId2	Level	Comment
HostName	flundra	ndbd	1				Read only
HostName	tonfisk	ndbd	2				Read only
HostName	grindval	ndb_mgmd	49				Read only
HostName	haj	mysqld	50				Read only
HostName	torsk	mysqld	51				Read only

```

5 rows in set (0.04 sec)

```

```
mcm> get H* yourcluster;
```

Name	Value	Process1	NodeId1	Process2	NodeId2	Level	Comment
HostName	tonfisk	ndb_mgmd	49				Read only
HostName	flundra	ndb_mgmd	53				Read only
HeartbeatIntervalDbApi	1500	ndbmttd	1			Process	
HeartbeatIntervalDbDb	1500	ndbmttd	1			Process	
HostName	tonfisk	ndbmttd	1				Read only

The `get` Command

```
| Name      | Value      | Process1 | Id1 | Process2 | Id2 | Level | Comment |
+-----+-----+-----+-----+-----+-----+-----+-----+
| DataDir   | /opt/c2data | ndbd     | 1   |           |     |       |         |
| HostName  | tonfisk     | ndbd     | 1   |           |     |       | Read only |
| DataDir   | /opt/c3data | ndbd     | 2   |           |     |       |         |
| HostName  | flundra     | ndbd     | 2   |           |     |       | Read only |
+-----+-----+-----+-----+-----+-----+-----+-----+
4 rows in set (1.36 sec)
```

In the example just shown, each filter includes a process type specifier. If you omit this specifier from one of the filters, you obtain a result that you might not expect:

```
mcm> get HostName,DataDir:ndbd mycluster;
+-----+-----+-----+-----+-----+-----+-----+-----+
| Name      | Value      | Process1 | Id1 | Process2 | Id2 | Level | Comment |
+-----+-----+-----+-----+-----+-----+-----+-----+
| HostName  | grindval   | ndb_mgmd | 49  |           |     |       | Read only |
| DataDir   | /opt/c2data | ndbd     | 1   |           |     |       |         |
| HostName  | tonfisk     | ndbd     | 1   |           |     |       | Read only |
| DataDir   | /opt/c3data | ndbd     | 2   |           |     |       |         |
| HostName  | flundra     | ndbd     | 2   |           |     |       | Read only |
| HostName  | haj        | mysql    | 50  |           |     |       | Read only |
| HostName  | torsk      | mysql    | 51  |           |     |       | Read only |
+-----+-----+-----+-----+-----+-----+-----+-----+
6 rows in set (0.58 sec)
```

The filter list `HostName,DataDir:ndbd` is perfectly valid. However, it actually consists of the filters `HostName` and `DataDir:ndbd`—in other words, it means “the `HostName` for all processes, and the `DataDir` for `ndbd` processes”.

Suppose you wish to obtain the values for `HostName` for just the `ndb_mgmd` and `mysql` processes in `mycluster`. You might be tempted to try using something like `HostName:ndb_mgmd,mysql` for the filter list, but this does not work, as you can see here:

```
mcm> get HostName:ndb_mgmd,mysql mycluster;
ERROR 6003 (00MGR): No such config variable mysql for process
```

This is due to the fact that each filter in the filter list must be a valid filter, and must include an attribute name. (In the filter list just shown, MySQL Cluster Manager tries to interpret the first string following the comma as an attribute name.) The correct filter list to use in a `get` command for retrieving the `HostName` for the `ndb_mgmd` and `mysql` processes in `mycluster` is shown in this example:

```
mcm> get HostName:ndb_mgmd,HostName:mysql mycluster;
+-----+-----+-----+-----+-----+-----+-----+-----+
| Name      | Value      | Process1 | Id1 | Process2 | Id2 | Level | Comment |
+-----+-----+-----+-----+-----+-----+-----+-----+
| HostName  | grindval   | ndb_mgmd | 49  |           |     |       | Read only |
| HostName  | haj        | mysql    | 50  |           |     |       | Read only |
| HostName  | torsk      | mysql    | 51  |           |     |       | Read only |
+-----+-----+-----+-----+-----+-----+-----+-----+
2 rows in set (0.21 sec)
```

It is also possible to obtain a list of attributes and their values for a given process type or instance of a process. For a given process type, use a filter having the form `:process_name`. For example, to retrieve all non-default attributes applying to `ndbd` processes in a cluster named `mycluster`, you can use the filter `:ndbd`, as shown here:

```
mcm> get :ndbd mycluster;
+-----+-----+-----+-----+-----+-----+-----+-----+
| Name      | Value      | Process1 | Id1 | Process2 | Id2 | Level | Comment |
+-----+-----+-----+-----+-----+-----+-----+-----+
| DataDir   | /opt/c2data | ndbd     | 1   |           |     |       |         |
| HostName  | tonfisk     | ndbd     | 1   |           |     |       | Read only |
| NodeId    | 1           | ndbd     | 1   |           |     |       | Read only |
| DataDir   | /opt/c3data | ndbd     | 2   |           |     |       |         |
| HostName  | flundra     | ndbd     | 2   |           |     |       | Read only |
| NodeId    | 2           | ndbd     | 2   |           |     |       | Read only |
+-----+-----+-----+-----+-----+-----+-----+-----+
```

```
6 rows in set (0.77 sec)
```

(The example just shown assumes that no attributes are set to non-default values.)

To get a list of all non-default attributes for a single instance of a process, use a filter having the form `:process_name:process_id`, as shown in this example, which retrieves all non-default attributes for the `ndbd` process having `2` as its process ID:

```
mcm> get :ndbd:2 mycluster;
```

Name	Value	Process1	Id1	Process2	Id2	Level	Comment
DataDir	/opt/c2data	ndbd	2				
HostName	flundra	ndbd	2				Read only
NodeId	2	ndbd	2				Read only

```
4 rows in set (0.32 sec)
```

If you try to obtain values for an attribute that you know is supported by your MySQL NDB Cluster version, but the result is empty, this almost certainly means that it is a default attribute that either has not been changed since the cluster was created or has been reset. In order to view default attributes using `get`, you must execute the command using the `--include-defaults` option (short form: `-d`).

Suppose you want to see how much `DataMemory` is configured for the `ndbd` processes in the cluster named `mycluster`, and you execute what appears to be the correct `get` command, but an empty result is returned, as shown here:

```
mcm> get DataMemory:ndbd mycluster;
Empty set (1.19 sec)
```

This means that the `DataMemory` attribute has its default value for all data nodes in the cluster. If you do not recall what this value is, you can determine it easily by repeating the same command with the addition of the `--include-defaults` (`-d`) option:

```
mcm> get --include-defaults DataMemory:ndbd mycluster;
```

Name	Value	Process1	Id1	Process2	Id2	Level	Comment
DataMemory	83886080	ndbd	1			Default	
DataMemory	83886080	ndbd	2			Default	

```
2 rows in set (0.62 sec)
```

Now suppose that you increase the `DataMemory` to 500 megabytes per data node, then repeat the `get` command to verify the new value:

```
mcm> set DataMemory:ndbd=500M mycluster;
```

```
Command result
```

```
Cluster reconfigured successfully
```

```
1 row in set (7.77 sec)
```

```
mcm> get --include-defaults DataMemory:ndbd mycluster;
```

Name	Value	Process1	Id1	Process2	Id2	Level	Comment
DataMemory	500M	ndbd	1			Process	
DataMemory	500M	ndbd	2			Process	

```
2 rows in set (1.46 sec)
```

You can see that, not only has the `Value` column in the `get` command output been updated to the new value, but the `Level` column has also been updated from `Default` to `Process`. This means that you no longer need the `--include-defaults` option to view this attribute, as shown here:

```
mcm> get DataMemory:ndbd mycluster;
```

```

+-----+-----+-----+-----+-----+-----+-----+-----+
| Name      | Value | Process1 | Id1 | Process2 | Id2 | Level | Comment |
+-----+-----+-----+-----+-----+-----+-----+-----+
| DataMemory | 500M | ndbd     | 1   |           |     | Process |         |
| DataMemory | 500M | ndbd     | 2   |           |     | Process |         |
+-----+-----+-----+-----+-----+-----+-----+-----+
2 rows in set (0.63 sec)

```

However, if you reset `DataMemory` (also on the process level), this is no longer the case. Then, `DataMemory` once again assumes its default value, after which you must use the `--include-defaults` option to retrieve it, as shown in this example:

```

mcm> reset DataMemory:ndbd mycluster;
+-----+-----+-----+-----+-----+-----+-----+-----+
| Command result |
+-----+-----+-----+-----+-----+-----+-----+-----+
| Cluster reconfigured successfully |
+-----+-----+-----+-----+-----+-----+-----+-----+
1 row in set (7.65 sec)

mcm> get DataMemory:ndbd mycluster;
Empty set (1.76 sec)

mcm> get --include-defaults DataMemory:ndbd mycluster;
+-----+-----+-----+-----+-----+-----+-----+-----+
| Name      | Value      | Process1 | Id1 | Process2 | Id2 | Level | Comment |
+-----+-----+-----+-----+-----+-----+-----+-----+
| DataMemory | 83886080 | ndbd     | 1   |           |     | Default |         |
| DataMemory | 83886080 | ndbd     | 2   |           |     | Default |         |
+-----+-----+-----+-----+-----+-----+-----+-----+
2 rows in set (1.01 sec)

```

For more information about these commands, see [Section 4.5.3, “The `set` Command”](#), and [Section 4.5.2, “The `reset` Command”](#).

The `get` command also tags multi-entry replication attributes as so in the `Comment` column; for example:

```

mcm> get replicate_ignore_table:mysqlqd mycluster;
+-----+-----+-----+-----+-----+-----+-----+-----+
| Name      | Value      | Process1 | NodeId1 | Process2 | NodeId2 | Level | Comment |
+-----+-----+-----+-----+-----+-----+-----+-----+
| replicate_ignore_table | mydb.t1      | mysqlqd  | 50      |           |         |       | Multi-ent |
| replicate_ignore_table | mydb.t50     | mysqlqd  | 50      |           |         |       | Multi-ent |
| replicate_ignore_table | mydb.mytable | mysqlqd  | 50      |           |         | Process | Multi-ent |
| replicate_ignore_table | mydb.t51     | mysqlqd  | 51      |           |         |       | Multi-ent |
| replicate_ignore_table | mydb.mytable | mysqlqd  | 51      |           |         | Process | Multi-ent |
+-----+-----+-----+-----+-----+-----+-----+-----+
5 rows in set (0.05 sec)

```

On how to reset multi-entry attributes, see [Section 4.5.2, “The `reset` Command”](#).

The `get` command does not normally display configuration attributes applying to TCP or SHM connections. However, such attributes can be set in the MySQL Cluster Manager client (using the `set` command); and once they have been set, they are displayed by applicable `get` commands. See [Setting TCP Connection Attributes](#), which provides an example of this.

4.5.2 The `reset` Command

```

reset [--sequential-restart] filter_specification_list cluster_name

filter_specification_list:
    filter_specification[,filter_specification][,...]

filter_specification:
    attribute_name[:process_specification][+process_specification]

process_specification:
    [process_name][:process_id]

```

```
process_name:
{ndb_mgmd|ndbd|ndbmtid|mysqld|ndbapi}
```

This command resets an attribute to its default value. Attributes can be set on either the process level or instance level. To reset an attribute on the process level, use a filter specification having the form `attribute_name:process_name`, where `attribute_name` is the name of the attribute to be reset, and `process_name` is the name of a MySQL NDB Cluster process. To reset a configuration attribute on the instance level, use a filter specification of the form `attribute_name:process_name:process_id`, where `process_id` is the process ID.

You cannot issue a `reset` command that resets all values for a given configuration attribute regardless of process type; each `reset` command must specify a process type or instance of a process. Otherwise, the command fails, as shown here:

```
mcm> reset DataMemory mycluster;
ERROR 3 (00MGR): Illegal syntax
```

You also cannot revert all configuration attributes for a given process type or instance of a process using a single filter specification; you must always include the name of the attribute to be reset. Otherwise, the `reset` command fails, as shown here:

```
mcm> reset :ndbd mycluster;
ERROR 3 (00MGR): Illegal syntax

mcm> reset :ndbd:3 mycluster;
ERROR 3 (00MGR): Illegal syntax
```

Suppose that the data memory for all `ndbd` processes in the cluster named `mycluster` has been set to 500 MB, as shown in the output of this `get` command:

```
mcm> get DataMemory mycluster;
+-----+-----+-----+-----+-----+-----+-----+
| Name      | Value | Process1 | Id1 | Process2 | Id2 | Level  | Comment |
+-----+-----+-----+-----+-----+-----+-----+
| DataMemory | 500M  | ndbd     | 2   |           |     | Process |         |
| DataMemory | 500M  | ndbd     | 3   |           |     | Process |         |
+-----+-----+-----+-----+-----+-----+-----+
2 rows in set (1.91 sec)
```

We can see from the entries in the `Level` column that the `DataMemory` setting for both `ndbd` processes applies on the process level. A process-level setting cannot be reset on the instance level, as shown here:

```
mcm> reset DataMemory:ndbd:2 mycluster;
ERROR 6010 (00MGR): No matching user defined setting was
found for config attribute DataMemory
mcm> reset DataMemory:ndbd:3 mycluster;
ERROR 6010 (00MGR): No matching user defined setting was
found for config attribute DataMemory
```

The following `reset` command also does not work, although you might think that it would do so, since it attempts to reset the attribute's value for both `ndbd` processes:

```
mcm> reset DataMemory:ndbd:2,DataMemory:ndbd:3 mycluster;
ERROR 6010 (00MGR): No matching user defined setting was
found for config attribute DataMemory
```

The previous command fails because MySQL Cluster Manager regards this as an attempt to apply two instance-level configuration changes. Because the `DataMemory` setting is a process-level setting, you must instead reset `DataMemory` to its default value on the process level; you can do this by using the filter specification `DataMemory:ndbd` in the `reset` command, as shown here:

```
mcm> reset DataMemory:ndbd mycluster;
+-----+-----+-----+-----+-----+-----+-----+
| Name      | Value | Process1 | Id1 | Process2 | Id2 | Level  | Comment |
+-----+-----+-----+-----+-----+-----+-----+
| DataMemory | 500M  | ndbd     | 2   |           |     | Process |         |
| DataMemory | 500M  | ndbd     | 3   |           |     | Process |         |
+-----+-----+-----+-----+-----+-----+-----+
2 rows in set (1.91 sec)
```

The `reset` Command

```
| Command result |
+-----+
| Cluster reconfigured successfully |
+-----+
1 row in set (6.16 sec)
```

If you execute the same `get` command as shown previously, the result is now empty:

```
mcm> get DataMemory mycluster;
Empty set (0.74 sec)
```

This is because the `get` command by default does not report default values. To retrieve the `DataMemory` values after resetting them, you must invoke `get` using the `--include-defaults` (short form: `-d`) option:

```
mcm> get --include-defaults DataMemory mycluster;
+-----+-----+-----+-----+-----+-----+-----+-----+
| Name      | Value   | Process1 | Id1 | Process2 | Id2 | Level  | Comment |
+-----+-----+-----+-----+-----+-----+-----+-----+
| DataMemory | 83886080 | ndbd     | 2   |           |     | Default |         |
| DataMemory | 83886080 | ndbd     | 3   |           |     | Default |         |
+-----+-----+-----+-----+-----+-----+-----+-----+
2 rows in set (1.21 sec)
```

The `DataMemory` values are now included in the output, and are marked with the word `Default` in the `Comments` column.

Now suppose that the `mysqld` configuration attribute `wait_timeout` for the `mysqld` process having the ID 4 in the cluster named `mycluster` has previously been set to the value 200 as shown here, and that no other changes have been to this attribute:

```
mcm> set wait_timeout:mysqld:4=200 mycluster;
+-----+
| Command result |
+-----+
| Cluster reconfigured successfully |
+-----+
1 row in set (7.78 sec)

mcm> get -d wait_timeout:mysqld:4 mycluster;
+-----+-----+-----+-----+-----+-----+-----+-----+
| Name      | Value | Process1 | Id1 | Process2 | Id2 | Level | Comment |
+-----+-----+-----+-----+-----+-----+-----+-----+
| wait_timeout | 200 | mysqld   | 4   |           |     |       |         |
+-----+-----+-----+-----+-----+-----+-----+-----+
1 row in set (0.98 sec)
```

Because the `Level` column is empty, we know that this setting applies on the instance level. If you try to reset it on the process level, the attempt fails, as shown here:

```
mcm> reset wait_timeout:mysqld mycluster2;
ERROR 6010 (OOMGR): No matching user defined setting was
found for config attribute wait_timeout
```

If you wish to reset this attribute to its default value, you must use the `reset` command with the instance-level filter specification `wait_timeout:mysqld:4`, as shown here:

```
mcm> reset wait_timeout:mysqld:4 mycluster;
+-----+
| Command result |
+-----+
| Cluster reconfigured successfully |
+-----+
1 row in set (7.61 sec)
```

Once you have reset `wait_timeout`, it no longer appears in the output of the earlier `get` command:

```
mcm> get wait_timeout:mysqld mycluster;
Empty set (1.42 sec)
```

This is because the default behavior of the `get` command is to display only those values that have been set either by the MySQL Cluster Manager or by the user. Since `wait_timeout` has been allowed to revert to its default value, you must use the `--include-defaults` (short form: `-d`) option to retrieve it, as shown here:

```
mcm> get -d wait_timeout:mysqlqd mycluster;
```

Name	Value	Process1	Id1	Process2	Id2	Level	Comment
wait_timeout	28800	mysqlqd	4			Default	

1 row in set (1.66 sec)

Now consider a situation in which process-level and instance-level settings have been made to a configuration attribute; in this example, we use `IndexMemory`. First, verify that `IndexMemory` is set to its default value for all data node processes (in this case, there are two of them):

```
mcm> get -d IndexMemory mycluster;
```

Name	Value	Process1	Id1	Process2	Id2	Level	Comment
IndexMemory	18874368	ndbd	2			Default	
IndexMemory	18874368	ndbd	3			Default	

2 rows in set (1.24 sec)

Now apply both a process-level change and an instance-level change to this attribute. You can do this with a single `set` command, as shown here:

```
mcm> set IndexMemory:ndbd=500M,IndexMemory:ndbd:3=750M mycluster;
```

Command result
Cluster reconfigured successfully

1 row in set (7.29 sec)

Because the process-level change was specified first, it is overridden for the `ndbd` process by the instance-level change specified second. The output from the following `get` command confirms that this is the case:

```
mcm> get IndexMemory mycluster;
```

Name	Value	Process1	Id1	Process2	Id2	Level	Comment
IndexMemory	500M	ndbd	2			Process	
IndexMemory	750M	ndbd	3				

2 rows in set (0.85 sec)

If the instance-level `IndexMemory` setting for the `ndbd` process with process ID 3 is reset, the process-level setting still applies, as shown here:

```
mcm> reset IndexMemory:ndbd:3 mycluster;
```

Command result
Cluster reconfigured successfully

1 row in set (6.41 sec)

```
mcm> get IndexMemory mycluster;
```

Name	Value	Process1	Id1	Process2	Id2	Level	Comment
IndexMemory	500M	ndbd	2			Process	
IndexMemory	500M	ndbd	3			Process	

2 rows in set (1.09 sec)

Now, re-apply the instance-level `IndexMemory` setting, and verify using `get` that it has taken effect:

```
mcm> set IndexMemory:ndbd:3=750M mycluster;
+-----+
| Command result |
+-----+
| Cluster reconfigured successfully |
+-----+
1 row in set (6.79 sec)

mcm> get IndexMemory mycluster;
+-----+-----+-----+-----+-----+-----+-----+-----+
| Name          | Value | Process1 | Id1 | Process2 | Id2 | Level | Comment |
+-----+-----+-----+-----+-----+-----+-----+-----+
| IndexMemory  | 500M  | ndbd     | 2   |           |     | Process |         |
| IndexMemory  | 750M  | ndbd     | 3   |           |     |         |         |
+-----+-----+-----+-----+-----+-----+-----+-----+
2 rows in set (1.76 sec)
```

If you reset the process-level setting, the instance-level setting remains, and only the `ndbd` process having process ID 2 has its `IndexMemory` reset to the default value; the instance-level setting remains in effect, as you can see from the following sequence of commands:

```
mcm> reset IndexMemory:ndbd mycluster;
+-----+
| Command result |
+-----+
| Cluster reconfigured successfully |
+-----+
1 row in set (7.36 sec)

mcm> get -d IndexMemory mycluster;
+-----+-----+-----+-----+-----+-----+-----+-----+
| Name          | Value   | Process1 | Id1 | Process2 | Id2 | Level   | Comment |
+-----+-----+-----+-----+-----+-----+-----+-----+
| IndexMemory  | 18874368 | ndbd     | 2   |           |     | Default |         |
| IndexMemory  | 750M     | ndbd     | 3   |           |     |         |         |
+-----+-----+-----+-----+-----+-----+-----+-----+
2 rows in set (0.10 sec)
```



Note

If the order of the specifiers in the original command that set `IndexMemory` had been reversed as `IndexMemory:ndbd:3=750M, IndexMemory:ndbd=500M`, the instance-level change would have been overridden by the process-level change, and the resulting `IndexMemory` setting for both `ndbd` processes would be `500M`. As discussed elsewhere, a process-level setting made after an instance-level setting that affects the same process completely removes the instance-level setting; the instance-level setting is not preserved, and resetting the attribute on the process level merely restores the default setting for all processes of that type. See [Section 4.5, “MySQL Cluster Manager Configuration Commands”](#), for more information.

The `get` and `reset` commands fully support multi-entry replication attributes; for example, if the `replicate_ignore_table` attribute has multiple entries:

```
mcm> get replicate_ignore_table:mysqld mycluster;
+-----+-----+-----+-----+-----+-----+-----+-----+
| Name          | Value          | Process1 | NodeId1 | Process2 | NodeId2 | Level   | Comment |
+-----+-----+-----+-----+-----+-----+-----+-----+
| replicate_ignore_table | mydb.t1       | mysqld   | 50      |           |         |         | Multi-ent |
| replicate_ignore_table | mydb.t50      | mysqld   | 50      |           |         |         | Multi-ent |
| replicate_ignore_table | mydb.mytable  | mysqld   | 50      |           |         | Process | Multi-ent |
| replicate_ignore_table | mydb.t51      | mysqld   | 51      |           |         |         | Multi-ent |
| replicate_ignore_table | mydb.mytable  | mysqld   | 51      |           |         | Process | Multi-ent |
+-----+-----+-----+-----+-----+-----+-----+-----+
5 rows in set (0.05 sec)
```

Without specifying a node ID, all the attribute's entries associated with the specified process type are reset with the following command:

```
mcm> reset replicate_ignore_table:mysql mycluster; # removes all process level entries
+-----+
| Command result |
+-----+
| Cluster reconfigured successfully |
+-----+
1 row in set (0.47 sec)

mcm> get replicate_ignore_table:mysql mycluster;
+-----+-----+-----+-----+-----+-----+-----+-----+
| Name | Value | Process1 | NodeId1 | Process2 | NodeId2 | Level | Comment |
+-----+-----+-----+-----+-----+-----+-----+-----+
| replicate_ignore_table | mydb.t1 | mysql | 50 | | | | Multi-entry |
| replicate_ignore_table | mydb.t50 | mysql | 50 | | | | Multi-entry |
| replicate_ignore_table | mydb.t51 | mysql | 51 | | | | Multi-entry |
+-----+-----+-----+-----+-----+-----+-----+-----+
3 rows in set (0.08 sec)
```

With a node ID specified, only the instance entries associated with the node ID are reset by the following command:

```
mcm> reset replicate_ignore_table:mysql:51 mycluster; # removes all instance level entries for nodeid 51
+-----+
| Command result |
+-----+
| Cluster reconfigured successfully |
+-----+
1 row in set (0.57 sec)

mcm> get replicate_ignore_table:mysql mycluster;
+-----+-----+-----+-----+-----+-----+-----+-----+
| Name | Value | Process1 | NodeId1 | Process2 | NodeId2 | Level | Comment |
+-----+-----+-----+-----+-----+-----+-----+-----+
| replicate_ignore_table | mydb.t1 | mysql | 50 | | | | Multi-entry |
| replicate_ignore_table | mydb.t50 | mysql | 50 | | | | Multi-entry |
+-----+-----+-----+-----+-----+-----+-----+-----+
2 rows in set (0.09 sec)
```

`reset` commands are executed whether or not the cluster has been started. In a cluster that is not running, the MySQL Cluster Manager merely updates the configuration files. However, in a running cluster, the MySQL Cluster Manager in addition automatically performs any node restarts or rolling restarts (see [Performing a Rolling Restart of an NDB Cluster](#)) that are required to cause the attribute changes to take effect (for MySQL Cluster Manager 1.4.8 and later, use the `--sequential-restart` option to make the rolling restart a [sequential](#) one). However, since restart operations—particularly rolling restarts—can take a great deal of time, it is preferable to make configuration changes before starting the cluster and putting it into use.

Resetting TCP Connection Attributes. Certain configuration attributes, such as those relating to TCP connections, apply to connections between processes rather than to individual processes or individual process types. As shown elsewhere (see [Setting TCP Connection Attributes](#)), when you set such an attribute on the process level using MySQL Cluster Manager, this means that the attribute applies to all connections between the two types of processes specified when issuing the `set` command. It is also possible to set such an attribute on the instance level, in which case it applies only to a single connection between two process instances.

Similarly, it is possible to reset such an attribute on either the process or instance level, depending on the level or levels at which it was set. In either case, an extended form of the process specifier is required, just as it is when setting an attribute that applies to a connection between processes. Assume that the `SendBufferMemory` attribute has previously been set for all connections between the two `ndbd` processes and the two `mysqld` processes that are found in a MySQL NDB Cluster named `mycluster2`, as shown in the output of this `get` command:

```
mcm> get SendBufferMemory mycluster2;
+-----+-----+-----+-----+-----+-----+-----+-----+
```

Name	Value	Process1	Id1	Process2	Id2	Level	Comment
SendBufferMemory	4M	ndbd	2	mysqld	4	Process	
SendBufferMemory	4M	ndbd	2	mysqld	5	Process	
SendBufferMemory	4M	ndbd	3	mysqld	4	Process	
SendBufferMemory	8M	ndbd	3	mysqld	5		

4 rows in set (0.59 sec)

Suppose that you wish to reset `SendBufferMemory` only for the connection between the `ndbd` process having process ID 3 and the `mysqld` process having process ID 5. The `SendBufferMemory` setting that applies to this connection is specified on the instance level, as you can see because the `Level` column value corresponding to this connection is empty; this means that it is possible to reset this value on the instance level. You can do this using the `reset` command shown here:

```
mcm> reset SendBufferMemory:ndbd:3+mysqld:5 mycluster2;
+-----+
| Command result |
+-----+
| Cluster reconfigured successfully |
+-----+
1 row in set (7.03 sec)
```

You can verify that the attribute was reset using the `get` command. However, as noted previously, once the instance-level setting has been removed, the process-level setting for this attribute again takes effect, so that the same setting applies to all connections between `ndbd` and `mysqld` processes, as shown here:

```
mcm> get SendBufferMemory mycluster2;
```

Name	Value	Process1	Id1	Process2	Id2	Level	Comment
SendBufferMemory	4M	ndbd	2	mysqld	4	Process	
SendBufferMemory	4M	ndbd	2	mysqld	5	Process	
SendBufferMemory	4M	ndbd	3	mysqld	4	Process	
SendBufferMemory	4M	ndbd	3	mysqld	5	Process	

4 rows in set (0.87 sec)

To reset this attribute on the process level, you can use the following `reset` command:

```
mcm> reset SendBufferMemory:ndbd+mysqld mycluster2;
+-----+
| Command result |
+-----+
| Cluster reconfigured successfully |
+-----+
1 row in set (8.01 sec)
```

You can verify that the attribute has been reset for all connection between `ndbd` processes and `mysqld` processes, by using the `get` command, as shown here:

```
mcm> get -d SendBufferMemory mycluster2;
Empty set (1.39 sec)
```

As noted elsewhere in this manual (see [Section 4.5.1, "The get Command"](#)), the empty result set is to be expected in this case, even when `get` is invoked using the `--include-defaults` (or `-d`) option, because the MySQL Cluster Manager client does not display attributes that appear in the `[tcp]` or `[shm]` sections of the `config.ini` configuration file if they have not been explicitly set by the user.

4.5.3 The `set` Command

```
set [--sequential-restart] attribute_assignment_list cluster_name

attribute_assignment_list:
    attribute_assignment[,attribute_assignment][,...]
```

```

attribute_assignment:
  attribute_name:process_specification[+process_specification][=value]

process_specification:
  [process_name][:process_id]

process_name:
  {ndb_mgmd|ndbd|ndbmtid|mysqld|ndbapi}
    
```

This command is used to set values for one or more configuration attributes. Attributes can be set on either the process level or instance level.

`set` commands are executed whether or not the cluster has been started. In a cluster that is not running, the MySQL Cluster Manager merely updates the configuration files. However, in a running cluster, the MySQL Cluster Manager in addition automatically performs any node restarts or rolling restarts (see [Performing a Rolling Restart of an NDB Cluster](#)) that are required to cause the attribute changes to take effect (for MySQL Cluster Manager 1.4.8 and later, use the `--sequential-restart` option to make the rolling restart a `sequential` one). However, since restart operations—particularly rolling restarts—can take a great deal of time, it is preferable to make configuration changes before starting the cluster and putting it into use.

To set an attribute on the process level, use a `set` statement that contains an attribute assignment having the form `attribute_name:process_name=value`.

For example, to set `DataMemory` to 500 MB on the `ndbd` process level, so that the new value applies to all `ndbd` processes in the cluster, you can issue a `set` command containing the attribute assignment `DataMemory:ndbd=500M`, as shown here:

```

mcm> set DataMemory:ndbd=500M mycluster;
+-----+
| Command result |
+-----+
| Cluster reconfigured successfully |
+-----+
1 row in set (5.68 sec)
    
```

To verify that the new setting is being used, you can issue the following `get` command:

```

mcm> get DataMemory mycluster;
+-----+-----+-----+-----+-----+-----+-----+-----+
| Name      | Value | Process1 | Id1  | Process2 | Id2  | Level  | Comment |
+-----+-----+-----+-----+-----+-----+-----+-----+
| DataMemory | 500M  | ndbd     | 1    |           |      | Process |         |
| DataMemory | 500M  | ndbd     | 2    |           |      | Process |         |
+-----+-----+-----+-----+-----+-----+-----+-----+
2 rows in set (0.79 sec)
    
```



Note

For more information about this command, see [Section 4.5.1, “The get Command”](#).

To set an attribute for a specific process instance, include the process ID in the attribute assignment; the form of such an attribute assignment is `attribute_name:process_name:process_id=value`. For example, to set the `wait_timeout` attribute for the `mysqld` process that has process ID 50 to 200, you would issue a `set` command that contains the attribute assignment `wait_timeout:mysqld:50=200`, like this:

```

mcm> set wait_timeout:mysqld:50=200 mycluster;
+-----+
| Command result |
+-----+
| Cluster reconfigured successfully |
+-----+
1 row in set (6.18 sec)
    
```

You can verify that the setting has taken effect using an applicable `get` command:

```
mcm> get wait_timeout mycluster;
+-----+-----+-----+-----+-----+-----+-----+
| Name          | Value | Process1 | Id1 | Process2 | Id2 | Level | Comment |
+-----+-----+-----+-----+-----+-----+-----+
| wait_timeout  | 200   | mysql    | 50  |           |     |       |         |
+-----+-----+-----+-----+-----+-----+
1 row in set (0.50 sec)
```

Attributes that are marked `Read only` cannot be set. Attempting to do so fails with an error, as shown here:

```
mcm> get :ndbd mycluster;
+-----+-----+-----+-----+-----+-----+-----+
| Name          | Value          | Process1 | Id1 | Process2 | Id2 | Level | Comment |
+-----+-----+-----+-----+-----+-----+-----+
| DataDir       | /opt/c2data    | ndbd     | 1   |           |     |       |         |
| HostName      | tonfisk        | ndbd     | 1   |           |     |       | Read only |
| NodeId        | 2              | ndbd     | 1   |           |     |       | Read only |
| DataDir       | /opt/c3data    | ndbd     | 2   |           |     |       |         |
| HostName      | grindval       | ndbd     | 2   |           |     |       | Read only |
| NodeId        | 3              | ndbd     | 2   |           |     |       | Read only |
+-----+-----+-----+-----+-----+-----+
6 rows in set (1.42 sec)
mcm> set HostName:ndbd:1=lax mycluster;
ERROR 6008 (00MGR): Config attribute HostName is read only and cannot be changed
```

However, you can set mandatory attributes, such as in the example shown previously in this section where the `DataDir` configuration attribute was set to a user-defined value.



Warning

The mandatory `NoOfReplicas` attribute must be set on the process level only. Attempting to set it on the instance level may leave the cluster, the MySQL Cluster Manager, or both in an unusable configuration.

Unlike the case with the `get` command, you cannot issue a `set` acting on a “global” scope—that is, you cannot, in a single attribute assignment, set a single value for an attribute such that the new attribute value applies to all processes regardless of process type, even if the attribute having that name can be applied to all process types. Nor can you specify multiple process types in a single attribute assignment. Attempting to do either of these things causes an error, as shown here:

```
mcm> set DataDir=/var/cluster-data mycluster;
ERROR 3 (00MGR): Illegal syntax

mcm> set DataDir:ndb_mgmd,ndbd,mysqld=/var/cluster-data mycluster;
ERROR 3 (00MGR): Illegal syntax
```

Instead, you must use a process-level attribute assignment for each process type. However, you are not necessarily required to issue a separate `set` command for each process type. Instead, you can also make multiple attribute assignments in a single `set` command, supplying the assignments as a comma-separated list. This `set` command assigns `/var/cdata` as the data directory (`DataDir`) for all MySQL NDB Cluster processes in the cluster named `mycluster`:

```
mcm> set DataDir:ndb_mgmd=/var/cdata, \
        DataDir:ndbd=/var/cdata, \
        DataDir:mysqld=/var/cdata mycluster;
+-----+-----+
| Command result |
+-----+-----+
| Cluster reconfigured successfully |
+-----+-----+
1 row in set (7.66 sec)
```

```
mcm> get DataDir mycluster;
```

Name	Value	Process1	NodeId1	Process2	NodeId2	Level	Comment
DataDir	/var/cdata	ndbmttd	1				
DataDir	/var/cdata	ndbmttd	2				
DataDir	/var/cdata	ndb_mgmd	49				
datadir	/var/cdata	mysqld	50				
datadir	/var/cdata	mysqld	51				

```
5 rows in set (0.08 sec)
```

As you can see from the `get` command just shown, the attribute assignments were successful, and took effect on the process level.



Note

In MySQL Cluster Manager, configuration attribute names are not case-sensitive. See [Case Sensitivity in String Searches](#) for more information about case-sensitivity issues in MySQL Cluster Manager.

Similarly, you cannot reference multiple process IDs in a single attribute assignment, even if they are processes of the same type; the following command does *not* work:

```
mcm> set DataMemory:ndbd:1,2=750M mycluster;
ERROR 3 (00MGR): Illegal syntax
```

Instead, you would need to use the following command:

```
mcm> set DataMemory:ndbd:1=750M,DataMemory:ndbd:2=750M mycluster;
```

Command result
Cluster reconfigured successfully

```
1 row in set (7.70 sec)
```

(Of course, if these are the only two data nodes in `mycluster`, then the command `set DataMemory:ndbd=750M mycluster` also accomplishes the same task.)



Note

A few configuration attributes apply to connections between processes and so require you to refer to both processes in the course of setting them. In such cases, you must use a special process specification syntax; see [Setting TCP Connection Attributes](#), for information about how this is done.

You also cannot set values for multiple attributes in a single attribute assignment; this means that the following commands do *not* work:

```
mcm> set UndoDataBuffer=32M,UndoIndexBuffer=8M:ndbd mycluster;
ERROR 3 (00MGR): Illegal syntax

mcm> set DataMemory,IndexMemory:ndbd=1G mycluster;
ERROR 3 (00MGR): Illegal syntax
```

However, if you write a complete and valid attribute assignment for each attribute whose value you wish to update, you can rewrite these two commands so that they execute successfully, as shown here:

```
mcm> set UndoDataBuffer:ndbd=32M,UndoIndexBuffer:ndbd=8M mycluster;
```

Command result
Cluster reconfigured successfully

```

+-----+
1 row in set (6.62 sec)

mcm> set DataMemory:ndbd=1G,IndexMemory:ndbd=1G mycluster;
+-----+
| Command result          |
+-----+
| Cluster reconfigured successfully |
+-----+
1 row in set (7.04 sec)

```

In fact, there is no reason that you cannot perform all four assignments in a single `set` command, using a list of four attribute assignments, like this:

```

mcm> set UndoDataBuffer:ndbd=32M,UndoIndexBuffer:ndbd=8M, \
      DataMemory:ndbd=1G, IndexMemory:ndbd=1G mycluster;
+-----+
| Command result          |
+-----+
| Cluster reconfigured successfully |
+-----+
1 row in set (6.24 sec)

```

However, it is a good idea not to perform too many attribute assignments in any single `set` command, since this makes it more difficult to spot errors.

On Windows, when setting attributes whose values contain paths (such as `DataDir`), you must replace any backslash characters in the path with forward slashes. Suppose that you want to use `c:\temp\node50` for the `tmpdir` attribute of the `mysqld` process having node ID 50 in a MySQL NDB Cluster named `mycluster` that is running on Windows. The original value for this attribute can be seen using the appropriate `get` command:

```

mcm> get tmpdir mycluster;
+-----+-----+-----+-----+-----+-----+-----+-----+
| Name   | Value           | Process1 | Id1  | Process2 | Id2  | Level | Comment |
+-----+-----+-----+-----+-----+-----+-----+-----+
| tmpdir | c:\c50data\tmp | mysql    | 50   |           |      |      |         |
+-----+-----+-----+-----+-----+-----+-----+-----+
1 row in set (0.22 sec)

```

The correct `set` command to make the desired configuration change is shown here:

```

mcm> set tmpdir:mysql:50=c:/temp/node50 mycluster;
+-----+
| Command result          |
+-----+
| Cluster reconfigured successfully |
+-----+
1 row in set (2.62 sec)

```

When you check the value using `get`—even though it was originally shown using backslashes—the forward slashes are used when displaying the new value:

```

mcm> get tmpdir mycluster;
+-----+-----+-----+-----+-----+-----+-----+-----+
| Name   | Value           | Process1 | Id1  | Process2 | Id2  | Level | Comment |
+-----+-----+-----+-----+-----+-----+-----+-----+
| tmpdir | c:/temp/node50 | mysql    | 50   |           |      |      |         |
+-----+-----+-----+-----+-----+-----+-----+-----+
1 row in set (0.22 sec)

```

However, if you try to use backslashes in the path when issuing the `set` command, the command fails:

```

mcm> set tmpdir:mysql:4=c:\temp\4 mycluster;
Outfile disabled.
ERROR:
Unknown command '\4'.
ERROR 6014 (00MGR): Path name for parameter tmpdir must be absolute.
The value 'c:mp4' is illegal.

```

Setting TCP Connection Attributes. For a few attributes that apply only when using TCP connections (such as the `SendBufferMemory` and `ReceiveBufferMemory` attributes), it is necessary to use a modified syntax for attribute value assignments. In this case, the attribute assignment contains two process specifications, one for each process type or instance to which the setting applies, joined with a plus sign (+). For the following example, consider the cluster named `mycluster2`, consisting of the processes shown here:

```
mcm> list processes mycluster2;
+-----+-----+-----+
| Id    | Name      | Host      |
+-----+-----+-----+
| 49    | ndb_mgmd  | grindval  |
| 1     | ndbd      | tonfisk   |
| 2     | ndbd      | flundra   |
| 50    | mysqld    | haj       |
| 51    | mysqld    | torsk     |
+-----+-----+-----+
5 rows in set (0.16 sec)
```

(See [Section 4.6.3, “The `list processes` Command”](#), for more information about this command.)

TCP connection attributes are not shown in the output from the `get` command unless they have been set. This means that, prior to setting `SendBufferMemory` for the first time, you obtain an empty result if you try to retrieve its value, as shown here:

```
mcm> get SendBufferMemory mycluster2;
Empty set (0.18 sec)

mcm> get --include-defaults SendBufferMemory mycluster2;
Empty set (0.93 sec)
```

To set the `SendBufferMemory` to 4 MB for all TCP connections between data nodes and SQL nodes, you can use the command shown here:

```
mcm> set SendBufferMemory:ndbd+mysqld=4M mycluster2;
+-----+-----+
| Command result |
+-----+-----+
| Cluster reconfigured successfully |
+-----+-----+
1 row in set (6.44 sec)
```

If you check the attribute's value afterwards using `get`, you can see that the value is applied to all possible connections between each of the two `ndbd` processes and each of the two `mysqld` processes in `mycluster2`, thus there are four rows in the output:

```
mcm> get SendBufferMemory mycluster2;
+-----+-----+-----+-----+-----+-----+-----+
| Name          | Value | Process1 | Id1 | Process2 | Id2 | Level | Comment |
+-----+-----+-----+-----+-----+-----+-----+
| SendBufferMemory | 4M    | ndbd      | 2   | mysqld    | 4   | Process |
| SendBufferMemory | 4M    | ndbd      | 2   | mysqld    | 5   | Process |
| SendBufferMemory | 4M    | ndbd      | 3   | mysqld    | 4   | Process |
| SendBufferMemory | 4M    | ndbd      | 3   | mysqld    | 5   | Process |
+-----+-----+-----+-----+-----+-----+-----+
4 rows in set (1.63 sec)
```

To override this setting for only the connection between the data node with process ID 2 and the `mysqld` process (process ID 4), you can include the process ID in each of the two parts of the process specification, as shown here:

```
mcm> set SendBufferMemory:ndbd:2+mysqld:4=8M mycluster2;
+-----+-----+
| Command result |
+-----+-----+
| Cluster reconfigured successfully |
+-----+-----+
1 row in set (7.95 sec)
```

When you check the result using a `get` command, you can see that the new setting applies on the instance level, and only to the connection between processes having IDs 2 and 4; the process-level setting made previously still applies to the remaining 3 connections:

```
mcm> get SendBufferMemory mycluster2;
```

Name	Value	Process1	Id1	Process2	Id2	Level	Comment
SendBufferMemory	8M	ndbd	2	mysqld	50		
SendBufferMemory	4M	ndbd	2	mysqld	51	Process	
SendBufferMemory	4M	ndbd	3	mysqld	50	Process	
SendBufferMemory	4M	ndbd	3	mysqld	51	Process	

```
4 rows in set (0.24 sec)
```

You cannot set a connection attribute on the process level in one part of the process specification (that is, for one end of the connection) and on the instance level in the other. Attempting to do so fails with an error, as shown here:

```
mcm> set SendBufferMemory:ndbd+mysqld:4=2M mycluster2;
ERROR 3 (00MGR): Illegal syntax
mcm> set SendBufferMemory:ndbd:2+mysqld=2M mycluster2;
ERROR 3 (00MGR): Illegal syntax
```

Setting Attributes for `mysqld` nodes. When a dynamic variable is set, `mcmd` sends out a SET GLOBAL statement to the `mysqld` to apply the value and also saves the value to the `mysqld` configuration file, so the value can be applied again at the next `mysqld` restart; however, an immediate restart is triggered when a non-dynamic variable is set.

Setup for `mysqld` connection pooling. Enabling connection pooling for `mysqld` can be done by setting the `ndb-cluster-connection-pool` attribute to the desired number of connections, but also requires an extra step in creating the cluster.

Because the `mysqld` process attempts to make multiple connections to the cluster when connection pooling is enabled, the cluster must be configured with “spare” or “empty” connections. You can do this by adding (otherwise) unused `ndbapi` entries in the `process_host` list used in the `create cluster` command, as shown here:

```
mcm> create cluster -P mypackage
> -R ndb_mgmd@10.100.10.97,ndbd@10.100.10.98,ndbd@10.100.10.99, \
    mysqld@10.100.10.100,ndbapi@10.100.10.100, \
    ndbapi@10.100.10.100,ndbapi@10.100.10.100
> mycluster;
```

Command result
Cluster created successfully

```
1 row in set (6.58 sec)
```

After this, you can use a `set` command like this one to set the size of the connection pool according to the number of excess connections available in the `config.ini` file:

```
mcm> set ndb_cluster_connection_pool:mysqld=4;
```

user attribute not supported for `mysqld`. Trying to set the `user` attribute for a `mysqld` process is not currently supported, and results in a warning being written to the MySQL Cluster Manager log.

4.6 MySQL Cluster Manager Process Commands

This section contains information about MySQL Cluster Manager client commands used to start and stop MySQL NDB Cluster processes, and to determine which processes are currently running.

MySQL Cluster Manager, `ndb_mgm`, and starting or stopping processes. For a MySQL NDB Cluster under MySQL Cluster Manager control, it is recommended *not* to use the `ndb_mgm` command-

line client that comes with the MySQL NDB Cluster distribution to perform operations that involve starting or stopping nodes. These include but are not limited to the following `ndb_mgm` client commands (Commands in the NDB Cluster Management Client):

- `START`
- `STOP`
- `RESTART`
- `SHUTDOWN`

4.6.1 The `add process` Command

```
add process {--processhosts=-R }process_host_list
  [--set=attribute_assignment_list] [--verbose | -v] [--sequential-restart] cluster_name

process_host_list:
  process_name[:node_id]@host[,process_name@host[,...]]

process_name:
  {ndb_mgmd|ndbd|ndbmt|mysqld|ndbapi}

attribute_assignment_list:
  attribute_assignment[,attribute_assignment][,...]

attribute_assignment:
  attribute_name:process_name[=value]
```

This command adds to an existing cluster one or more processes, which are specified using a `process_host_list` with the `--processhosts` option, the format of which is the same as that used with the `create cluster` command. Any hosts referenced in the list must be members of the site to which the cluster belongs. In addition, all hosts must be resolvable.

For example, the following `add process` command adds two `mysqld` processes on hosts `tonfisk` and `flundra` to the cluster named `mycluster`:

```
mcm> add process --processhosts=mysqld@tonfisk,mysqld@flundra mycluster;
+-----+
| Command result          |
+-----+
| Processes added successfully |
+-----+
1 row in set (2 min 10.39 sec)
```

With the `--verbose` option, the command shows an updated process list, after the new processes have been added:

```
mcm> add process --processhosts=ndbmt@tonfisk,ndbmt@flundra --verbose mycluster;
+-----+
| NodeId | Name      | Host      |
+-----+
| 49     | ndb_mgmd | tonfisk   |
| 53     | ndb_mgmd | flundra   |
| 1      | ndbmt    | tonfisk   |
| 2      | ndbmt    | flundra   |
| 3      | ndbmt    | tonfisk   |
| 4      | ndbmt    | flundra   |
| 50     | mysqld   | tonfisk   |
| 51     | mysqld   | flundra   |
| 52     | ndbapi   | *         |
+-----+
9 rows in set (2 min 7.57 sec)
```

You can also manually assign a node ID to the new process you are adding to the cluster by adding “:*node_ID*.” after the `process_name`. For MySQL Cluster Manager 1.3.3 and earlier, trying to

manually assign node IDs less than 49 for `ndb_mgmd`, `mysqld`, or `ndbapi` fails with an error; the restriction, however, has been lifted since MySQL Cluster Manager 1.3.4. Nevertheless, you are still recommended to follow the best practice of reserving node ID 1 to 48 for data nodes. The following command adds two `ndbd` processes with node IDs 10 and 11 on hosts `tonfisk` and `flundra`, respectively, to `mycluster`:

```
mcm> add process --processhosts=ndbd:10@tonfisk,ndbd:11@flundra mycluster;
+-----+
| Command result |
+-----+
| Processes added successfully |
+-----+
1 row in set (2 min 13.40 sec)
```



Note

NDB 8.0 supports up to 144 data nodes (for release 8.0.18 and later). So, when managing an NDB 8.0 cluster, you are recommended to follow the best practice of reserving node ID 1 to 144 for data nodes.

If the cluster is not running when you run the `add process` command, it is recommended that you start all the new processes added by this command together using the `start process --added` command or start them together with the whole cluster using the `start cluster` command—besides starting the nodes, either of the two commands also initializes the added nodes and causes new cluster nodegroups to be formed by issuing a `CREATE NODEGROUP` command to the cluster. If the added nodes are started with `start process --initial` instead, you are then required to run `CREATE NODEGROUP` manually via the `ndb_mgm` client.

If the cluster is running when you run the `add process` command, a rolling restart for the cluster is performed at the end of the `add process` command. For *MySQL Cluster Manager 1.4.8 and later*, use the `--sequential-restart` option to make the rolling restart a `sequential` one.

Adding Free Processes

Using the `add process` command, you can add unmanaged `mysqld` processes, or `ndbapi` slots for `ndbapi` applications such as `ndb_restore`. To add an unmanaged `mysqld` process, prefix the hostname with the wildcard `*` (asterisk character):

```
mcm> add process --processhosts=mysqld*@tonfisk,mysqld*@flundra mycluster;
+-----+
| Command result |
+-----+
| Processes added successfully |
+-----+
1 row in set (2 min 3.14 sec)
```

To allow the unmanaged `mysqld` nodes to connect from any host, use the wildcard `*` (asterisk character) in place of the hostname or IP address:

```
mcm> add process --processhosts=mysqld@*,mysqld@* mycluster;
+-----+
| Command result |
+-----+
| Processes added successfully |
+-----+
1 row in set (2 min 3.14 sec)
```

The same applies to `ndbapi` slots for `ndbapi` applications such as `ndb_restore`: prefix the hostname with the wildcard character to limit connectivity to a specific host, or use only a wildcard, without hostname, to allow `ndbapi` applications from any host:

```
mcm> add process --processhosts=ndbapi*@tonfisk,ndbapi@* mycluster;
+-----+
| Command result |
+-----+
| Processes added successfully |
+-----+
```

```
+-----+
1 row in set (2 min 8.13 sec)
```

Because “free” processes are not managed by MySQL Cluster Manager, there is no need to run the `start process --added` command after they have been successfully added to the cluster.

Using `add process` to Simplify `create cluster` Commands

Processes added before the cluster is started for the first time are started with the cluster. This makes it possible to use this command to break down what would otherwise be very long `create cluster` commands. Consider the following set of commands that creates and then starts a cluster named `mycluster`:

```
create cluster --processhosts=ndb_mgmd@host1,ndbd@host1,ndbd@host2, \
  mysqld@host3,mysqld@host4 mycluster;
start cluster mycluster;
```

The long `create cluster` command can be divided into a shorter (and more manageable) version of itself, plus several `add process` commands. This set of commands performs the same task as the previous set, creating `mycluster` with exactly the same processes and hosts as before, and then starting it:

```
create cluster --processhosts=ndb_mgmd@host1 mycluster;
add process --processhosts=ndbd@host1,ndbd@host2 mycluster;
add process --processhosts=mysqld@host3,mysqld@host4 mycluster;
start cluster mycluster;
```

Notice that a process that is added to a cluster that was created using `create cluster --import` and before the import takes place is added with status `import`, which means it cannot be started or stopped using `start process` or `stop process` before an import has taken place.

Configuring a New Process when Adding it

A newly added process inherits its configuration attribute settings from those in effect for its process type on the parent cluster, or assume the default settings for that process type if none apply. Existing attribute settings in the cluster must have process-level scope to be inherited by new processes added later; instance-level settings set for existing process instances prior to adding any new ones do not apply to any of the added processes. (See [Configuration attributes](#), for more information about the scope of attribute settings.)

Inherited attribute settings can be overridden when adding processes; to do this, use the `add process` command's `--set` option. This option takes as its argument an attribute assignment list similar in format to that used with the `get` and `set` commands. Suppose that the current `ndbd` process-level setting in the cluster named `mycluster` for the `DataDir` attribute is `/home/users/ndb/cluster-data`, but you wish to add two new `ndbd` processes that use `/tmp/cluster/data` instead. You can do this using the following command:

```
mcm> add process --set=ndbd:DataDir=/tmp/cluster/data
> --processhosts=mysqld@tonfisk,mysqld@flundra
> mycluster;
```



Note

Unlike the way you use the `set` command, an equal sign (=) immediately following the `--set` option is required.

When setting attributes this way, which involves specifying paths for processes running on Windows, you must replace any backslashes (\) used with forward slashes (/), just as with the `set` command. See [Setting Attributes Containing Paths on Windows \[129\]](#), for more information.

After a process has been added using `add process`, you can also use the `set` command to modify its configuration attribute settings (or specify additional ones) as you would with any other cluster process being managed with MySQL Cluster Manager.

**Note**

When IPv6-enabled Windows systems are used as MySQL NDB Cluster hosts under MySQL Cluster Manager, you must reference these hosts using IPv4 addresses. Otherwise, MySQL Cluster Manager is unable to connect to the agent processes on those hosts. See [Section 5.1, “MySQL Cluster Manager Usage and Design Limitations”](#).

4.6.2 The `change process` Command

```
change process [--sequential-restart] old_proc_type[:proc-id]=new_proc_type cluster_name
old_proc_type | new_proc_type:
{ndbd|ndbmtd}
```

This command is used (MySQL NDB Cluster 7.0 and later) to change the process type for a given MySQL NDB Cluster process or group of MySQL NDB Cluster processes from one process type (*old-process-type*) to another process type (*new-process-type*).

Currently, the only two process types available for use with this command are `ndbd` and `ndbmtd`. This means that `change process` can be used to change the data node process running on one or more data nodes from the single-threaded data node daemon (`ndbd`) to the multithreaded data node daemon (`ndbmtd`) or vice versa.

By default, `change process` affects all data nodes running the *old-process-type*. By specifying an optional *process_id*, its action can be restricted to the data node having that process ID.

Suppose you have a cluster that is named `mycluster` and has two data nodes using `ndbd` processes, as reflected in the output of the following `show status` command:

```
mcm> show status --process mycluster;
+-----+-----+-----+-----+-----+
| NodeId | Process | Host   | Status | Nodegroup |
+-----+-----+-----+-----+-----+
| 49     | ndb_mgmd | flundra | running |           |
| 1      | ndbd     | tonfisk | running | n/a       |
| 2      | ndbd     | grindval | running | n/a       |
| 50     | mysqld   | haj     | running |           |
| 51     | mysqld   | torsk   | running |           |
| 52     | ndbapi   | *      | running |           |
+-----+-----+-----+-----+-----+
6 rows in set (0.06 sec)
```

To change both data nodes to so that they use multithreaded (`ndbmtd`) processes, issue the command shown here, without any *process_id* specifier:

```
mcm> change process ndbd=ndbmtd mycluster;
+-----+-----+
| Command result |
+-----+-----+
| Process changed successfully |
+-----+-----+
1 row in set (2 min 17.51 sec)
```

After the command has executed, you can verify that both data nodes are now using `ndbmtd` by checking the output of the appropriate `show status` command, as shown here:

```
mcm> show status --process mycluster;
+-----+-----+-----+-----+-----+
| NodeId | Process | Host   | Status | Nodegroup |
+-----+-----+-----+-----+-----+
| 49     | ndb_mgmd | flundra | running |           |
| 1      | ndbmtd   | tonfisk | running | n/a       |
| 2      | ndbmtd   | grindval | running | n/a       |
| 50     | mysqld   | haj     | running |           |
| 51     | mysqld   | torsk   | running |           |
| 52     | ndbapi   | *      | running |           |
+-----+-----+-----+-----+-----+
```

```

+-----+-----+-----+-----+-----+
6 rows in set (0.09 sec)

```

A rolling restart for the cluster is performed at the end of the `change process` command. For MySQL Cluster Manager 1.4.8 and later, use the `--sequential-restart` option to make the rolling restart a sequential one.



Note

The `change process` command can be used whether or not the cluster or the data node or data nodes to be changed are running. However, the command executes much more quickly if the data node or data nodes to be changed are not running. The next set of examples illustrates this.

It is possible (and sometimes desirable) to use `ndbd` and `ndbmtid` data node processes concurrently; thus, it is also possible using the `change process` command to change a single data node process from single-threaded to multithreaded, or from multithreaded to single-threaded. To do this, you must specify the data node process using its process ID.

First, we stop the cluster and verify that all processes are no longer running, as shown here:

```

mcm> stop cluster mycluster;
+-----+-----+-----+-----+
| Command result |
+-----+-----+-----+
| Cluster stopped successfully |
+-----+-----+-----+
1 row in set (22.93 sec)

mcm> show status --process mycluster;
+-----+-----+-----+-----+-----+
| NodeId | Process | Host | Status | Nodegroup |
+-----+-----+-----+-----+-----+
| 49 | ndb_mgmd | flundra | stopped | |
| 1 | ndbmtid | tonfisk | stopped | n/a |
| 2 | ndbmtid | grindval | stopped | n/a |
| 50 | mysqld | haj | stopped | |
| 51 | mysqld | torsk | stopped | |
| 52 | ndbapi | * | stopped | |
+-----+-----+-----+-----+-----+
6 rows in set (0.05 sec)

```

The following command changes only the node having the process ID 2 from using the multithreaded data node daemon to the single-threaded version:

```

mcm> change process ndbmtid:2=ndbd mycluster;
+-----+-----+-----+-----+
| Command result |
+-----+-----+-----+
| Process changed successfully |
+-----+-----+-----+
1 row in set (6.52 sec)

```

As you can see, `change process` operates much more quickly when the process to be changed is not running. As before, you can verify that the command succeeded using `show status`:

```

mcm> show status --process mycluster;
+-----+-----+-----+-----+-----+
| NodeId | Process | Host | Status | Nodegroup |
+-----+-----+-----+-----+-----+
| 49 | ndb_mgmd | flundra | stopped | |
| 1 | ndbmtid | tonfisk | stopped | n/a |
| 2 | ndbd | grindval | stopped | n/a |
| 50 | mysqld | haj | stopped | |
| 51 | mysqld | torsk | stopped | |
| 52 | ndbapi | * | stopped | |
+-----+-----+-----+-----+-----+
6 rows in set (0.07 sec)

```

To complete the example, we start the cluster again, using `start cluster`, then change node number 2 back from `ndbd` (single-threaded) to `ndbmttd` (multithreaded) using `change process`, then verify the change using `show status`:

```
mcm> start cluster mycluster;
+-----+
| Command result |
+-----+
| Cluster started successfully |
+-----+
1 row in set (36.43 sec)

mcm> change process ndbd:2=ndbmttd mycluster;
+-----+
| Command result |
+-----+
| Process changed successfully |
+-----+
1 row in set (2 min 10.41 sec)

mcm> show status --process mycluster;
+-----+-----+-----+-----+-----+
| NodeId | Process | Host | Status | Nodegroup |
+-----+-----+-----+-----+-----+
| 49 | ndb_mgmd | flundra | running | |
| 1 | ndbmttd | tonfisk | running | n/a |
| 2 | ndbmttd | grindval | running | n/a |
| 50 | mysqld | haj | running | |
| 51 | mysqld | torsk | running | |
| 52 | ndbapi | * | running | |
+-----+-----+-----+-----+-----+
6 rows in set (0.11 sec)
```

You can see that it can require much less time to stop the cluster, change a data node process, and then start the cluster again than it is to change the process while the cluster is running. However, if you do this, the cluster is not available while it is stopped.

As noted previously, `change process` works only with `ndbd` and `ndbmttd` processes; attempting to use any other process type causes the command to fail with an error, as shown here:

```
mcm> change process ndb_mgmd=mysqld mycluster;
ERROR 7009 (00MGR): Processes ndb_mgmd and mysqld are not interchangeable in this package
mcm> change process ndbd=mysqld mycluster;
ERROR 7009 (00MGR): Processes ndbd and mysqld are not interchangeable in this package
```

4.6.3 The `list processes` Command

```
list processes cluster_name
```

This command displays all processes making up a given cluster. The following example demonstrates how to list all processes that are part of the cluster named `mycluster`:

```
mcm> list processes mycluster;
+-----+-----+-----+
| NodeId | Name | Host |
+-----+-----+-----+
| 49 | ndb_mgmd | flundra |
| 1 | ndbd | tonfisk |
| 2 | ndbd | grindval |
| 50 | mysqld | haj |
| 51 | mysqld | torsk |
| 52 | ndbapi | * |
+-----+-----+-----+
6 rows in set (0.03 sec)
```

The `cluster_name` argument is required. If this argument is omitted, the command fails with an error, as shown here:

```
mcm> list processes;
```

```
ERROR 6 (00MGR): Illegal number of operands
```

4.6.4 The `start process` Command

```
start process [--initial|-i] process_id | --added} cluster_name
```

This command starts the MySQL NDB Cluster process having the process ID `process_id` in the cluster named `cluster_name`. The status of the process to be started, as shown by `show status --process`, must be `added`, `stopped`, or `failed` (only if the failed process has exited properly can it be restarted with the command).

This example demonstrates how to start the process having the process ID `1` belonging to the cluster `mycluster`:

```
mcm> start process 1 mycluster;
+-----+
| Command result |
+-----+
| Process started successfully |
+-----+
1 row in set (13.93 sec)
```

You can obtain process IDs for all processes in a given cluster using `show status --process` or `list processes`. These are the same as the node IDs for these processes as shown in the output of other `mcm` client commands such as `get` or in the output of `ndb_mgm -e "show"` (see [ndb_mgm — The NDB Cluster Management Client](#)).

When the `--initial` option (short form: `-i`) is used, the following happens:

- For a data node process, MySQL Cluster Manager starts it with the `--initial` option, causing the data node to rebuild its file system.
- For an SQL node process (only for MySQL Cluster Manager 1.4.2 and later), MySQL Cluster Manager rebuilds the `mysqld` data directory with the `mysqld --initialize-insecure` command for MySQL 5.7 and with the `mysql_install_db` command for MySQL 5.6. The node's data directory must be empty, or the reinitialization will not be attempted.

Invoking this command with the `--added` option rather than with a process ID starts all nodes that were added previously to the cluster using `add process` but not yet started. For the added data and `mysqld` nodes, the use of the `--added` option also implies the use of the `--initial` option, meaning that `mcmd` will attempt to initialize the added nodes (see description for the `--initial` option above). Also, when the `--added` option is used, once all the added nodes are running, a `CREATE NODEGROUP` command is issued to the management node for the creation of new nodegroups.

You cannot use this command to start a `mysqld` process in a stopped or unavailable cluster—trying to do so will cause an error. This applies, for example, to the case in which a cluster has been created for a cluster import, but the import is not yet completed (see [Section 4.4.1, “The create cluster Command”](#), and [Section 3.5, “Importing MySQL NDB Clusters into MySQL Cluster Manager”](#)).

4.6.5 The `stop process` Command

```
stop process process_id cluster_name
```

This command stops the running MySQL NDB Cluster process with the process ID `process_id` in the cluster named `cluster_name`. The status of the process as shown by `show status --process` must be `running`.

Suppose that the process ID of a data node in the cluster named `mycluster` is `3`. Then this data node can be stopped as shown here:

```
mcm> stop process 3 mycluster;
+-----+
| Command result |
+-----+
```

```

+-----+
| Process stopped successfully |
+-----+
1 row in set (33.07 sec)

```

You can use `show status --process` or `list processes` to obtain process IDs for all processes in a given cluster.

In the event of a disk failure where MySQL Cluster Manager loses its manager directory (including its repository), the agent is able to recover information from other agents, but it does not actually control processes any longer, although it can detect them. This is due to the fact that the MySQL Cluster Manager agent cannot access the PID files. In this case, `stop process` no longer works, and you must kill such processes manually. Keep in mind that, if `StopOnError` is set to 0, the MySQL Cluster Manager agent restarts the data node process automatically; if `StopOnError` is 1 (the default), then you must execute the `start process` command manually.

This command does not work with processes in a cluster created for import where the import has not yet actually been completed. See [Section 4.4.1, “The `create cluster` Command”](#), and [Section 3.5, “Importing MySQL NDB Clusters into MySQL Cluster Manager”](#), for more information.

4.6.6 The `update process` Command

```
update process [--remove-angel] --pid=os_pid process_id cluster_name
```

For MySQL Cluster Manager 1.4.2 and later: This command updates the status of the MySQL NDB Cluster process having the process ID `process_id` in the cluster named `cluster_name` when the status of the process is no longer reflected correctly in the output of the `show status --process` command. This typically happens in the following cases:

- The process is a data node configured with `StopOnError=true`, so that it would not be automatically restarted by `mcmd` after it has stopped. Instead of using the `start process` command to restart the process, a user might have restarted the process manually, which would have restored the process but left `mcmd` without the knowledge of the restore. An `update process` is then needed to restore the control of the process by `mcmd`.
- The process is a node that has been stopped by `mcmd` but, for some reasons, its PID remains valid with the operating system. In some cases, the process might even be running again, without `mcmd` knowing or being able to control it.
- `mcmd` cannot connect to an SQL node due to various reasons (for example, there are already too many connections to the node); process status for the node becomes `failed`, while the PID file continues to exist.

The command works by importing the process into the control of `mcmd` again. Checks performed on a process by `mcmd` during a cluster import are performed for the `update process` command. Both the process's ID in the cluster (`process_id`) and its PID on the operating system (specified with the `--pid` option) are required. Suppose that the process ID of a data node in the cluster named `mycluster` is 3 and its PID on the operating system is 9846, the data node can be updated as shown here:

```

mcm> update process --pid=9846 3 mycluster;
+-----+
| Command result          |
+-----+
| Process updated successfully |
+-----+
1 row in set (33.07 sec)

```

For a data node or an SQL node, the command only works if there is at least 1 replica per nodegroup running.

For MySQL Cluster Manager 1.4.7 and later, `update process` supports a `--remove-angel` option, which should be used when updating data nodes: it kills any running `angel process` for a data node and updates its PID file prior to the actual update; those steps are necessary for the update process.

4.6.7 The `remove process` Command

```
remove process [--removedirs] process_id_list cluster_name
```

```
process_id_list:
    process_id[, process_id[, ...]]
```

This command removes permanently the processes in the `process_id_list` from the cluster named `cluster_name`. It provides a means to scale down a cluster offline.

If the `--removedirs` option is used, all data for the specified processes will be deleted.

The following restrictions apply when using this command:

1. The cluster must be in the status of `created` or `stopped`.
2. The processes to be removed must be in the status of `stopped`, `added`, or `import`.
3. The command cannot remove all processes from a cluster in the `created` status; at least one process must be left.
4. The command cannot remove all process of the same type from a cluster in the `stopped` status; at least one process must be left in the cluster for each type of nodes (management, data, and API).
5. The command cannot remove a data node that is in the `stopped` status if it is already a member of a node group (i.e., if it has ever been started and was fully functional).

You can use the `show status --process` or `list processes` command to obtain the process IDs for all the processes in a given cluster:

```
mcm> show status --process mycluster;
+-----+-----+-----+-----+-----+-----+
| NodeId | Process | Host   | Status | Nodegroup | Package |
+-----+-----+-----+-----+-----+-----+
| 49     | ndb_mgmd | flundra | added  |           | mypackage |
| 1      | ndbmttd  | flundra | added  | n/a       | mypackage |
| 2      | ndbmttd  | flundra | added  | n/a       | mypackage |
| 50     | mysqld   | flundra | added  |           | mypackage |
| 51     | mysqld   | flundra | added  |           | mypackage |
| 52     | ndbapi   | *      | added  |           |           |
| 53     | ndbapi   | *      | added  |           |           |
+-----+-----+-----+-----+-----+-----+
7 rows in set (0.03 sec)
```

The process IDs are the same as the node IDs for the processes shown in the output of the above or some other `mcm` client commands, or in the output of the `ndb_mgm -e "show"` command (see [ndb_mgm — The NDB Cluster Management Client](#)). In the above example, the SQL node with the process ID 50 in `mycluster` can be removed by the following command:

```
mcm> remove process 50 mycluster;
+-----+-----+
| Command result |
+-----+-----+
| Process removed successfully |
+-----+-----+
1 row in set (0.48 sec)
```

And in this case, since the cluster was never started, we may also remove both data nodes:

```
mcm> remove process 1,2 mycluster;
+-----+-----+
| Command result |
+-----+-----+
| Process removed successfully |
+-----+-----+
1 row in set (0.40 sec)
```

4.7 MySQL Cluster Manager Backup and Restore Commands

This section contains information about MySQL Cluster Manager client commands relating to backing up a MySQL NDB Cluster and restoring it from backup.

4.7.1 The `abort backup` Command

```
abort backup --backupid=backup_id cluster_name
```

This command aborts a backup of cluster *cluster_name* having the ID *backup_id* specified with the `--backupid` option. You can obtain a list of backups and their IDs known to this MySQL Cluster Manager instance using the `list backups`. If the backup is not actually in progress, the command has no effect.

4.7.2 The `backup cluster` Command

```
backup cluster [--backupid=backup_id]
               [--snapshotstart | --snapshotend]
               [--waitstarted | --waitcompleted]
               cluster_name
```

This command creates a backup of the MySQL NDB Cluster named *cluster_name*. `backup cluster` takes a backup of the cluster's NDB tables only; tables using other MySQL storage engines (such as InnoDB or MyISAM) are ignored.

By default, this command uses the backup ID assigned and returned by `ndb_mgmd` (see the discussions on *backup_id* in [Using The NDB Cluster Management Client to Create a Backup](#) for more information); you can override this behavior by specifying a backup ID using the `--backupid` option.

The `--snapshotstart` option causes the backup to match the state of the cluster when the backup began.

The `--snapshotend` option causes the backup to reflect the state of the cluster when the backup was finished. If neither option is specified, the MySQL Cluster Manager client acts as though `--snapshotend` had been used.

When the `--waitstarted` option is used, the MySQL Cluster Manager client waits until the backup has started before returning control to the user, after which the user can check the backup process's status with the `show status` command and the `--backup` option.

When the `--waitcompleted` option is used, the MySQL Cluster Manager client waits until the backup process is complete before returning control to the user. If neither `--waitstarted` or `--waitcompleted` is specified, the client behaves as if `--waitcompleted` had been used.

```
mcm> backup cluster mycluster;
+-----+
| Command result |
+-----+
| Backup completed successfully |
+-----+
1 row in set (33.50 sec)
```

You can verify that the backup was performed by checking the output of `list backups`, as shown here:

```
mcm> list backups mycluster;
+-----+-----+-----+-----+-----+
| BackupId | NodeId | Host | Timestamp | Comment |
+-----+-----+-----+-----+-----+
| 1 | 1 | tonfisk | 2016-10-24 22:24:54Z | |
| 1 | 2 | tonfisk | 2016-10-24 22:24:54Z | |
```

```

| 2 | 1 | tonfisk | 2016-10-24 22:24:54Z |
| 2 | 2 | tonfisk | 2016-10-24 22:24:54Z |
+-----+-----+-----+-----+
4 rows in set (0.02 sec)

```

Each row in the output represents a backup *image*—that is, a set of backup files specific to a given backup of a named cluster on a given data node. `Timestamp` values are in UTC. The backup image is saved in the folder `BackupDataDir/BACKUP/BACKUP-Id`, where `BackupDataDir` is a cluster parameter. If `BackupDataDir` is not specified, it takes up the value of `DataDir`, so that the image is stored in the directory `DataDir/BACKUP/BACKUP-backup_id`.

It is possible to remove an unwanted backup from a given node by deleting this image directory and its contents. To remove a given backup completely, you must remove the corresponding image from each data node's `BACKUP` directory. You can do this as long as a backup or restore operation is not in progress. It is not necessary to stop the cluster or MySQL Cluster Manager agent prior to removing the images.

The `BackupId` is used with `abort backup` and `restore cluster`.

Logical Backup for NDB Table Metadata

To allow more flexibility for cluster reconfiguration during a restore, starting with release 1.4.1, the `backup cluster` command also creates a logical backup for the metadata of the NDB tables in the cluster. Use the `--all` option with the `list backups` command to list all backups, including the logical backups for the NDB tables' metadata, which are marked by the comment "Schema":

```

mcm> list backups --all newcluster;
+-----+-----+-----+-----+-----+
| BackupId | NodeId | Host      | Timestamp           | Comment |
+-----+-----+-----+-----+-----+
| 1 | 1 | tonfisk | 2016-08-12 16:55:52Z | |
| 1 | 2 | tonfisk | 2016-08-12 16:55:52Z |
| 1 | 3 | tonfisk | 2016-08-12 16:55:52Z |
| 1 | 4 | tonfisk | 2016-08-12 16:55:52Z |
| 1 | 50 | tonfisk | 2016-08-12 16:55:55Z | Schema |
+-----+-----+-----+-----+-----+
5 rows in set (0.02 sec)

```

The logical backup was created using the `mysqldump` utility. The backup is saved with the file name `BACKUP-BackupID.mysql_nodeid.schema.sql` extension, to be found in the folder `backupdatadir/BACKUP/BACKUP-id`, where `backupdatadir` (notice that the name is in lowercase) is a `mysqld` parameter used only for specifying the location of the logical backup created by MySQL Cluster Manager. If `backupdatadir` is not specified using the `set` command with the `mcm` client, the default value of `/mcm_data_repository/clusters/clustername/mysqld_nodeid/` is used, so that the logical backup is saved in the folder `/mcm_data_repository/clusters/clustername/mysqld_nodeid/BACKUP/BACKUP-Id`.

The following restrictions apply for the creation of the logical backups for NDB table metadata:

- At least one `mysqld` node must be running on the cluster for the logical backup to be performed
- No backup was created for any `mysqld` node that was not running.
- Metadata for non-NDB tables are not backed up.
- The logical backup is NOT a proper point-in-time backup—no DDL operations should be performed on the cluster when the backup process is running on the cluster, or the backed-up metadata will become inconsistent with the backed-up data.

The backup for the NDB table metadata is helpful for restoring data from a cluster to another one with a different configuration (for example, when the target cluster for restore has a different number of data

nodes); see [Section 3.6.2.4, “Partial restore—data nodes added”](#) and [Section 3.6.2.5, “Restoring a Backup to a Cluster with Fewer Data Nodes”](#) for some use cases.

4.7.3 The `list backups` Command

```
list backups [--backupid=I] [backup_id] [--all|-a] cluster_name
```

```
list backups [--backupid=I] [backup_id] [--agent|-A] site_name
```

Without the `--agent` option, the command lists all backups of the MySQL NDB Cluster named `cluster_name` that are known to this instance of MySQL Cluster Manager. The output includes the backup and node ID as well as a UTC timestamp for each backup, as shown here:

```
mcm> list backups mycluster;
```

BackupId	NodeId	Host	Timestamp	Comment
1	1	tonfisk	2016-10-24 22:24:54Z	
1	2	tonfisk	2016-10-24 22:24:54Z	
2	1	tonfisk	2016-10-24 22:24:54Z	
2	2	tonfisk	2016-10-24 22:24:54Z	

```
4 rows in set (0.02 sec)
```

The `Timestamp` column shows the timestamp (in UTC) of the first file to appear in any backup instance folder. There are 3 files in each backup filesset: `*.ctl`, `*.data`, and `*.log`. If the backup instance folder is empty, the timestamp of the folder itself is shown.

With the `--backupid` option used, the commands only list backups with the specified ID:

```
mcm> list backups --backupid=2 mycluster;
```

BackupId	NodeId	Host	Timestamp	Comment
2	1	tonfisk	2016-10-24 22:24:54Z	
2	2	tonfisk	2016-10-24 22:24:54Z	

```
2 rows in set (0.02 sec)
```

For MySQL Cluster Manager 1.4.1 and later. The `backup cluster` command also creates backups of the metadata for a cluster's NDB tables, which are listed by the `list backups` command when the `--all` option is used. The metadata backups are marked by the comment `Schema` in the backup listing:

```
mcm> list backups --all newcluster;
```

BackupId	NodeId	Host	Timestamp	Comment
1	1	tonfisk	2016-08-12 16:55:52Z	
1	2	tonfisk	2016-08-12 16:55:52Z	
1	3	tonfisk	2016-08-12 16:55:52Z	
1	4	tonfisk	2016-08-12 16:55:52Z	
1	50	tonfisk	2016-08-12 16:55:55Z	Schema

```
5 rows in set (0.02 sec)
```

See [Logical Backup for NDB Table Metadata](#) for details about the metadata backup.

For MySQL Cluster Manager 1.4.6 and later. When the `--agent` option is used and a `site_name` is specified, the command lists agent backups created for a specific site:

```
mcm> list backups --agent mysite;
```

BackupId	Agent	Host	Timestamp	Comment
1522914101	0	tonfisk	2018-04-05 07:41:41Z	
1522914105	0	tonfisk	2018-04-05 07:41:45Z	
1522914121	0	tonfisk	2018-04-05 07:42:01Z	

```
+-----+-----+-----+-----+-----+
3 rows in set (0.00 sec)
```

The backup IDs reflect the Unix Epoch times at which the backups were taken.

The output can be filtered with the `--backupid` option:

```
mcm> list backups --agent --backupid=1522914121 mysite;
+-----+-----+-----+-----+-----+
| BackupId | Agent | Host   | Timestamp                | Comment |
+-----+-----+-----+-----+-----+
| 1522914121 | 0     | tonfisk | 2018-04-05 07:42:01Z    |         |
+-----+-----+-----+-----+-----+
1 row in set (0.07 sec)
```

4.7.4 The `restore cluster` Command

```
restore cluster
  {--backupid=-I }backup_id
  [--disable-indexes|-x]
  [--disable-metadata|-M]
  [--epoch|-e]
  [--exclude-databases=db_name]
  [--exclude-intermediate-sql-tables]
  [--exclude-missing-columns]
  [--exclude-missing-tables]
  [--exclude-tables=db_name.tbl_name[,db_name.tbl_name][,...]]
  [--include-databases=db_name]
  [--include-stored-grants]
  [--include-tables=db_name.tbl_name[,db_name.tbl_name][,...]]
  [--lossy-conversions]
  [--no-binlog|-l]
  [--no-restore-disk-objects]
  [{--parallelism=-p }#]
  [--privilege-tables|-P]
  [--progress-frequency]
  [--promote-attributes]
  [--rewrite-database]
  [--skip-broken-objects]
  [{--skip-nodeid=-s }id_list]
  [--skip-table-check]
  [--skip-unknown-objects]
  cluster_name
```

This command restores a cluster from a backup having the specified backup ID (`--backupid` option; short form: `-I`) to the MySQL NDB Cluster named `cluster_name`. In its simplest form, it can be used as shown here, to restore the cluster named `mycluster` to the state saved in the backup having backup ID 3:

```
mcm> restore cluster --backupid=3 mycluster;
+-----+-----+
| Command result |
+-----+-----+
| Restore completed successfully |
+-----+-----+
1 row in set (18.60 sec)
```

If you are restoring an existing cluster to a known good state, you must wipe any existing data first. Stop the cluster using `stop cluster`, then restart it using `start cluster` with the `--initial` option, which causes the data node file systems to be cleared. (Note that for MySQL NDB Cluster 7.5 and earlier, Disk Data files must be removed manually.) Following this, you can restore the cluster from the desired backup using `restore cluster`.



Important

In order to restore a backup using `restore cluster`, the cluster must have an unused slot for an `ndbapi` process in its configuration. Otherwise, the

command fails with the error `Unable to perform restore - no vacant ndbapi slots in config for cluster cluster_name`. See [Adding Free Processes](#), for information on how to add a free `ndbapi` slot to your cluster.

Additional options that can be employed with this command include:

--disable-indexes and --disable-metadata. To cause indexes to be ignored when restoring the table data, use the `--disable-indexes` option (short form: `-x`; *only for MySQL Cluster Manager 1.4.7 and earlier*). Doing this can decrease the time required to restore a large data set, particularly where many indexes were in use. Similarly, you can cause metadata to be ignored during the restoration process by using the `--disable-metadata` option (short form: `-M`).

--epoch. When the `--epoch` option (short form: `-e`) is used, epoch information is restored to the cluster replication status table (`mysql.ndb_apply_status`), which can be useful for replicas in MySQL NDB Cluster replication.

--exclude-databases and --exclude-tables. Prevent one or more databases or tables from being restored using the options `--exclude-databases` and `--exclude-tables`. `--exclude-databases` takes a comma-delimited list of one or more databases that should not be restored. `--exclude-tables` takes a comma-delimited list of one or more tables (using the `database.table` format) that should not be restored. When `--exclude-databases` or `--exclude-tables` is used, only those databases or tables named by the option are excluded; all other databases and tables are restored.

--exclude-missing-columns. When this option is used, `restore cluster` ignores any columns missing from tables being restored as compared to the versions of those tables found in the backup.

--exclude-missing-tables. When this option is used, `restore cluster` ignores any tables from the backup that are not found in the target database.

--exclude-intermediate-sql-tables[=TRUE|FALSE]. When performing `ALTER TABLE` operations, `mysqld` creates intermediate tables (whose names are prefixed with `#sql-`). When `TRUE`, the `--exclude-intermediate-sql-tables` option keeps `restore cluster` from restoring such tables that may have been left over from such operations. This option is `TRUE` by default.

--include-databases and --include-tables. Use the `--include-databases` option or the `--include-tables` option for restoring only specific databases or tables, respectively. `--include-databases` takes a comma-delimited list of databases to be restored. `--include-tables` takes a comma-delimited list of tables (in the `database.table` format) to be restored. When `--include-databases` or `--include-tables` is used, only those databases or tables named by the option are restored; all other databases and tables are excluded by `restore cluster`, and are not restored.

--include-stored-grants. *When managing NDB Cluster 8.0.19 and later*, the `restore cluster` command does not restore [shared users and grants](#) to the `mysql.ndb_sql_metadata` table by default; use the `--include-stored-grants` option (*available only for MySQL Cluster Manager 1.4.8 and later*) to override this behavior and enable the restore of shared user and grant data and metadata.

--lossy-conversions. Using `--lossy-conversions` allows lossy conversions of column values (type demotions or changes in sign) when restoring data from backup. With some exceptions, the rules governing demotion are the same as for MySQL replication; see [Replication of Columns Having Different Data Types](#), for information about specific type conversions currently supported by attribute demotion. `restore cluster` reports any truncation of data that it performs during lossy conversions once per attribute and column.

--no-binlog. The `--no-binlog` option (short form: `-l`) stops any SQL nodes (`mysqld` processes) in the cluster from writing data from the restore into their binary logs.

--no-restore-disk-objects. This option stops `restore cluster` from restoring any MySQL NDB Cluster Disk Data objects, such as tablespaces and log file groups; see [NDB Cluster Disk Data Tables](#), for more information about these objects.

--parallelism=#. The `--parallelism` option (short form: `-p`) sets the maximum number of parallel transactions that the `restore cluster` command attempts to use. The default value is 128; the maximum is 1024, and the minimum is 1.

--privilege-tables. The `--privilege-tables` option (short form: `-P`) causes restoration of tables required for distributed grants (see [Distributed Privileges Using Shared Grant Tables](#)).

--progress-frequency=N. Print a status report each `N` seconds to a temporary stdout dump file `mcm` creates at `mcm_data/clusters/cluster_name/nodeid/tmp` while the backup is in progress. 0 (the default) causes no status reports to be printed. The maximum is 65535.

--promote-attributes. Allow attributes to be promoted when MySQL Cluster Manager restores data from a backup. See the [discussion on attribute promotion in the MySQL NDB Cluster manual](#) for more details.

--rewrite-database=old_dbname,new_dbname. This option causes a database with the name `old_dbname` in the backup to be restored under the name `new_dbname`.

--skip-nodeid. The `--skip-nodeid` option (short form: `-s`) takes a comma-separated list of node IDs. The nodes whose IDs are listed may include of data nodes, SQL nodes, or both. Nodes having these IDs are skipped by the restoration process.

--skip-broken-objects. This option causes `restore cluster` to ignore corrupt tables while reading a backup, and to continue restoring any remaining tables (that are not also corrupted). Currently, the `--skip-broken-objects` option works only in the case of missing blob parts tables.

--skip-table-check. It is possible to restore data without restoring table metadata. The default behavior when doing this is for `restore cluster` to fail with an error if table data do not match the table schema; this can be overridden using the `--skip-table-check` option.

--skip-unknown-objects. This option causes `restore cluster` to ignore any schema objects it does not recognize while reading a backup. This can be used for restoring, for example, a backup made from a newer version of MySQL NDB Cluster to an older version.

4.7.5 The `backup agents` Command

```
backup agents [--hosts=host_list] [site_name]
```

```
host_list:
  host[, host[, ...]]
```

This command backs up the configuration data for the `mcmd` agents on the hosts specified in `host_list` with the `--hosts` option (short form: `-h`) for the site named `site_name`. If no host names are specified, all agents of the site are backed up. If no `site_name` is given, only the agent that the `mcm` client is connected to is backed up.

The backup for each agent is created in a subfolder named `rep_backup/timestamp` under the agent repository (the `mcm_data` folder), with `timestamp` reflecting the time the backup began. If you want the backup to be at another place, create a soft link from `mcm_data/rep_backup` to your desired storage location.

For MySQL Cluster Manager 1.4.2 and later: An empty file named `INCOMPLETE` is created in the folder in which the backup is created when the backup 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.

Notice that the `backup agents` command works differently from the `backup cluster` command, which backs up cluster data; the `backup agents` command, on the other hand, backs up agent configuration data. Using together the backups created by both commands, you can restore not just the cluster, but the complete cluster-plus-manager setup. See [Section 3.7, “Backing Up and Restoring MySQL Cluster Manager Agents”](#) for more details about backing up and restoring `mcmd` agents.

4.8 MySQL Cluster Manager Cluster Importation Commands

This section contains descriptions of MySQL Cluster Manager commands used to perform operations connected with importing clusters into MySQL Cluster Manager. These operations include migration of cluster processes and copying of configuration data.

4.8.1 The `import cluster` Command

```
import cluster [--dryrun|-y] [--remove-angel] cluster_name
```

This command imports a MySQL NDB Cluster created independently of MySQL Cluster Manager into a cluster named `cluster_name` that has been created in MySQL Cluster Manager. You are strongly advised to create `cluster_name` using the `create cluster` command's `--import` option; see that command's description for more information about the `--import` option and its effects.

`import cluster` requires a single argument, the name of the cluster created using MySQL Cluster Manager (`cluster_name`) into which you wish to import a MySQL NDB Cluster created externally into MySQL Cluster Manager and bring it under MySQL Cluster Manager control. The cluster named in the command must already exist in MySQL Cluster Manager.

`import cluster` also supports a `--dryrun` option. When this option is used, only the checks required for importation are performed against the existing cluster. This makes it possible to test a given configuration without actually placing any cluster processes under MCM control. `-y` is supported as a short form of this option.

For MySQL Cluster Manager 1.4.7 and later, `import cluster` supports a `--remove-angel` option. When this option is used, any running `angel processes` for the data nodes of the cluster to be imported are stopped by `mcmd` prior to the actual import, which is a necessary step for a cluster import unless the angel processes have already been stopped manually. When this option is used together with the `--dryrun` option, no removals of angel processes will actually be performed, but the checks for angel processes (which occur when the `--dryrun` option is used alone) will be skipped. It is recommended that you use the two options separately: perform the checks with the `--dryrun` option only, and once the only errors observed are with the angel processes, run `import cluster` again with the `--remove-angel` option only to complete the import.

For more information about importing clusters into MySQL Cluster Manager, including examples, see [Section 3.5, “Importing MySQL NDB Clusters into MySQL Cluster Manager”](#).

4.8.2 The `import config` Command

```
import config [--dryrun|-y] cluster_name
```

This command imports the configuration of an autonomous or “wild” cluster into the cluster named `cluster_name`.

`import config` requires a single argument, `cluster_name`, which is the name of the cluster created using MySQL Cluster Manager into which you wish to import the configuration of a MySQL NDB Cluster created externally. The cluster named in the command must already exist in MySQL Cluster Manager; you are also strongly advised to use `create cluster --import` when creating `cluster_name`.

`import config` also supports a `--dryrun` option (short form: `-y`). When this option is used, the checks required for importing the configuration data are performed, and the `set` commands

for performing the actual import are written to the file `/path-to-mcm-data-repository/clusters/clustername/tmp/import_config.message_id.mcm` for your examination. This makes it possible to test the configuration import without actually copying any of the settings into the cluster controlled by MySQL Cluster Manager. You can then import all the settings using the `import config` command (without the `--dryrun` option), or adjust some of the settings in the `/path-to-mcm-data-repository/clusters/clustername/tmp/import_config.message_id.mcm` file and then import the settings by executing the file with the `mcm` agent. See [Section 3.5.2.3, “Creating and Configuring the Target Cluster”](#) for examples on using the `import config` command.



Note

When importing configuration attributes for a `mysqld` node, if a relative path is used for the `socket` value or for any directory value (for example, `plugin_dir`), the import will be rejected by the `mcm` client. Make sure an absolute path is used in those cases and, if necessary, make adjustments to the attributes in the `.mcm` file produced by the `import config --dryrun` command and then import the settings by executing the file with the `mcm` client.

Chapter 5 MySQL Cluster Manager Limitations and Known Issues

Table of Contents

5.1 MySQL Cluster Manager Usage and Design Limitations	149
5.2 MySQL Cluster Manager 1.4.8 Limitations Relating to the MySQL Server	149
5.3 MySQL Cluster Manager Limitations Relating to MySQL NDB Cluster	150
5.4 Syntax and Related Issues in MySQL Cluster Manager	151

In this chapter we discuss limitations of and known issues in MySQL Cluster Manager version 1.4.8.

5.1 MySQL Cluster Manager Usage and Design Limitations

The limitations discussed in this section occur by intention or design in MySQL Cluster Manager 1.4.8. Some of these items may become obsolete in future versions; we will update this section accordingly if and as those changes come about.

change process command. Currently, the `change process` command can be used only to exchange an `ndbd` process for an `ndbmta` process, or the reverse. That is, in effect, it can be used only to switch a data node between a single-threaded process and a multithreaded process. It cannot be used for changing a cluster node's type (for example, you cannot change a data node to an SQL node, management node, or NDB API application node).

Because of this restriction, and due to the fact that multithreaded data nodes are not supported in MySQL NDB Cluster 6.3 and earlier, the `change process` command is useful only with clusters running MySQL NDB Cluster 7.0 or newer.

License keys and operating platforms. License keys are not required to use MySQL Cluster Manager 1.4.8.

Concurrent client sessions unsupported. Currently there is no negotiation or arbitration between multiple `mcm` clients. While it is possible to use the client from multiple locations, we do not support *concurrent* client sessions. You should be careful always to allow a command issued in one `mcm` client session to finish executing before issuing a new command in a different client session. This is true especially when using the `-B` or `--background` option with `mcm` client commands; see [Backgrounded commands and error handling](#).

IPv6 and host names (Windows). When IPv6 support is enabled on Windows systems, host names other than `localhost` are resolved using IPv6. When an IPv6-enabled Windows system is used as a MySQL NDB Cluster host under MySQL Cluster Manager, you must reference it using its IPv4 address. Otherwise, `mcm` will be unable to connect to the agent process on that host.

This applies to host names used with the MySQL Cluster Manager client commands `create cluster`, `create site`, `add hosts`, `add package`, `delete package`, `stop agents`, and `add process`.

Use of antivirus software on Windows platforms. On-access scanning by antivirus software on Windows platforms might cause access to the cluster configuration file being denied for the `mcmd` agent, causing updates for the cluster configuration to fail sometimes.

5.2 MySQL Cluster Manager 1.4.8 Limitations Relating to the MySQL Server

The limitations described in this section relate to functionality in the MySQL Server that is unsupported or reduced, or otherwise differs when using it with MySQL Cluster Manager.

Replication. Replication is currently not directly supported by MySQL Cluster Manager. See [Section 5.3, “MySQL Cluster Manager Limitations Relating to MySQL NDB Cluster”](#), for more information.

Limited `mysqld` option modifier support. MySQL Cluster Manager does not recognize the `--loose`, `--maximum`, `--enable`, and `--disable` prefixes for `mysqld` options used as MySQL Cluster Manager configuration attributes (for a description of these modifiers, see [Program Option Modifiers](#)). For example, the command `set loose-skip-innodb:mysqld=true mycluster;` fails with the error `No such config variable loose-skip-innodb for process mysqld`.

The `--skip` option modifier is supported in some but not all cases, so that commands such as `set skip-innodb:mysqld=true mycluster;` and `set skip-grant-tables:mysqld=true mycluster;` can be used with MySQL Cluster Manager, while `set skip-column-names:mysqld=true mycluster;` cannot. (Bug #48559, Bug #47779)

Dashes and underscores in MySQL option and variable names. When using the `mysql` client or other MySQL client applications, many MySQL system options and variables can be named using either dashes or underscores in their names. For example, you can use either `ndb_batch_size` or `ndb-batch-size` with the MySQL Server, and the variable is set correctly. This is not the case in MySQL Cluster Manager, where only the forms using underscores are accepted as attribute names. For example, assuming that `mycluster` is a viable cluster, the command `set ndb_batch_size:mysqld=65536 mycluster;` works to set the size of `ndb_batch_size` on all `mysqld` processes in the cluster, but `set ndb-batch-size:mysqld=65536 mycluster;` fails.

Dependencies between MySQL Cluster Manager `mysqld` attributes and MySQL server variables. MySQL Cluster Manager does not track dependencies between `mysqld` attributes (MySQL server options and system variables). That means MySQL Cluster Manager might have `mysqld` started successfully and report so, even though the server has ended up in a non-functional state because dependent attributes were set inconsistently. It is therefore a good idea for users to check the `mysqld` attributes before starting the node and the `mysql` log for status of the node after it has been started.

MySQL Cluster Manager `mysqld` attributes and MySQL user variables. MySQL user variables are not accessible as MySQL Cluster Manager configuration attributes.

5.3 MySQL Cluster Manager Limitations Relating to MySQL NDB Cluster

This section describes limitations relating to MySQL NDB Cluster functionality that is unsupported or curtailed by MySQL Cluster Manager 1.4.8.

MySQL Cluster Manager and replication. MySQL Cluster Manager currently does not provide any explicit support for MySQL NDB Cluster Replication. However, you should still be able to perform manual setup of replication of a MySQL NDB Cluster that is managed by MySQL Cluster Manager.

Backup and restore operations. MySQL Cluster Manager provides integrated backup and restore functionality. You can back up [NDB](#) databases and tables using the `mcm` client `backup cluster` command, and restore them using the `restore cluster` client command. MySQL Cluster Manager also supports restoration of distributed privileges.

You can also back up [NDB](#) databases and tables using the `ndb_mgm` client `START BACKUP` command, and restore them using the `ndb_restore` program; however MySQL Cluster Manager is not aware of backups that it was not employed to create. Both of the programs just mentioned are supplied with the MySQL NDB Cluster distribution.



Note

Backups of tables using storage engines other than [NDB](#), as well as of all other database objects that are not tables, cannot be made using MySQL Cluster Manager, and must be made using some other method, such as `mysqldump`.

Rolling restarts. Currently, all cluster nodes must be running in order to perform a rolling restart using MySQL Cluster Manager. However, MySQL NDB Cluster itself requires only that at least one management server and all data nodes are running (in other words, any `mysqld` processes and any additional `ndb_mgmd` processes can be stopped). In such cases, you can perform the rolling restart manually, after stopping the MySQL Cluster Manager agent.

When making changes in configuration attributes only those nodes requiring a restart to make the change take effect are actually restarted. `ndbapi` nodes are never restarted by MySQL Cluster Manager.

Cluster Imports.

MySQL Cluster Manager will reject an import if it cannot access the process information of the cluster being imported. Therefore, the MySQL Cluster Manager agents must be run by a sufficiently privileged user—normally the same user that runs the cluster.

Cluster Reconfiguration. Cluster configuration updates (using the `set` or `reset` command) that would trigger a rolling restart of the nodes are not executed by MySQL Cluster Manager unless there are more than one data node defined for each node group; to perform such updates when the requirement is not met, a user should, using MySQL Cluster Manager, stop the cluster, use the `set` or `reset` command to change the cluster configurations, and then start the cluster again. If, however, your configuration changes require an `initial restart` of your cluster, you will need to backup the data, recreate your cluster from scratch with the new settings, and then restore your old data onto it.

5.4 Syntax and Related Issues in MySQL Cluster Manager

This section covers MySQL Cluster Manager issues relating to limitations in SQL and other syntax.

Backgrounded commands and error handling.

MySQL Cluster Manager client commands that are run in the background (that is, when they are invoked with the `--background` or `-B` option) do not issue any error messages in the client. If you run a command in the background, you must check the agent log or use the `show status` command to verify whether the backgrounded command was successful.

Appendix A Attribute Summary Tables

Table of Contents

A.1 Management Node Configuration Parameters	153
A.2 Data Node Configuration Parameters	154
A.3 API Node Configuration Parameters	160
A.4 Other Node Configuration Parameters	160
A.5 MySQL Server Option and Variable Reference for MySQL Cluster	162

This appendix provides tables of configuration attributes, grouped according to their process type or by the section of the MySQL NDB Cluster configuration file in which they appear. This information is current for MySQL NDB Cluster 7.5 and 7.6.

Each table provides the following information:

- *Name*: The name of the attribute. The name of the attribute is linked to the attribute's full description in the online MySQL NDB Cluster documentation.
- *Type/Units*: The data type or unit by which the attribute is measured.
- *Range*: The default value of the attribute, if not set by the user, and the minimum and maximum values that can be set for the attribute.
- *Restart Type*: The type of restart required for a change in value in this attribute to be applied in a running MySQL NDB Cluster. The restart type is indicated in this column by an **N** for a node restart, or an **S** for a system restart. *Data node attributes*: The presence of an **I** in this column indicates that a data node must be restarted using the `--initial` option for a change to take effect.

Attributes having restart type **N** can be changed using a rolling restart of the cluster, and thus can be changed at any time, even if the cluster is running. Changing an attribute whose restart type is **S** requires a complete shutdown of all cluster nodes, followed by a restart of the nodes once all of them have been stopped. Currently, such attributes can be set only before starting a cluster for the first time.

A.1 Management Node Configuration Parameters

- [ArbitrationDelay](#): When asked to arbitrate, arbitrator waits this long before voting (milliseconds).
- [ArbitrationRank](#): If 0, then management node is not arbitrator. Kernel selects arbitrators in order 1, 2.
- [DataDir](#): Data directory for this node.
- [ExecuteOnComputer](#): String referencing earlier defined COMPUTER.
- [ExtraSendBufferMemory](#): Memory to use for send buffers in addition to any allocated by TotalSendBufferMemory or SendBufferMemory. Default (0) allows up to 16MB.
- [HeartbeatIntervalMgmdMgmd](#): Time between management-node-to-management-node heartbeats; connection between management nodes is considered lost after 3 missed heartbeats.
- [HeartbeatThreadPriority](#): Set heartbeat thread policy and priority for management nodes; see manual for allowed values.
- [HostName](#): Host name or IP address for this management node.
- [Id](#): Number identifying management node. Now deprecated; use NodeId instead.

- [LogDestination](#): Where to send log messages: console, system log, or specified log file.
- [NodeId](#): Number uniquely identifying management node among all nodes in cluster.
- [PortNumber](#): Port number to send commands to and fetch configuration from management server.
- [PortNumberStats](#): Port number used to get statistical information from management server.
- [TotalSendBufferMemory](#): Total memory to use for all transporter send buffers.
- [wan](#): Use WAN TCP setting as default.

A.2 Data Node Configuration Parameters

- [Arbitration](#): How arbitration should be performed to avoid split-brain issues in event of node failure.
- [ArbitrationTimeout](#): Maximum time (milliseconds) database partition waits for arbitration signal.
- [BackupDataBufferSize](#): Default size of databuffer for backup (in bytes).
- [BackupDataDir](#): Path to where to store backups. Note that string '/BACKUP' is always appended to this setting, so that *effective* default is FileSystemPath/BACKUP.
- [BackupDiskWriteSpeedPct](#): Sets percentage of data node's allocated maximum write speed (MaxDiskWriteSpeed) to reserve for LCPs when starting backup.
- [BackupLogBufferSize](#): Default size of log buffer for backup (in bytes).
- [BackupMaxWriteSize](#): Maximum size of file system writes made by backup (in bytes).
- [BackupMemory](#): Total memory allocated for backups per node (in bytes).
- [BackupReportFrequency](#): Frequency of backup status reports during backup in seconds.
- [BackupWriteSize](#): Default size of file system writes made by backup (in bytes).
- [BatchSizePerLocalScan](#): Used to calculate number of lock records for scan with hold lock.
- [BuildIndexThreads](#): Number of threads to use for building ordered indexes during system or node restart. Also applies when running `ndb_restore --rebuild-indexes`. Setting this parameter to 0 disables multithreaded building of ordered indexes.
- [CompressedBackup](#): Use zlib to compress backups as they are written.
- [CompressedLCP](#): Write compressed LCPs using zlib.
- [ConnectCheckIntervalDelay](#): Time between data node connectivity check stages. Data node is considered suspect after 1 interval and dead after 2 intervals with no response.
- [CrashOnCorruptedTuple](#): When enabled, forces node to shut down whenever it detects corrupted tuple.
- [DataDir](#): Data directory for this node.
- [DataMemory](#): Number of bytes on each data node allocated for storing data; subject to available system RAM and size of IndexMemory.
- [DefaultHashMapSize](#): Set size (in buckets) to use for table hash maps. Three values are supported: 0, 240, and 3840.
- [DictTrace](#): Enable DBDICT debugging; for NDB development.
- [DiskIOThreadPool](#): Number of unbound threads for file access, applies to disk data only.

- [Diskless](#): Run without using disk.
- [DiskPageBufferEntries](#): Memory to allocate in DiskPageBufferMemory; very large disk transactions may require increasing this value.
- [DiskPageBufferMemory](#): Number of bytes on each data node allocated for disk page buffer cache.
- [DiskSyncSize](#): Amount of data written to file before synch is forced.
- [EnablePartialLcp](#): Enable partial LCP (true); if this is disabled (false), all LCPs write full checkpoints.
- [EnableRedoControl](#): Enable adaptive checkpointing speed for controlling redo log usage.
- [EventLogBufferSize](#): Size of circular buffer for NDB log events within data nodes.
- [ExecuteOnComputer](#): String referencing earlier defined COMPUTER.
- [ExtraSendBufferMemory](#): Memory to use for send buffers in addition to any allocated by TotalSendBufferMemory or SendBufferMemory. Default (0) allows up to 16MB.
- [FileSystemPath](#): Path to directory where data node stores its data (directory must exist).
- [FileSystemPathDataFiles](#): Path to directory where data node stores its Disk Data files. Default value is FileSystemPathDD, if set; otherwise, FileSystemPath is used if it is set; otherwise, value of DataDir is used.
- [FileSystemPathDD](#): Path to directory where data node stores its Disk Data and undo files. Default value is FileSystemPath, if set; otherwise, value of DataDir is used.
- [FileSystemPathUndoFiles](#): Path to directory where data node stores its undo files for Disk Data. Default value is FileSystemPathDD, if set; otherwise, FileSystemPath is used if it is set; otherwise, value of DataDir is used.
- [FragmentLogFileSize](#): Size of each redo log file.
- [HeartbeatIntervalDbApi](#): Time between API node-data node heartbeats. (API connection closed after 3 missed heartbeats).
- [HeartbeatIntervalDbDb](#): Time between data node-to-data node heartbeats; data node considered dead after 3 missed heartbeats.
- [HeartbeatOrder](#): Sets order in which data nodes check each others' heartbeats for determining whether given node is still active and connected to cluster. Must be zero for all data nodes or distinct nonzero values for all data nodes; see documentation for further guidance.
- [HostName](#): Host name or IP address for this data node.
- [IndexMemory](#): Number of bytes on each data node allocated for storing indexes; subject to available system RAM and size of DataMemory.
- [IndexStatAutoCreate](#): Enable/disable automatic statistics collection when indexes are created.
- [IndexStatAutoUpdate](#): Monitor indexes for changes and trigger automatic statistics updates.
- [IndexStatSaveScale](#): Scaling factor used in determining size of stored index statistics.
- [IndexStatSaveSize](#): Maximum size in bytes for saved statistics per index.
- [IndexStatTriggerPct](#): Threshold percent change in DML operations for index statistics updates. Value is scaled down by IndexStatTriggerScale.

- [IndexStatTriggerScale](#): Scale down IndexStatTriggerPct by this amount, multiplied by base 2 logarithm of index size, for large index. Set to 0 to disable scaling.
- [IndexStatUpdateDelay](#): Minimum delay between automatic index statistics updates for given index. 0 means no delay.
- [InitFragmentLogFiles](#): Initialize fragment log files, using sparse or full format.
- [InitialLogFileGroup](#): Describes log file group that is created during initial start. See documentation for format.
- [InitialNoOfOpenFiles](#): Initial number of files open per data node. (One thread is created per file).
- [InitialTablespace](#): Describes tablespace that is created during initial start. See documentation for format.
- [InsertRecoveryWork](#): Percentage of RecoveryWork used for inserted rows; has no effect unless partial local checkpoints are in use.
- [LateAlloc](#): Allocate memory after connection to management server has been established.
- [LcpScanProgressTimeout](#): Maximum time that local checkpoint fragment scan can be stalled before node is shut down to ensure systemwide LCP progress. Use 0 to disable.
- [LockExecuteThreadToCPU](#): Comma-delimited list of CPU IDs.
- [LockMaintThreadsToCPU](#): CPU ID indicating which CPU runs maintenance threads.
- [LockPagesInMainMemory](#): 0=disable locking, 1=lock after memory allocation, 2=lock before memory allocation.
- [LogLevelCheckpoint](#): Log level of local and global checkpoint information printed to stdout.
- [LogLevelCongestion](#): Level of congestion information printed to stdout.
- [LogLevelConnection](#): Level of node connect/disconnect information printed to stdout.
- [LogLevelError](#): Transporter, heartbeat errors printed to stdout.
- [LogLevelInfo](#): Heartbeat and log information printed to stdout.
- [LogLevelNodeRestart](#): Level of node restart and node failure information printed to stdout.
- [LogLevelShutdown](#): Level of node shutdown information printed to stdout.
- [LogLevelStartup](#): Level of node startup information printed to stdout.
- [LogLevelStatistic](#): Level of transaction, operation, and transporter information printed to stdout.
- [LongMessageBuffer](#): Number of bytes allocated on each data node for internal long messages.
- [MaxAllocate](#): No longer used; has no effect.
- [MaxBufferedEpochs](#): Allowed number of epochs that subscribing node can lag behind (unprocessed epochs). Exceeding causes lagging subscribers to be disconnected.
- [MaxBufferedEpochBytes](#): Total number of bytes allocated for buffering epochs.
- [MaxDiskWriteSpeed](#): Maximum number of bytes per second that can be written by LCP and backup when no restarts are ongoing.
- [MaxDiskWriteSpeedOtherNodeRestart](#): Maximum number of bytes per second that can be written by LCP and backup when another node is restarting.

- [MaxDiskWriteSpeedOwnRestart](#): Maximum number of bytes per second that can be written by LCP and backup when this node is restarting.
- [MaxFKBuildBatchSize](#): Maximum scan batch size to use for building foreign keys. Increasing this value may speed up builds of foreign keys but impacts ongoing traffic as well.
- [MaxDMLOperationsPerTransaction](#): Limit size of transaction; aborts transaction if it requires more than this many DML operations.
- [MaxLCPStartDelay](#): Time in seconds that LCP polls for checkpoint mutex (to allow other data nodes to complete metadata synchronization), before putting itself in lock queue for parallel recovery of table data.
- [MaxNoOfAttributes](#): Suggests total number of attributes stored in database (sum over all tables).
- [MaxNoOfConcurrentIndexOperations](#): Total number of index operations that can execute simultaneously on one data node.
- [MaxNoOfConcurrentOperations](#): Maximum number of operation records in transaction coordinator.
- [MaxNoOfConcurrentScans](#): Maximum number of scans executing concurrently on data node.
- [MaxNoOfConcurrentSubOperations](#): Maximum number of concurrent subscriber operations.
- [MaxNoOfConcurrentTransactions](#): Maximum number of transactions executing concurrently on this data node, total number of transactions that can be executed concurrently is this value times number of data nodes in cluster.
- [MaxNoOfFiredTriggers](#): Total number of triggers that can fire simultaneously on one data node.
- [MaxNoOfLocalOperations](#): Maximum number of operation records defined on this data node.
- [MaxNoOfLocalScans](#): Maximum number of fragment scans in parallel on this data node.
- [MaxNoOfOpenFiles](#): Maximum number of files open per data node.(One thread is created per file).
- [MaxNoOfOrderedIndexes](#): Total number of ordered indexes that can be defined in system.
- [MaxNoOfSavedMessages](#): Maximum number of error messages to write in error log and maximum number of trace files to retain.
- [MaxNoOfSubscribers](#): Maximum number of subscribers.
- [MaxNoOfSubscriptions](#): Maximum number of subscriptions (default 0 = MaxNoOfTables).
- [MaxNoOfTables](#): Suggests total number of NDB tables stored in database.
- [MaxNoOfTriggers](#): Total number of triggers that can be defined in system.
- [MaxNoOfUniqueHashIndexes](#): Total number of unique hash indexes that can be defined in system.
- [MaxParallelCopyInstances](#): Number of parallel copies during node restarts. Default is 0, which uses number of LDMS on both nodes, to maximum of 16.
- [MaxParallelScansPerFragment](#): Maximum number of parallel scans per fragment. Once this limit is reached, scans are serialized.
- [MaxReorgBuildBatchSize](#): Maximum scan batch size to use for reorganization of table partitions. Increasing this value may speed up table partition reorganization but impacts ongoing traffic as well.
- [MaxStartFailRetries](#): Maximum retries when data node fails on startup, requires StopOnError = 0. Setting to 0 causes start attempts to continue indefinitely.

- **MaxUIBuildBatchSize**: Maximum scan batch size to use for building unique keys. Increasing this value may speed up builds of unique keys but impacts ongoing traffic as well.
- **MemReportFrequency**: Frequency of memory reports in seconds; 0 = report only when exceeding percentage limits.
- **MinDiskWriteSpeed**: Minimum number of bytes per second that can be written by LCP and backup.
- **MinFreePct**: Percentage of memory resources to keep in reserve for restarts.
- **NodeGroup**: Node group to which data node belongs; used only during initial start of cluster.
- **NodeId**: Number uniquely identifying data node among all nodes in cluster.
- **NoOfFragmentLogFiles**: Number of 16 MB redo log files in each of 4 file sets belonging to data node.
- **NoOfReplicas**: Number of copies of all data in database.
- **Numa**: (Linux only; requires libnuma) Controls NUMA support. Setting to 0 permits system to determine use of interleaving by data node process; 1 means that it is determined by data node.
- **ODirect**: Use O_DIRECT file reads and writes when possible.
- **ODirectSyncFlag**: O_DIRECT writes are treated as synchronized writes; ignored when ODirect is not enabled, InitFragmentLogFiles is set to SPARSE, or both.
- **RealtimeScheduler**: When true, data node threads are scheduled as real-time threads. Default is false.
- **RecoveryWork**: Percentage of storage overhead for LCP files: greater value means less work in normal operations, more work during recovery.
- **RedoBuffer**: Number of bytes on each data node allocated for writing redo logs.
- **RedoOverCommitCounter**: When RedoOverCommitLimit has been exceeded this many times, transactions are aborted, and operations are handled as specified by DefaultOperationRedoProblemAction.
- **RedoOverCommitLimit**: Each time that flushing current redo buffer takes longer than this many seconds, number of times that this has happened is compared to RedoOverCommitCounter.
- **ReservedSendBufferMemory**: This parameter is present in NDB code but is not enabled.
- **RestartOnErrorInsert**: Control type of restart caused by inserting error (when StopOnError is enabled).
- **SchedulerExecutionTimer**: Number of microseconds to execute in scheduler before sending.
- **SchedulerResponsiveness**: Set NDB scheduler response optimization 0-10; higher values provide better response time but lower throughput.
- **SchedulerSpinTimer**: Number of microseconds to execute in scheduler before sleeping.
- **ServerPort**: Port used to set up transporter for incoming connections from API nodes.
- **SharedGlobalMemory**: Total number of bytes on each data node allocated for any use.
- **StartFailRetryDelay**: Delay in seconds after start failure prior to retry; requires StopOnError = 0.
- **StartFailureTimeout**: Milliseconds to wait before terminating. (0=Wait forever).

- [StartNoNodeGroupTimeout](#): Time to wait for nodes without nodegroup before trying to start (0=forever).
- [StartPartialTimeout](#): Milliseconds to wait before trying to start without all nodes. (0=Wait forever).
- [StartPartitionedTimeout](#): Milliseconds to wait before trying to start partitioned. (0=Wait forever).
- [StartupStatusReportFrequency](#): Frequency of status reports during startup.
- [StopOnError](#): When set to 0, data node automatically restarts and recovers following node failures.
- [StringMemory](#): Default size of string memory (0 to 100 = % of maximum, 101+ = actual bytes).
- [TcpBind_INADDR_ANY](#): Bind IP_ADDR_ANY so that connections can be made from anywhere (for autogenerated connections).
- [TimeBetweenEpochs](#): Time between epochs (synchronization used for replication).
- [TimeBetweenEpochsTimeout](#): Timeout for time between epochs. Exceeding causes node shutdown.
- [TimeBetweenGlobalCheckpoints](#): Time between group commits of transactions to disk.
- [TimeBetweenGlobalCheckpointsTimeout](#): Minimum timeout for group commit of transactions to disk.
- [TimeBetweenInactiveTransactionAbortCheck](#): Time between checks for inactive transactions.
- [TimeBetweenLocalCheckpoints](#): Time between taking snapshots of database (expressed in base-2 logarithm of bytes).
- [TimeBetweenWatchDogCheck](#): Time between execution checks inside data node.
- [TimeBetweenWatchDogCheckInitial](#): Time between execution checks inside data node (early start phases when memory is allocated).
- [TotalSendBufferMemory](#): Total memory to use for all transporter send buffers..
- [TransactionBufferMemory](#): Dynamic buffer space (in bytes) for key and attribute data allocated for each data node.
- [TransactionDeadlockDetectionTimeout](#): Time transaction can spend executing within data node. This is time that transaction coordinator waits for each data node participating in transaction to execute request. If data node takes more than this amount of time, transaction is aborted.
- [TransactionInactiveTimeout](#): Milliseconds that application waits before executing another part of transaction. This is time transaction coordinator waits for application to execute or send another part (query, statement) of transaction. If application takes too much time, then transaction is aborted. Timeout = 0 means that application never times out.
- [TwoPassInitialNodeRestartCopy](#): Copy data in 2 passes during initial node restart, which enables multithreaded building of ordered indexes for such restarts.
- [UndoDataBuffer](#): Unused; has no effect.
- [UndoIndexBuffer](#): Unused; has no effect.
- [UseShm](#): Use shared memory connections between this data node and API node also running on this host.

- [MaxNoOfExecutionThreads](#): For ndbmysd only, specify maximum number of execution threads.
- [NoOfFragmentLogParts](#): Number of redo log file groups belonging to this data node.
- [ThreadConfig](#): Used for configuration of multithreaded data nodes (ndbmysd). Default is empty string; see documentation for syntax and other information.

A.3 API Node Configuration Parameters

- [ApiVerbose](#): Enable NDB API debugging; for NDB development.
- [ArbitrationDelay](#): When asked to arbitrate, arbitrator waits this many milliseconds before voting.
- [ArbitrationRank](#): If 0, then API node is not arbitrator. Kernel selects arbitrators in order 1, 2.
- [AutoReconnect](#): Specifies whether an API node should reconnect fully when disconnected from cluster.
- [BatchByteSize](#): Default batch size in bytes.
- [BatchSize](#): Default batch size in number of records.
- [ConnectBackoffMaxTime](#): Specifies longest time in milliseconds (~100ms resolution) to allow between connection attempts to any given data node by this API node. Excludes time elapsed while connection attempts are ongoing, which in worst case can take several seconds. Disable by setting to 0. If no data nodes are currently connected to this API node, [StartConnectBackoffMaxTime](#) is used instead.
- [ConnectionMap](#): Specifies which data nodes to connect.
- [DefaultHashMapSize](#): Set size (in buckets) to use for table hash maps. Three values are supported: 0, 240, and 3840.
- [DefaultOperationRedoProblemAction](#): How operations are handled in event that [RedoOverCommitCounter](#) is exceeded.
- [ExecuteOnComputer](#): String referencing earlier defined COMPUTER.
- [ExtraSendBufferMemory](#): Memory to use for send buffers in addition to any allocated by [TotalSendBufferMemory](#) or [SendBufferMemory](#). Default (0) allows up to 16MB.
- [HeartbeatThreadPriority](#): Set heartbeat thread policy and priority for API nodes; see manual for allowed values.
- [HostName](#): Host name or IP address for this SQL or API node.
- [Id](#): Number identifying MySQL server or API node (Id). Now deprecated; use [NodeId](#) instead.
- [MaxScanBatchSize](#): Maximum collective batch size for one scan.
- [NodeId](#): Number uniquely identifying SQL node or API node among all nodes in cluster.
- [StartConnectBackoffMaxTime](#): Same as [ConnectBackoffMaxTime](#) except that this parameter is used in its place if no data nodes are connected to this API node.
- [TotalSendBufferMemory](#): Total memory to use for all transporter send buffers.
- [wan](#): Use WAN TCP setting as default.

A.4 Other Node Configuration Parameters

- [HostName](#): Host name or IP address of this computer.

- [Id](#): Unique identifier for this computer.
- [Checksum](#): If checksum is enabled, all signals between nodes are checked for errors.
- [Group](#): Used for group proximity; smaller value is interpreted as being closer.
- [HostName1](#): Name or IP address of first of two computers joined by TCP connection.
- [HostName2](#): Name or IP address of second of two computers joined by TCP connection.
- [NodeId1](#): ID of node (data node, API node, or management node) on one side of connection.
- [NodeId2](#): ID of node (data node, API node, or management node) on one side of connection.
- [NodeIdServer](#): Set server side of TCP connection.
- [OverloadLimit](#): When more than this many unsent bytes are in send buffer, connection is considered overloaded.
- [PortNumber](#): Port used for TCP transporter.
- [PreSendChecksum](#): If this parameter and Checksum are both enabled, perform pre-send checksum checks, and check all TCP signals between nodes for errors.
- [Proxy](#):
- [ReceiveBufferMemory](#): Bytes of buffer for signals received by this node.
- [SendBufferMemory](#): Bytes of TCP buffer for signals sent from this node.
- [SendSignalId](#): Sends ID in each signal. Used in trace files. Defaults to true in debug builds.
- [TCP_MAXSEG_SIZE](#): Value used for TCP_MAXSEG.
- [TCP_RCV_BUF_SIZE](#): Value used for SO_RCVBUF.
- [TCP_SND_BUF_SIZE](#): Value used for SO_SNDBUF.
- [TcpBind_INADDR_ANY](#): Bind InAddrAny instead of host name for server part of connection.
- [Checksum](#): If checksum is enabled, all signals between nodes are checked for errors.
- [Group](#): Used for group proximity; smaller value is interpreted as being closer.
- [HostName1](#): Name or IP address of first of two computers joined by SHM connection.
- [HostName2](#): Name or IP address of second of two computers joined by SHM connection.
- [NodeId1](#): ID of node (data node, API node, or management node) on one side of connection.
- [NodeId2](#): ID of node (data node, API node, or management node) on one side of connection.
- [NodeIdServer](#): Set server side of SHM connection.
- [OverloadLimit](#): When more than this many unsent bytes are in send buffer, connection is considered overloaded.
- [PortNumber](#): Port used for SHM transporter.
- [PreSendChecksum](#): If this parameter and Checksum are both enabled, perform pre-send checksum checks, and check all SHM signals between nodes for errors.
- [SendBufferMemory](#): Bytes in shared memory buffer for signals sent from this node.

- [SendSignalId](#): Sends ID in each signal. Used in trace files.
- [ShmKey](#): Shared memory key; when set to 1, this is calculated by NDB.
- [ShmSpinTime](#): When receiving, number of microseconds to spin before sleeping.
- [ShmSize](#): Size of shared memory segment.
- [Signum](#): Signal number to be used for signalling.

A.5 MySQL Server Option and Variable Reference for MySQL Cluster

- [Com_show_ndb_status](#): Count of SHOW NDB STATUS statements.
- [Handler_discover](#): Number of times that tables have been discovered.
- [ndb-batch-size](#): Size (in bytes) to use for NDB transaction batches.
- [ndb-blob-read-batch-bytes](#): Specifies size in bytes that large BLOB reads should be batched into. 0 = no limit.
- [ndb-blob-write-batch-bytes](#): Specifies size in bytes that large BLOB writes should be batched into. 0 = no limit.
- [ndb-cluster-connection-pool](#): Number of connections to cluster used by MySQL.
- [ndb-cluster-connection-pool-nodeids](#): Comma-separated list of node IDs for connections to cluster used by MySQL; number of nodes in list must match value set for --ndb-cluster-connection-pool.
- [ndb-connectstring](#): Address of NDB management server distributing configuration information for this cluster.
- [ndb-default-column-format](#): Use this value (FIXED or DYNAMIC) by default for COLUMN_FORMAT and ROW_FORMAT options when creating or adding table columns.
- [ndb-deferred-constraints](#): Specifies that constraint checks on unique indexes (where these are supported) should be deferred until commit time. Not normally needed or used; for testing purposes only.
- [ndb-distribution](#): Default distribution for new tables in NDBCLUSTER (KEYHASH or LINHASH, default is KEYHASH).
- [ndb-log-apply-status](#): Cause MySQL server acting as replica to log mysql.ndb_apply_status updates received from its immediate source in its own binary log, using its own server ID. Effective only if server is started with --ndbcluster option.
- [ndb-log-empty-epochs](#): When enabled, causes epochs in which there were no changes to be written to ndb_apply_status and ndb_binlog_index tables, even when --log-slave-updates is enabled.
- [ndb-log-empty-update](#): When enabled, causes updates that produced no changes to be written to ndb_apply_status and ndb_binlog_index tables, even when --log-slave-updates is enabled.
- [ndb-log-exclusive-reads](#): Log primary key reads with exclusive locks; allow conflict resolution based on read conflicts.
- [ndb-log-fail-terminate](#): Terminate mysqld process if complete logging of all found row events is not possible.
- [ndb-log-orig](#): Log originating server id and epoch in mysql.ndb_binlog_index table.

- [ndb-log-transaction-id](#): Write NDB transaction IDs in binary log. Requires `--log-bin-v1-events=OFF`.
- [ndb-log-update-minimal](#): Log updates in minimal format.
- [ndb-log-updated-only](#): Log complete rows (ON) or updates only (OFF).
- [ndb-log-update-as-write](#): Toggles logging of updates on source between updates (OFF) and writes (ON).
- [ndb-mgmd-host](#): Set host (and port, if desired) for connecting to management server.
- [ndb-nodeid](#): NDB Cluster node ID for this MySQL server.
- [ndb-optimized-node-selection](#): Enable optimizations for selection of nodes for transactions. Enabled by default; use `--skip-ndb-optimized-node-selection` to disable.
- [ndb-transid-mysql-connection-map](#): Enable or disable `ndb_transid_mysql_connection_map` plugin; that is, enable or disable `INFORMATION_SCHEMA` table having that name.
- [ndb-wait-connected](#): Time (in seconds) for MySQL server to wait for connection to cluster management and data nodes before accepting MySQL client connections.
- [ndb-wait-setup](#): Time (in seconds) for MySQL server to wait for NDB engine setup to complete.
- [ndb-allow-copying-alter-table](#): Set to OFF to keep ALTER TABLE from using copying operations on NDB tables.
- [Ndb_api_adaptive_send_deferred_count](#): Number of adaptive send calls not actually sent by this MySQL Server (SQL node).
- [Ndb_api_adaptive_send_deferred_count_session](#): Number of adaptive send calls not actually sent in this client session.
- [Ndb_api_adaptive_send_deferred_count_slave](#): Number of adaptive send calls not actually sent by this replica.
- [Ndb_api_adaptive_send_forced_count](#): Number of adaptive sends with forced-send set sent by this MySQL Server (SQL node).
- [Ndb_api_adaptive_send_forced_count_session](#): Number of adaptive sends with forced-send set in this client session.
- [Ndb_api_adaptive_send_forced_count_slave](#): Number of adaptive sends with forced-send set sent by this replica.
- [Ndb_api_adaptive_send_unforced_count](#): Number of adaptive sends without forced-send sent by this MySQL Server (SQL node).
- [Ndb_api_adaptive_send_unforced_count_session](#): Number of adaptive sends without forced-send in this client session.
- [Ndb_api_adaptive_send_unforced_count_slave](#): Number of adaptive sends without forced-send sent by this replica.
- [Ndb_api_bytes_received_count](#): Quantity of data (in bytes) received from data nodes by this MySQL Server (SQL node).
- [Ndb_api_bytes_received_count_session](#): Quantity of data (in bytes) received from data nodes in this client session.
- [Ndb_api_bytes_received_count_slave](#): Quantity of data (in bytes) received from data nodes by this replica.

- [Ndb_api_bytes_sent_count](#): Quantity of data (in bytes) sent to data nodes by this MySQL Server (SQL node).
- [Ndb_api_bytes_sent_count_session](#): Quantity of data (in bytes) sent to data nodes in this client session.
- [Ndb_api_bytes_sent_count_slave](#): Quantity of data (in bytes) sent to data nodes by this replica.
- [Ndb_api_event_bytes_count](#): Number of bytes of events received by this MySQL Server (SQL node).
- [Ndb_api_event_bytes_count_injector](#): Number of bytes of event data received by NDB binary log injector thread.
- [Ndb_api_event_data_count](#): Number of row change events received by this MySQL Server (SQL node).
- [Ndb_api_event_data_count_injector](#): Number of row change events received by NDB binary log injector thread.
- [Ndb_api_event_nondata_count](#): Number of events received, other than row change events, by this MySQL Server (SQL node).
- [Ndb_api_event_nondata_count_injector](#): Number of events received, other than row change events, by NDB binary log injector thread.
- [Ndb_api_pk_op_count](#): Number of operations based on or using primary keys by this MySQL Server (SQL node).
- [Ndb_api_pk_op_count_session](#): Number of operations based on or using primary keys in this client session.
- [Ndb_api_pk_op_count_slave](#): Number of operations based on or using primary keys by this replica.
- [Ndb_api_pruned_scan_count](#): Number of scans that have been pruned to one partition by this MySQL Server (SQL node).
- [Ndb_api_pruned_scan_count_session](#): Number of scans that have been pruned to one partition in this client session.
- [Ndb_api_pruned_scan_count_slave](#): Number of scans that have been pruned to one partition by this replica.
- [Ndb_api_range_scan_count](#): Number of range scans that have been started by this MySQL Server (SQL node).
- [Ndb_api_range_scan_count_session](#): Number of range scans that have been started in this client session.
- [Ndb_api_range_scan_count_slave](#): Number of range scans that have been started by this replica.
- [Ndb_api_read_row_count](#): Total number of rows that have been read by this MySQL Server (SQL node).
- [Ndb_api_read_row_count_session](#): Total number of rows that have been read in this client session.
- [Ndb_api_read_row_count_slave](#): Total number of rows that have been read by this replica.
- [Ndb_api_scan_batch_count](#): Number of batches of rows received by this MySQL Server (SQL node).

- [Ndb_api_scan_batch_count_session](#): Number of batches of rows received in this client session.
- [Ndb_api_scan_batch_count_slave](#): Number of batches of rows received by this replica.
- [Ndb_api_table_scan_count](#): Number of table scans that have been started, including scans of internal tables, by this MySQL Server (SQL node).
- [Ndb_api_table_scan_count_session](#): Number of table scans that have been started, including scans of internal tables, in this client session.
- [Ndb_api_table_scan_count_slave](#): Number of table scans that have been started, including scans of internal tables, by this replica.
- [Ndb_api_trans_abort_count](#): Number of transactions aborted by this MySQL Server (SQL node).
- [Ndb_api_trans_abort_count_session](#): Number of transactions aborted in this client session.
- [Ndb_api_trans_abort_count_slave](#): Number of transactions aborted by this replica.
- [Ndb_api_trans_close_count](#): Number of transactions aborted (may be greater than sum of TransCommitCount and TransAbortCount) by this MySQL Server (SQL node).
- [Ndb_api_trans_close_count_session](#): Number of transactions aborted (may be greater than sum of TransCommitCount and TransAbortCount) in this client session.
- [Ndb_api_trans_close_count_slave](#): Number of transactions aborted (may be greater than sum of TransCommitCount and TransAbortCount) by this replica.
- [Ndb_api_trans_commit_count](#): Number of transactions committed by this MySQL Server (SQL node).
- [Ndb_api_trans_commit_count_session](#): Number of transactions committed in this client session.
- [Ndb_api_trans_commit_count_slave](#): Number of transactions committed by this replica.
- [Ndb_api_trans_local_read_row_count](#): Total number of rows that have been read by this MySQL Server (SQL node).
- [Ndb_api_trans_local_read_row_count_session](#): Total number of rows that have been read in this client session.
- [Ndb_api_trans_local_read_row_count_slave](#): Total number of rows that have been read by this replica.
- [Ndb_api_trans_start_count](#): Number of transactions started by this MySQL Server (SQL node).
- [Ndb_api_trans_start_count_session](#): Number of transactions started in this client session.
- [Ndb_api_trans_start_count_slave](#): Number of transactions started by this replica.
- [Ndb_api_uk_op_count](#): Number of operations based on or using unique keys by this MySQL Server (SQL node).
- [Ndb_api_uk_op_count_session](#): Number of operations based on or using unique keys in this client session.
- [Ndb_api_uk_op_count_slave](#): Number of operations based on or using unique keys by this replica.

- [Ndb_api_wait_exec_complete_count](#): Number of times thread has been blocked while waiting for operation execution to complete by this MySQL Server (SQL node).
- [Ndb_api_wait_exec_complete_count_session](#): Number of times thread has been blocked while waiting for operation execution to complete in this client session.
- [Ndb_api_wait_exec_complete_count_slave](#): Number of times thread has been blocked while waiting for operation execution to complete by this replica.
- [Ndb_api_wait_meta_request_count](#): Number of times thread has been blocked waiting for metadata-based signal by this MySQL Server (SQL node).
- [Ndb_api_wait_meta_request_count_session](#): Number of times thread has been blocked waiting for metadata-based signal in this client session.
- [Ndb_api_wait_meta_request_count_slave](#): Number of times thread has been blocked waiting for metadata-based signal by this replica.
- [Ndb_api_wait_nanos_count](#): Total time (in nanoseconds) spent waiting for some type of signal from data nodes by this MySQL Server (SQL node).
- [Ndb_api_wait_nanos_count_session](#): Total time (in nanoseconds) spent waiting for some type of signal from data nodes in this client session.
- [Ndb_api_wait_nanos_count_slave](#): Total time (in nanoseconds) spent waiting for some type of signal from data nodes by this replica.
- [Ndb_api_wait_scan_result_count](#): Number of times thread has been blocked while waiting for scan-based signal by this MySQL Server (SQL node).
- [Ndb_api_wait_scan_result_count_session](#): Number of times thread has been blocked while waiting for scan-based signal in this client session.
- [Ndb_api_wait_scan_result_count_slave](#): Number of times thread has been blocked while waiting for scan-based signal by this replica.
- [ndb_autoincrement_prefetch_sz](#): NDB auto-increment prefetch size.
- [ndb_cache_check_time](#): Number of milliseconds between checks of cluster SQL nodes made by MySQL query cache.
- [ndb_clear_apply_status](#): Causes RESET SLAVE/RESET REPLICA to clear all rows from `ndb_apply_status` table; ON by default.
- [Ndb_cluster_node_id](#): Node ID of this server when acting as NDB Cluster SQL node.
- [Ndb_config_from_host](#): NDB Cluster management server host name or IP address.
- [Ndb_config_from_port](#): Port for connecting to NDB Cluster management server.
- [Ndb_conflict_fn_epoch](#): Number of rows that have been found in conflict by `NDB$EPOCH()` NDB replication conflict detection function.
- [Ndb_conflict_fn_epoch2](#): Number of rows that have been found in conflict by NDB replication `NDB$EPOCH2()` conflict detection function.
- [Ndb_conflict_fn_epoch2_trans](#): Number of rows that have been found in conflict by NDB replication `NDB$EPOCH2_TRANS()` conflict detection function.
- [Ndb_conflict_fn_epoch_trans](#): Number of rows that have been found in conflict by `NDB$EPOCH_TRANS()` conflict detection function.
- [Ndb_conflict_fn_max](#): Number of times that NDB replication conflict resolution based on "greater timestamp wins" has been applied to update and delete operations.

- [Ndb_conflict_fn_max_del_win](#): Number of times that NDB replication conflict resolution based on outcome of NDB\$MAX_DELETE_WIN() has been applied to update and delete operations.
- [Ndb_conflict_fn_old](#): Number of times in NDB replication "same timestamp wins" conflict resolution has been applied.
- [Ndb_conflict_last_conflict_epoch](#): Most recent NDB epoch on this replica in which some conflict was detected.
- [Ndb_conflict_last_stable_epoch](#): Number of rows found to be in conflict by transactional conflict function.
- [Ndb_conflict_reflected_op_discard_count](#): Number of reflected operations that were not applied due error during execution.
- [Ndb_conflict_reflected_op_prepare_count](#): Number of reflected operations received that have been prepared for execution.
- [Ndb_conflict_refresh_op_count](#): Number of refresh operations that have been prepared.
- [Ndb_conflict_trans_conflict_commit_count](#): Number of epoch transactions committed after requiring transactional conflict handling.
- [Ndb_conflict_trans_detect_iter_count](#): Number of internal iterations required to commit epoch transaction. Should be (slightly) greater than or equal to [Ndb_conflict_trans_conflict_commit_count](#).
- [Ndb_conflict_trans_reject_count](#): Number of transactions rejected after being found in conflict by transactional conflict function.
- [Ndb_conflict_trans_row_conflict_count](#): Number of rows found in conflict by transactional conflict function. Includes any rows included in or dependent on conflicting transactions.
- [Ndb_conflict_trans_row_reject_count](#): Total number of rows realigned after being found in conflict by transactional conflict function. Includes [Ndb_conflict_trans_row_conflict_count](#) and any rows included in or dependent on conflicting transactions.
- [ndb_data_node_neighbour](#): Specifies cluster data node "closest" to this MySQL Server, for transaction hinting and fully replicated tables.
- [ndb_default_column_format](#): Sets default row format and column format (FIXED or DYNAMIC) used for new NDB tables.
- [ndb_deferred_constraints](#): Specifies that constraint checks should be deferred (where these are supported). Not normally needed or used; for testing purposes only.
- [ndb_distribution](#): Default distribution for new tables in NDBCLUSTER (KEYHASH or LINHASH, default is KEYHASH).
- [Ndb_epoch_delete_delete_count](#): Number of delete-delete conflicts detected (delete operation is applied, but row does not exist).
- [ndb_eventbuffer_free_percent](#): Percentage of free memory that should be available in event buffer before resumption of buffering, after reaching limit set by [ndb_eventbuffer_max_alloc](#).
- [ndb_eventbuffer_max_alloc](#): Maximum memory that can be allocated for buffering events by NDB API. Defaults to 0 (no limit).
- [Ndb_execute_count](#): Number of round trips to NDB kernel made by operations.
- [ndb_extra_logging](#): Controls logging of NDB Cluster schema, connection, and data distribution events in MySQL error log.

- [ndb_force_send](#): Forces sending of buffers to NDB immediately, without waiting for other threads.
- [ndb_fully_replicated](#): Whether new NDB tables are fully replicated.
- [ndb_index_stat_enable](#): Use NDB index statistics in query optimization.
- [ndb_index_stat_option](#): Comma-separated list of tunable options for NDB index statistics; list should contain no spaces.
- [ndb_join_pushdown](#): Enables pushing down of joins to data nodes.
- [Ndb_last_commit_epoch_server](#): Epoch most recently committed by NDB.
- [Ndb_last_commit_epoch_session](#): Epoch most recently committed by this NDB client.
- [ndb_log_apply_status](#): Whether or not MySQL server acting as replica logs `mysql.ndb_apply_status` updates received from its immediate source in its own binary log, using its own server ID.
- [ndb_log_bin](#): Write updates to NDB tables in binary log. Effective only if binary logging is enabled with `--log-bin`.
- [ndb_log_binlog_index](#): Insert mapping between epochs and binary log positions into `ndb_binlog_index` table. Defaults to ON. Effective only if binary logging is enabled.
- [ndb_log_empty_epochs](#): When enabled, epochs in which there were no changes are written to `ndb_apply_status` and `ndb_binlog_index` tables, even when `log_replica_updates` or `log_slave_updates` is enabled.
- [ndb_log_empty_update](#): When enabled, updates which produce no changes are written to `ndb_apply_status` and `ndb_binlog_index` tables, even when `log_replica_updates` or `log_slave_updates` is enabled.
- [ndb_log_exclusive_reads](#): Log primary key reads with exclusive locks; allow conflict resolution based on read conflicts.
- [ndb_log_orig](#): Whether id and epoch of originating server are recorded in `mysql.ndb_binlog_index` table. Set using `--ndb-log-orig` option when starting `mysqld`.
- [ndb_log_transaction_id](#): Whether NDB transaction IDs are written into binary log (Read-only).
- [Ndb_number_of_data_nodes](#): Number of data nodes in this NDB cluster; set only if server participates in cluster.
- [ndb-optimization-delay](#): Number of milliseconds to wait between processing sets of rows by `OPTIMIZE TABLE` on NDB tables.
- [ndb_optimized_node_selection](#): Determines how SQL node chooses cluster data node to use as transaction coordinator.
- [Ndb_pruned_scan_count](#): Number of scans executed by NDB since cluster was last started where partition pruning could be used.
- [Ndb_pushed_queries_defined](#): Number of joins that API nodes have attempted to push down to data nodes.
- [Ndb_pushed_queries_dropped](#): Number of joins that API nodes have tried to push down, but failed.
- [Ndb_pushed_queries_executed](#): Number of joins successfully pushed down and executed on data nodes.
- [Ndb_pushed_reads](#): Number of reads executed on data nodes by pushed-down joins.

- [ndb_read_backup](#): Enable read from any replica for all NDB tables; use `NDB_TABLE=READ_BACKUP={0|1}` with `CREATE TABLE` or `ALTER TABLE` to enable or disable for individual NDB tables.
- [ndb_recv_thread_activation_threshold](#): Activation threshold when receive thread takes over polling of cluster connection (measured in concurrently active threads).
- [ndb_recv_thread_cpu_mask](#): CPU mask for locking receiver threads to specific CPUs; specified as hexadecimal. See documentation for details.
- [ndb_report_thresh_binlog_epoch_slip](#): NDB 7.5 and later: Threshold for number of epochs completely buffered, but not yet consumed by binlog injector thread which when exceeded generates `BUFFERED_EPOCHS_OVER_THRESHOLD` event buffer status message; prior to NDB 7.5: Threshold for number of epochs to lag behind before reporting binary log status.
- [ndb_report_thresh_binlog_mem_usage](#): Threshold for percentage of free memory remaining before reporting binary log status.
- [ndb_row_checksum](#): When enabled, set row checksums; enabled by default.
- [Ndb_scan_count](#): Total number of scans executed by NDB since cluster was last started.
- [ndb_show_foreign_key_mock_tables](#): Show mock tables used to support `foreign_key_checks=0`.
- [ndb_slave_conflict_role](#): Role for replica to play in conflict detection and resolution. Value is one of `PRIMARY`, `SECONDARY`, `PASS`, or `NONE` (default). Can be changed only when replication SQL thread is stopped. See documentation for further information.
- [Ndb_slave_max_replicated_epoch](#): Most recently committed NDB epoch on this replica. When this value is greater than or equal to `Ndb_conflict_last_conflict_epoch`, no conflicts have yet been detected.
- [Ndb_system_name](#): Configured cluster system name; empty if server not connected to NDB.
- [ndb_table_no_logging](#): NDB tables created when this setting is enabled are not checkpointed to disk (although table schema files are created). Setting in effect when table is created with or altered to use `NDBCLUSTER` persists for table's lifetime.
- [ndb_table_temporary](#): NDB tables are not persistent on disk: no schema files are created and tables are not logged.
- [ndb_use_copying_alter_table](#): Use copying `ALTER TABLE` operations in NDB Cluster.
- [ndb_use_exact_count](#): Forces NDB to use a count of records during `SELECT COUNT(*)` query planning to speed up this type of query.
- [ndb_use_transactions](#): Set to `OFF`, to disable transaction support by NDB. Not recommended except in certain special cases; see documentation for details.
- [ndb_version](#): Shows build and NDB engine version as an integer.
- [ndb_version_string](#): Shows build information including NDB engine version in `ndb-x.y.z` format.
- [ndbcluster](#): Enable NDB Cluster (if this version of MySQL supports it). Disabled by `--skip-ndbcluster`.
- [ndbinfo_database](#): Name used for NDB information database; read only.
- [ndbinfo_max_bytes](#): Used for debugging only.
- [ndbinfo_max_rows](#): Used for debugging only.

- [ndbinfo_offline](#): Put ndbinfo database into offline mode, in which no rows are returned from tables or views.
- [ndbinfo_show_hidden](#): Whether to show ndbinfo internal base tables in mysql client; default is OFF.
- [ndbinfo_table_prefix](#): Prefix to use for naming ndbinfo internal base tables; read only.
- [ndbinfo_version](#): ndbinfo engine version; read only.
- [server_id_bits](#): Number of least significant bits in server_id actually used for identifying server, permitting NDB API applications to store application data in most significant bits. server_id must be less than 2 to power of this value.
- [skip-ndbcluster](#): Disable NDB Cluster storage engine.
- [slave_allow_batching](#): Turns update batching on and off for replica.
- [transaction_allow_batching](#): Allows batching of statements within one transaction. Disable AUTOCOMMIT to use.

Index

A

- abort backup command, 141
 - backupid option, 141
- add hosts command, 77
 - hosts option, 77
- add package command, 85, 88
 - hosts option, 85
- add process command, 132
 - sequential-restart option, 133
- added option
 - start process command, 138
- agent
 - backing up and restoring, 57, 146
 - configuration, 16
 - defined, 2, 2
 - distribution, 6, 6
 - distribution layout, 7
 - installing, 6
 - restoring, 59
 - starting, 30
 - starting (Linux), 30
 - starting (Windows), 31
 - stopping (Linux), 31
 - stopping (Windows), 33
- agent-uuid option
 - mcmd, 23
- architecture, 1, 2
- attributes
 - case-sensitivity, 67
 - summary table, 153
- autotune command, 101
 - dryrun option, 102
 - sequential-restart option, 102
 - template option, 101
 - writeload option, 101

B

- background option
 - common use, 70
 - limitations, 151
- backgrounding of commands, 151
- backup
 - commands, 141
- backup agents command, 146
 - hosts option, 146
- backup cluster command, 141
 - backupid option, 141
 - snapshotend option, 141
 - snapshotstart option, 141
 - waitcompleted option, 141
 - waitstarted option, 141
- backup images, 142
- backup option
 - show status command, 97
- backup status, 97

- backupid option
 - abort backup command, 141
 - backup cluster command, 141
 - restore cluster command, 144
- backups
 - creating, 57, 141, 146
 - removing, 142
 - restoring to a cluster with fewer data nodes, 55
- basedir option
 - add package command, 85
 - mcmd, 23
- bootstrap option
 - mcmd, 23

C

- change log-level command, 79
 - hosts option, 79
- change process command, 135
 - limitations, 149
 - sequential-restart option, 136
- changing data node processes, 135
- client
 - commands in, 66
 - defined, 2, 3
 - executing scripts with, 73
 - importing, 38
 - mysql client commands in, 72
 - starting, 33
- client commands, 65
 - abort backup, 141
 - add hosts, 77
 - add package, 85, 88
 - add process, 132
 - autotune, 101
 - backup, 141
 - backup agents, 146
 - backup cluster, 141
 - case-sensitivity in, 67
 - change log-level, 79
 - change process, 135
 - cluster, 89
 - collect logs, 80
 - command-specific, 72
 - common options, 70
 - configuration, 105
 - create cluster, 89
 - create site, 81
 - delete cluster, 93
 - delete package, 87
 - delete site, 82
 - get, 108
 - help, 77
 - identifiers in, 67
 - import cluster, 147
 - import config, 147
 - importing clusters, 147
 - information, 71
 - list backups, 143

- list clusters, 94
- list hosts, 83
- list nextnodeids, 94
- list processes, 137
- list sites, 82
- online help, 71
- options, 66, 69
- package, 85
- process, 131
- remove hosts, 78
- remove process, 140
- reset, 119
- restart cluster, 94
- restore, 141
- restore cluster, 144
- results returned by, 69
- rotate log, 79
- set, 125
- show settings, 83
- show status, 95
- show warning, 85
- site, 77
- start cluster, 99
- start process, 138
- status, 97
- stop agents, 84
- stop cluster, 101
- stop process, 138
- syntax, 66
- update process, 139
- upgrade cluster, 102
- version, 84
- cluster
 - defined, 1
- cluster backups
 - aborting, 141
 - creating, 50
 - listing, 143
 - restoring, 50
 - restoring from, 144
- cluster option
 - show status command, 96
- cluster processes
 - adding, 132
- cluster status, 96
- clusters
 - autotuning, 101
 - creating, 36, 36, 89
 - creating for import, 91
 - importing, 38, 147, 147
 - listing, 94
 - removing, 93
 - restarting, 94
 - starting, 99
 - stopping, 101
 - upgrading, 102
- collect logs command, 80
- command status, 97

- common terms, 1
- concurrent client sessions unsupported, 149
- configuration
 - derivation of attributes, 105
- configuration attributes, 105
 - command-line-only, 107
 - defined, 2
 - for TCP connections, 130
 - getting, 108
 - how determined by MySQL Cluster Manager, 106
 - levels applying, 105
 - mandatory, 106
 - read-only, 107
 - resetting, 119
 - setting, 125
- configuration commands, 105
- configuration data
 - importing, 147
- configuration file, 16
- configuration parameters (see [configuration attributes](#))
- configuration variables (see [configuration attributes](#))
- connecting to agent
 - with mcm client, 34
 - with mysql client, 35
- copy-port option
 - mcmd, 23
- create cluster command, 89
 - import option, 91
 - assignment of node IDs in, 91
- create site command, 81
 - hosts option, 81

D

- daemon option
 - mcmd, 24
- defaults-file option
 - mcmd, 24
- delete cluster command, 93
- delete package command, 87
 - hosts option, 87
- delete site command, 82
- deployment (example), 3
- disable-indexes option
 - restore cluster command, 145
- disable-metadata option
 - restore cluster command, 145
- downgrades
 - NDB Cluster downgrades involving logical downgrades of mysqld nodes, 62
- dryrun option
 - autotune command, 102
 - import cluster command, 147
 - import config command, 147

E

- epoch option
 - restore cluster command, 145

- event-threads option
 - mcmd, 24
- eventual consistency, 4
- exclude-databases option
 - restore cluster command, 145
- exclude-intermediate-sql-tables option
 - restore cluster command, 145
- exclude-missing-columns option
 - restore cluster command, 145
- exclude-missing-tables option
 - restore cluster command, 145
- exclude-tables option
 - restore cluster command, 145

F

- force option
 - common use, 70
 - upgrade cluster command, 104

G

- get command, 108
 - filtering output, 115
 - include-defaults option, 112
 - with multiple attributes, 116

H

- help command, 77
- help for commands, 72
- help option
 - common use, 71, 72
 - mcmd, 24
- help-all option
 - mcmd, 25
- help-manager option
 - mcmd, 25
- hostinfo option
 - show settings command, 84
- hosts
 - defined, 1
- hosts option
 - add hosts command, 77, 78
 - add package command, 85
 - backup agents command, 146
 - change log-level command, 79
 - create site command, 81
 - delete package command, 87
 - remove hosts command, 78
 - rotate log command, 80
 - stop agents command, 84

I

- identifiers, 67
 - case-sensitivity, 67
 - spaces in, 68
- import cluster command, 38, 40, 147
- import config command, 38, 40, 147
- import option

- create cluster command, 91
- importing clusters
 - basic procedure, 38
 - creating for import, 91
 - example, 40
 - limitations, 151
- include-databases option
 - restore cluster command, 145
- include-defaults option
 - get command, 112
- include-stored-grants option
 - restore cluster command, 145
- include-tables option
 - restore cluster command, 145
- initial option
 - mcmd, 26
 - start cluster command, 100
 - start process command, 138
- installation, 5
- IPv6
 - and hostnames (Windows), 149

K

- keepalive option
 - mcmd, 26

L

- license keys
 - and operating platform, 149
- license keys (obsolete), 6
- limitations, 149
 - background option, 151
 - backup, 150, 150
 - change process command, 149
 - client, 150, 150, 150
 - cluster imports, 151
 - concurrent usage, 149
 - IPv6, 149
 - license keys, 149
 - MySQL server variables, 150
 - MySQL user variables, 150
 - operating systems, 149
 - relating to MySQL Server, 149
 - replication, 150
 - restarts, 151
 - syntax, 151
 - Windows, 149
- list backups command, 143
- list clusters command, 94
- list commands command, 71
- list hosts command, 83
- list nextnodeids command, 94
- list processes command, 137
- list sites command, 82
- log-backtrace-on-crash option
 - mcmd, 26
- log-file option

- mcmd, 26
- log-level option
 - mcmd, 26
- log-use-syslog option
 - mcmd, 27
- logging
 - configuring, 79
- logs
 - obtaining, 80
 - rotating, 79
- lossy-conversions option
 - restore cluster command, 145

M

- management site
 - defined, 1
- management sites
 - adding hosts, 77
 - creating, 81
 - deleting, 82
 - listing, 82
 - listing hosts, 83
 - removing hosts, 78
- manager-directory option
 - mcmd, 27
- manager-password option
 - mcmd, 28
- manager-port option
 - mcmd, 28
- manager-username option
 - mcmd, 28
- max-open-files option
 - mcmd, 28
- mcm client
 - and mysql client, 33
- mcmd, 30
 - agent-uuid option, 23
 - basedir option, 23
 - bootstrap option, 23
 - copy-port option, 23
 - daemon option, 24
 - defaults-file option, 24
 - event-threads option, 24
 - help option, 24
 - help-all option, 25
 - help-manager option, 25
 - initial option, 26
 - keepalive option, 26
 - log-backtrace-on-crash option, 26
 - log-file option, 26
 - log-level option, 26
 - log-use-syslog option, 27
 - manager-directory option, 27
 - manager-password option, 28
 - manager-port option, 28
 - manager-username option, 28
 - max-open-files option, 28
 - pid-file option, 29

- plugin-dir option, 29
- plugins option, 29
- verbose-shutdown option, 29
- version option, 29
- xcom-port option, 30
- mcmd.exe, 31
- multiple client sessions, 149
- mysql client commands, 72
- mysql-cluster-manager (OBSOLETE, see mcmd), 30
- mysqld options, 150

N

- ndb_mgm (MySQL NDB Cluster command-line client)
 - using with MySQL Cluster Manager, 131
- no-binlog option
 - restore cluster command, 145
- no-restore-disk-objects option
 - restore cluster command, 146
- node IDs
 - and create cluster command, 91
 - listing, 94
- nodeid option
 - upgrade cluster command, 104

O

- obtaining MySQL Cluster Manager, 5
- operation option
 - show status command, 97

P

- package option
 - create cluster command, 90
 - upgrade cluster command, 102
- packages
 - defined, 2
 - listing, 85, 88
 - registering, 85
 - removing, 87
- parallelism option
 - restore cluster command, 146
- pid-file option
 - mcmd, 29
- plugin-dir option
 - mcmd, 29
- plugins option
 - mcmd, 29
- privilege-tables option
 - restore cluster command, 146
- process option
 - show status command, 98
- process status, 98
- processes
 - changing, 135
 - commands, 131
 - defined, 1
 - listing, 137
 - removing, 140

- starting, 138
- status, 98
- stopping, 138
- updating, 139
- processhosts option
 - add process command, 132
 - create cluster command, 90
- progress of action
 - check using the show status command, 99
- progress option
 - show status command, 99
- progress-frequency option
 - restore cluster command, 146
- progressbar option
 - show status command, 99
- promote-attributes option
 - restore cluster command, 146

R

- ReceiveBufferMemory, 125, 130
- remove hosts command, 78
 - hosts option, 78
- remove process command, 140
- remove-angel option
 - update process command, 139
- removeangel option
 - import cluster command, 147
- removedirs option
 - delete cluster command, 93
 - remove process command, 140
- replication, 150
 - setup, 59
- reset command, 119
 - and attribute name, 120
 - and TCP connections, 125
 - order of operands, 123
 - process level, 120
 - scope, 120
 - sequential-restart option, 124
- restart cluster command, 94
 - sequential-restart option, 95
- restore (from backup)
 - commands, 141
- restore cluster command, 144
 - backupid option, 144
 - disable-indexes option, 145
 - disable-metadata option, 145
 - epoch option, 145
 - exclude-databases option, 145
 - exclude-intermediate-sql-tables option, 145
 - exclude-missing-columns option, 145
 - exclude-missing-tables option, 145
 - exclude-tables option, 145
 - include-databases option, 145
 - include-stored-grants option, 145
 - include-tables option, 145
 - lossy-conversions option, 145
 - no-binlog option, 145

- no-restore-disk-objects option, 146
- parallelism option, 146
- privilege-tables option, 146
- progress-frequency option, 146
- promote-attributes option, 146
- rewrite-database option, 146
- skip-broken-objects option, 146
- skip-nodeid option, 146
- skip-table-check option, 146
- skip-unknown-objects option, 146
- retry option
 - upgrade cluster command, 104
- rewrite-database option
 - restore cluster command, 146
- rolling restarts, 151
- rotate log command, 79
 - hosts option, 80

S

- scripts (MySQL Cluster Manager client), 73
- SendBufferMemory, 125, 130
- sequential-restart option
 - add process command, 133
 - autotune command, 102
 - change process command, 136
 - reset command, 124
 - restart cluster command, 95
 - set command, 126
- set command, 125
 - and TCP connection attributes, 130
 - instance level, 126
 - paths used with (Windows), 129
 - scope, 126, 127
 - sequential-restart option, 126
 - undoing effects of, 119
 - verifying effects, 127
 - with multiple attributes, 128
 - with multiple processes, 128
- set option
 - add process command, 134
 - upgrade cluster command, 104
- show settings command, 83
- show status command, 95
 - backup option, 97
 - cluster option, 96
 - operation option, 97
 - process option, 98
 - progress option, 99
 - progressbar option, 99
- skip-broken-objects option
 - restore cluster command, 146
- skip-init option
 - start cluster command, 100
- skip-nodeid option
 - restore cluster command, 146
- skip-table-check option
 - restore cluster command, 146
- skip-unknown-objects option

- restore cluster command, 146
- snapshotend option
 - backup cluster command, 141
- snapshotstart option
 - backup cluster command, 141
- start cluster command, 99
 - initial option, 100
 - skip-init option, 100
- start process command, 138
- starting and stopping nodes
 - and ndb_mgm (MySQL NDB Cluster command-line client), 131
- stop agents command, 84
 - hosts option, 84
- stop cluster command, 101
- stop process command, 138
- syntax issues, 151

T

- template option
 - autotune command, 101
- terminology, 1

U

- update process command, 139
- upgrade cluster command, 102
- upgrades
 - MySQL NDB Cluster, 151

V

- verbose option
 - add process command, 132
 - create cluster command, 91
- verbose-shutdown option
 - mcmd, 29
- version command, 84
- version option
 - mcmd, 29

W

- waitcompleted option
 - backup cluster command, 141
- waitstarted option
 - backup cluster command, 141
- writeload option
 - autotune command, 101

X

- xcom-port option
 - mcmd, 30