MySQL Cluster

Document Status

Subject	Status	Section
Describe config.ini file parameters	Done	Cluster Configuration
Describe my.cnf file parameters	Done	Cluster Configuration
Connection string to be used	Done	Cluster Configuration
Starting and stopping the cluster	Done	Stopping and Starting the Cluster
Logging / debugging	Done	Logging and Debugging
Backing Up MySQL CLuster	Done	Backing Up MySQL Cluster
Restoring MySQL CLuster	WIP	

Presentation

MySQL Cluster is a high availability solution for MySQL which:

- uses a clustered database engine called NDBCLUSTER or simply NDB. The commands used to manipulate the NDB engine data can be found **here**,
- uses a different mysqld binary than that used by the standard MySQL,
- is compatible with Red Hat, Debian, SUSE, Oracle Linux, Mac OSX, Microsoft ™ Windows®, Solaris and generic Linux - see **this link**,
- can be downloaded from **this link**,

The cluster contains several different types of nodes also called services:

Node	Short Name	Service	Function
Cluster Management Server Node	MGM node	ndb_mgmd	Contains the configuration data, starts and stops the cluster and performs backups
MySQL Server node	SQL node or API node	mysqld	Access the cluster data via mysqld with options –ndbcluster and –ndb-connectstring

Node	Short Name	Service	Function
Data node	S/O	ndbd	contain the cluster data

It is possible to have several types of nodes (services) installed on the same physical or virtual host. In this case we refer to a **cluster host**. However it is not a good idea to have several instances of the ndbd service on the **same** host.

Each node in the cluster obtains its configuration from the Cluster Management node. This node is also responsable for logging.

When the configuration is changed a *rolling restart* is initiated where each node is restarted, one after the other.

When data is saved to disk, we refer to a **checkpoint** :

- **Local Checkpoint** (LCP): Data is saved on the Data node's disk.
- **Global Checkpoint** (GCP): Transactions are synchronised on all disks.

High Availability

In a cluster with 4 Data nodes and 1 replica or copy the number of node groups is 4. This configuration is not recommended since high availability is not achieved.

In a cluster with 4 Data nodes and 2 replicas or copies the number of node groups is 2. This configuration is recommended to ensure high availability.

Node groups are formed automatically. The two lowest Data node IDs form the first group followed by the next two. For example if there are four data nodes with IDs 1,2,3 and 4 respectively and 2 replicas, the first node group will contain data node 1 and data node 2 whilst the second group will contain date node 3 and data node 4.

<note important> Each node group must have the same number of nodes. </note>

Data is partionned into as many partitions as there are nodes.

Each node in a node group stores a **replica** (a copy of a partition). There are therefore as many replicas as there are nodes within the node group.

High availability is ensured as long as at least one node per node group is up.

Installation

Using a clean CentOS VM, create a full clone called **node1** using one processor and 512MB of RAM.

Attach a second .vmdk 2GB hard disk to the SATA controller of the **node1** VM, Call this disk **cluster**. Boot up the VM and create a partition on /dev/sdb :

```
[root@centos ~]# fdisk /dev/sdb
Device contains neither a valid DOS partition table, nor Sun, SGI or OSF disklabel
Building a new DOS disklabel with disk identifier 0xa6de8f27.
Changes will remain in memory only, until you decide to write them.
After that, of course, the previous content won't be recoverable.

Warning: invalid flag 0x0000 of partition table 4 will be corrected by w(rite)

WARNING: DOS-compatible mode is deprecated. It's strongly recommended to
         switch off the mode (command 'c') and change display units to
         sectors (command 'u').

Command (m for help): n
Command action
   e   extended
   p   primary partition (1-4)
p
Partition number (1-4): 1
First cylinder (1-261, default 1):
Using default value 1
```

```
Last cylinder, +cylinders or +size{K,M,G} (1-261, default 261):
Using default value 261

Command (m for help): w
The partition table has been altered!

Calling ioctl() to re-read partition table.
Syncing disks.
```

Now create an Ext4 filesystem on /dev/sdb1 :

```
[root@centos ~]# mkfs.ext4 /dev/sdb1
mke2fs 1.41.12 (17-May-2010)
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
Stride=0 blocks, Stripe width=0 blocks
131072 inodes, 524112 blocks
26205 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=536870912
16 block groups
32768 blocks per group, 32768 fragments per group
8192 inodes per group
Superblock backups stored on blocks:
    32768, 98304, 163840, 229376, 294912

Writing inode tables: done
Creating journal (8192 blocks): done
Writing superblocks and filesystem accounting information: done

This filesystem will be automatically checked every 37 mounts or
```

```
180 days, whichever comes first.  Use tune2fs -c or -i to override.
```

Mount /dev/sdb1 on /cluster and edit /etc/fstab accordingly :

```
[root@redhat ~]# mkdir /cluster
[root@centos ~]# mount /dev/sdb1 /cluster
[root@centos ~]# blkid /dev/sdb1
/dev/sdb1: UUID="52a48f95-1310-4321-a217-0efb3365247e" TYPE="ext4"
[root@centos ~]# vi /etc/fstab
```

```
#
# /etc/fstab
# Created by anaconda on Fri Oct 25 09:32:46 2013
#
# Accessible filesystems, by reference, are maintained under '/dev/disk'
# See man pages fstab(5), findfs(8), mount(8) and/or blkid(8) for more info
#
UUID=c7b1d3e8-6471-4cba-947b-430db974e774 /                       ext4    defaults        1 1
UUID=d8988475-7dc7-4a61-8081-6153b7c9551b /boot                   ext4    defaults        1 2
UUID=a1d6a043-6f10-4f60-bb9c-aaaac9632c57 swap                    swap    defaults        0 0
UUID=52a48f95-1310-4321-a217-0efb3365247e /cluster                ext4    defaults        0 0
tmpfs                   /dev/shm                tmpfs   defaults        0 0
devpts                  /dev/pts                devpts  gid=5,mode=620  0 0
sysfs                   /sys                    sysfs   defaults        0 0
proc                    /proc                   proc    defaults        0 0
```

Now download the tarball **mysql-cluster-gpl-7.4.4-linux-glibc2.5-i686.tar.gz** from **this URL** and move it to /cluster. Untar the tarball and rename the directory to **mysql_cluster** :

```
[root@centos ~]# mv /home/trainee/Download/mysql-cluster-gpl-7.4.4-linux-glibc2.5-i686.tar.gz /cluster
[root@centos ~]# cd /cluster
[root@centos cluster]# ls
lost+found  mysql-cluster-gpl-7.4.4-linux-glibc2.5-i686.tar.gz
[root@centos cluster]# tar xf mysql-cluster-gpl-7.4.4-linux-glibc2.5-i686.tar.gz
```

```
[root@centos cluster]# rm -rf mysql-cluster-gpl-7.4.4-linux-glibc2.5-i686.tar.gz
[root@centos cluster]# mv mysql-cluster-gpl-7.4.4-linux-glibc2.5-i686/ mysql_cluster
```

To facilitate our study of MySQL Cluster, we are going to remove anything that could cause problems. Firstly stop iptables and remove it using chkconfig :

```
[root@centos cluster]# service iptables stop
iptables: Setting chains to policy ACCEPT: filter          [  OK  ]
iptables: Flushing firewall rules:                         [  OK  ]
iptables: Unloading modules:                               [  OK  ]
[root@centos cluster]# chkconfig --del iptables
```

Then put SELinux into permissive mode by editing **/etc/selinux/config** :

```
[root@centos cluster]# vi /etc/selinux/config
[root@centos cluster]# cat /etc/selinux/config

# This file controls the state of SELinux on the system.
# SELINUX= can take one of these three values:
#     enforcing - SELinux security policy is enforced.
#     permissive - SELinux prints warnings instead of enforcing.
#     disabled - No SELinux policy is loaded.
SELINUX=permissive
# SELINUXTYPE= can take one of these two values:
#     targeted - Targeted processes are protected,
#     mls - Multi Level Security protection.
SELINUXTYPE=targeted
```

Stop the VM and create 3 full clones as follows, each having **two** network interfaces. The first should be natted whilst the second should be in an internal network called **intnet**. Do NOT reinitialize the MAC address of each eth0 :

Clone	**RAM**
Node2	512 MB
Node3	1500 MB

Clone	RAM
Node4	1500 MB

Add a second network interface to node1 and put it into the intnet network.

Boot node1, node2, node3 and node4. Change the **/etc/inittab** file on node2, node3 and node4 so that they start up in runlevel 3. Configure eth1 as 192.168.22.105, 192.168.22.106, 192.168.22.107 et 192.68.22.108 respectively in each node. Make sure sshd is started on each node.

Open a terminal on node1 and type **/cluster/mysql_cluster/bin/ndb_setup.py**:

```
[root@node1 ~]# /cluster/mysql_cluster/bin/ndb_setup.py
Running out of install dir: /cluster/mysql_cluster/bin
Starting web server on port 8081
deathkey=433749
The application should now be running in your browser.
(Alternatively you can navigate to http://localhost:8081/welcome.html to start it)
```

In your browser, navigate to: **http://localhost:8081/welcome.html**:

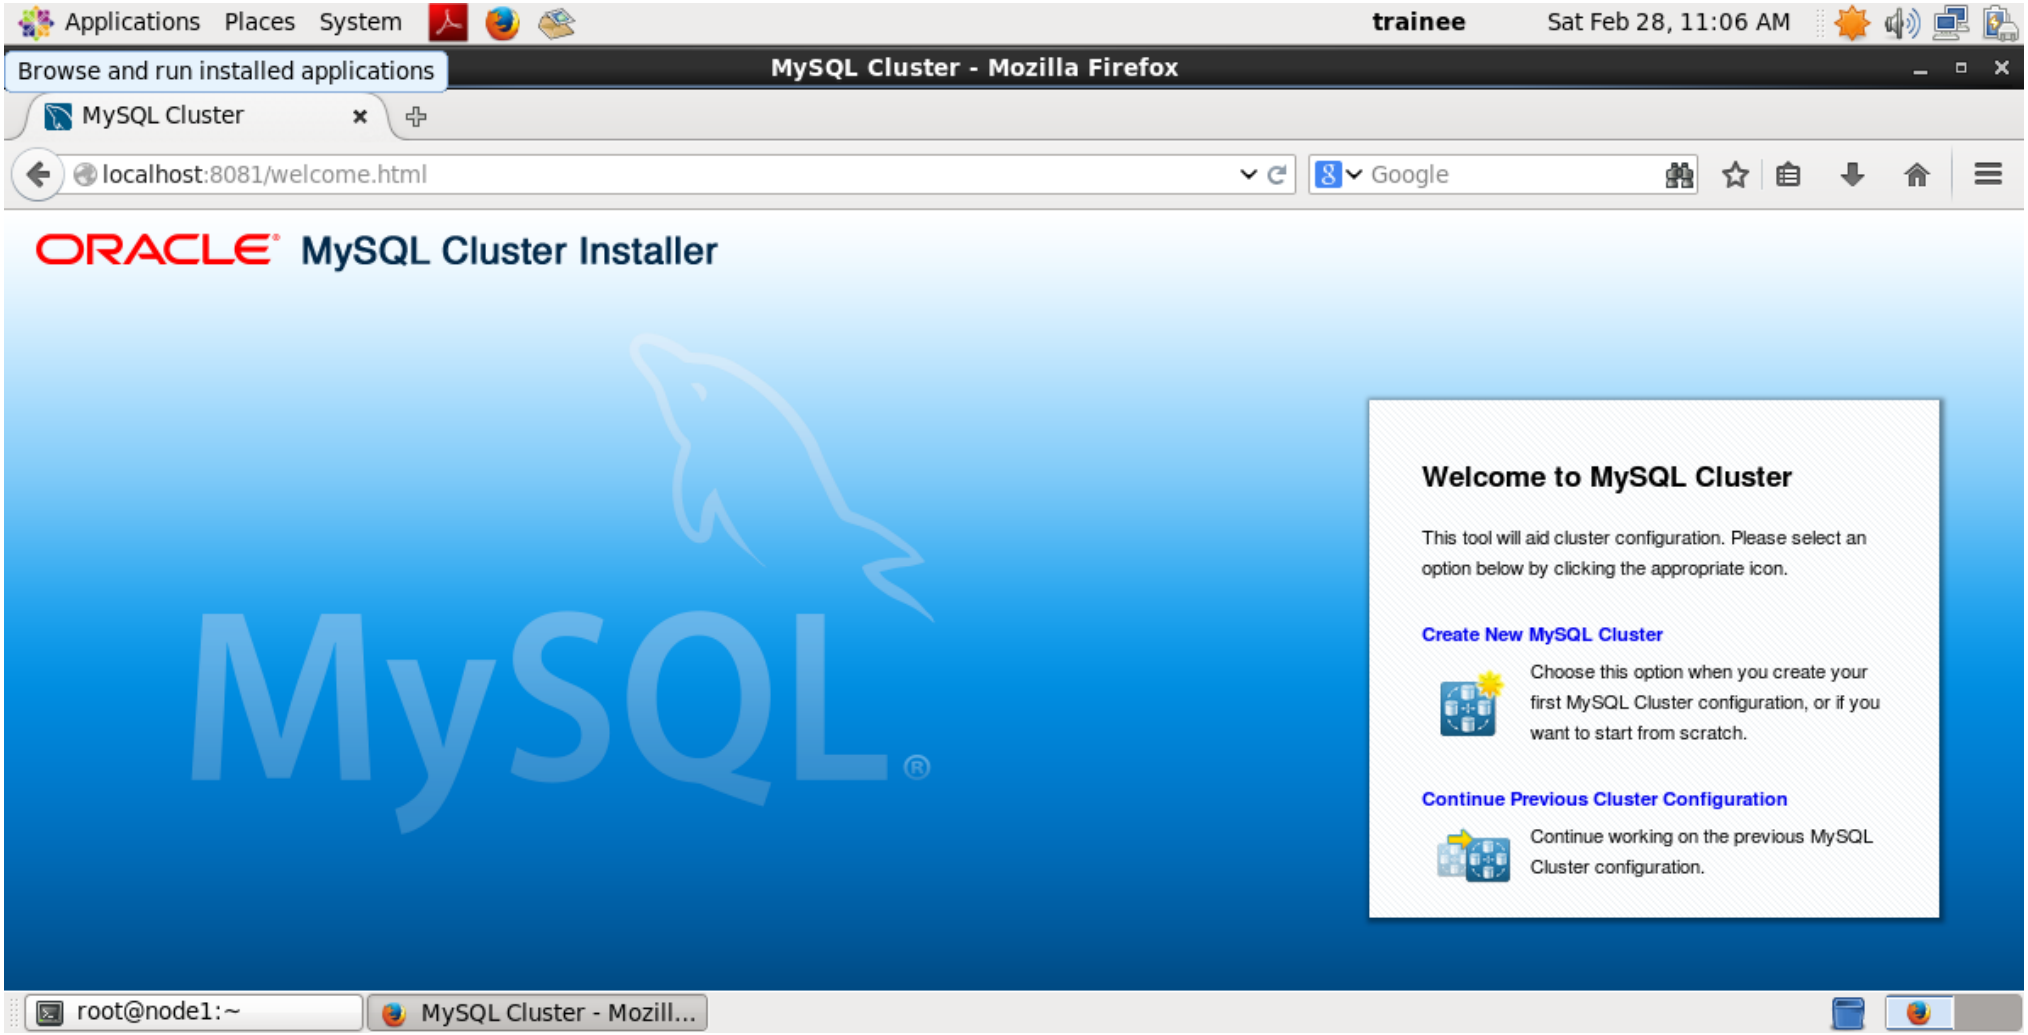

Click on the link **Create New MySQL Cluster** and then fill in the top as follows:

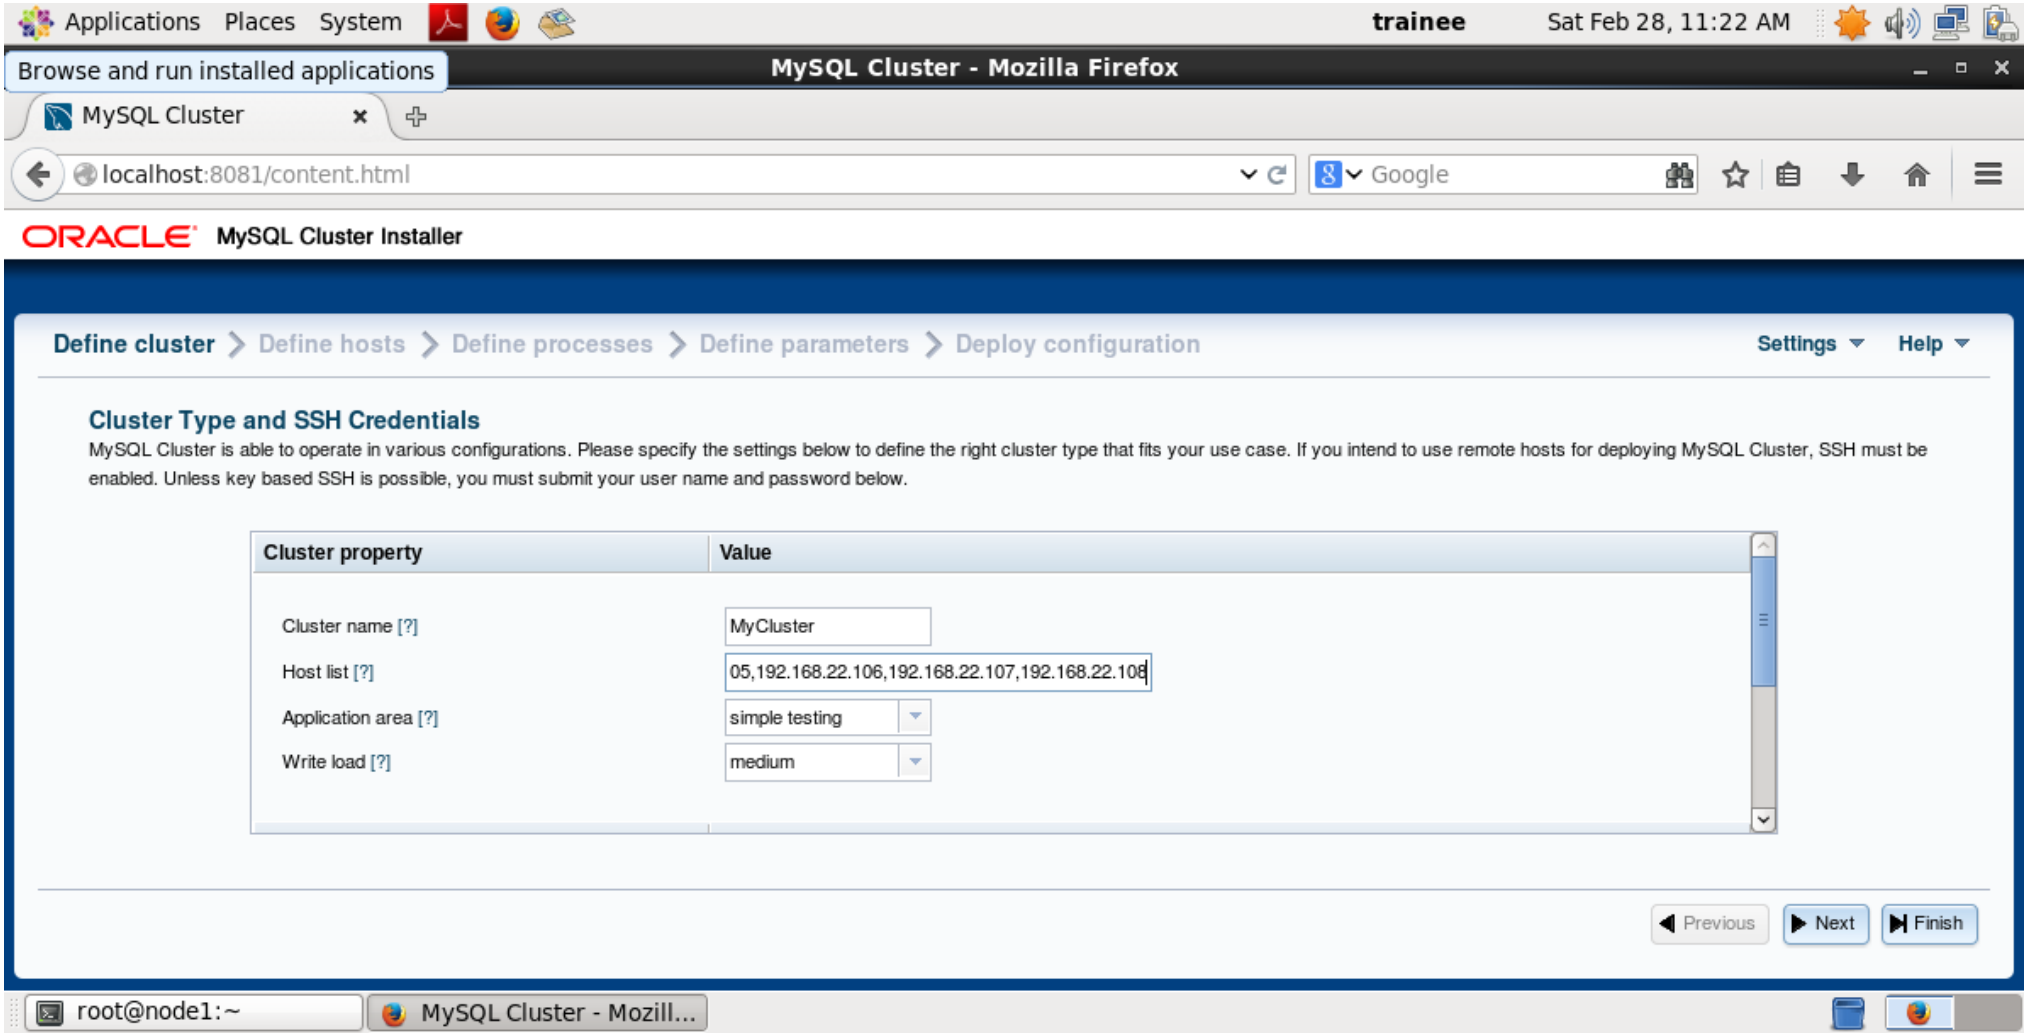

At the bottom of the page, untick **Key based SSH**, and input **root** and **fenestros**. Click on the **Next** button:

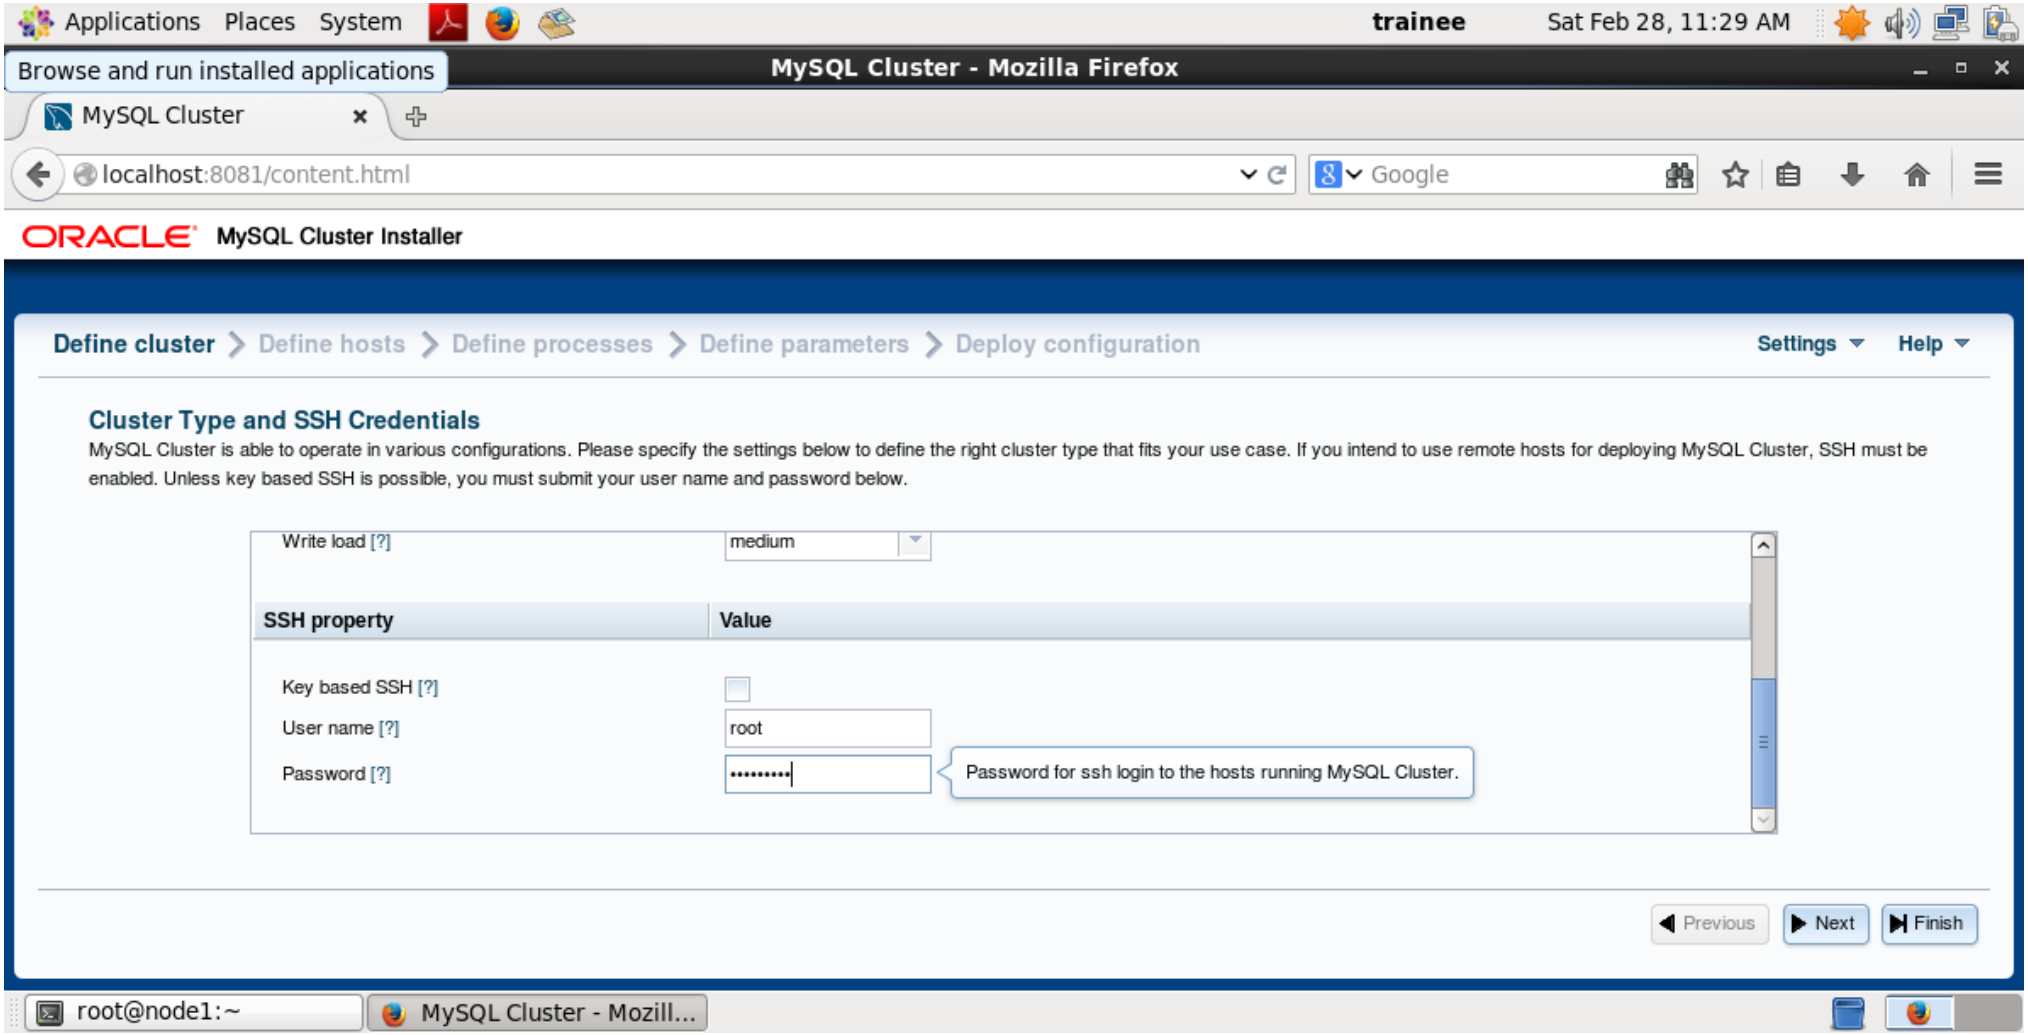

Your nodes should automatically be detected:

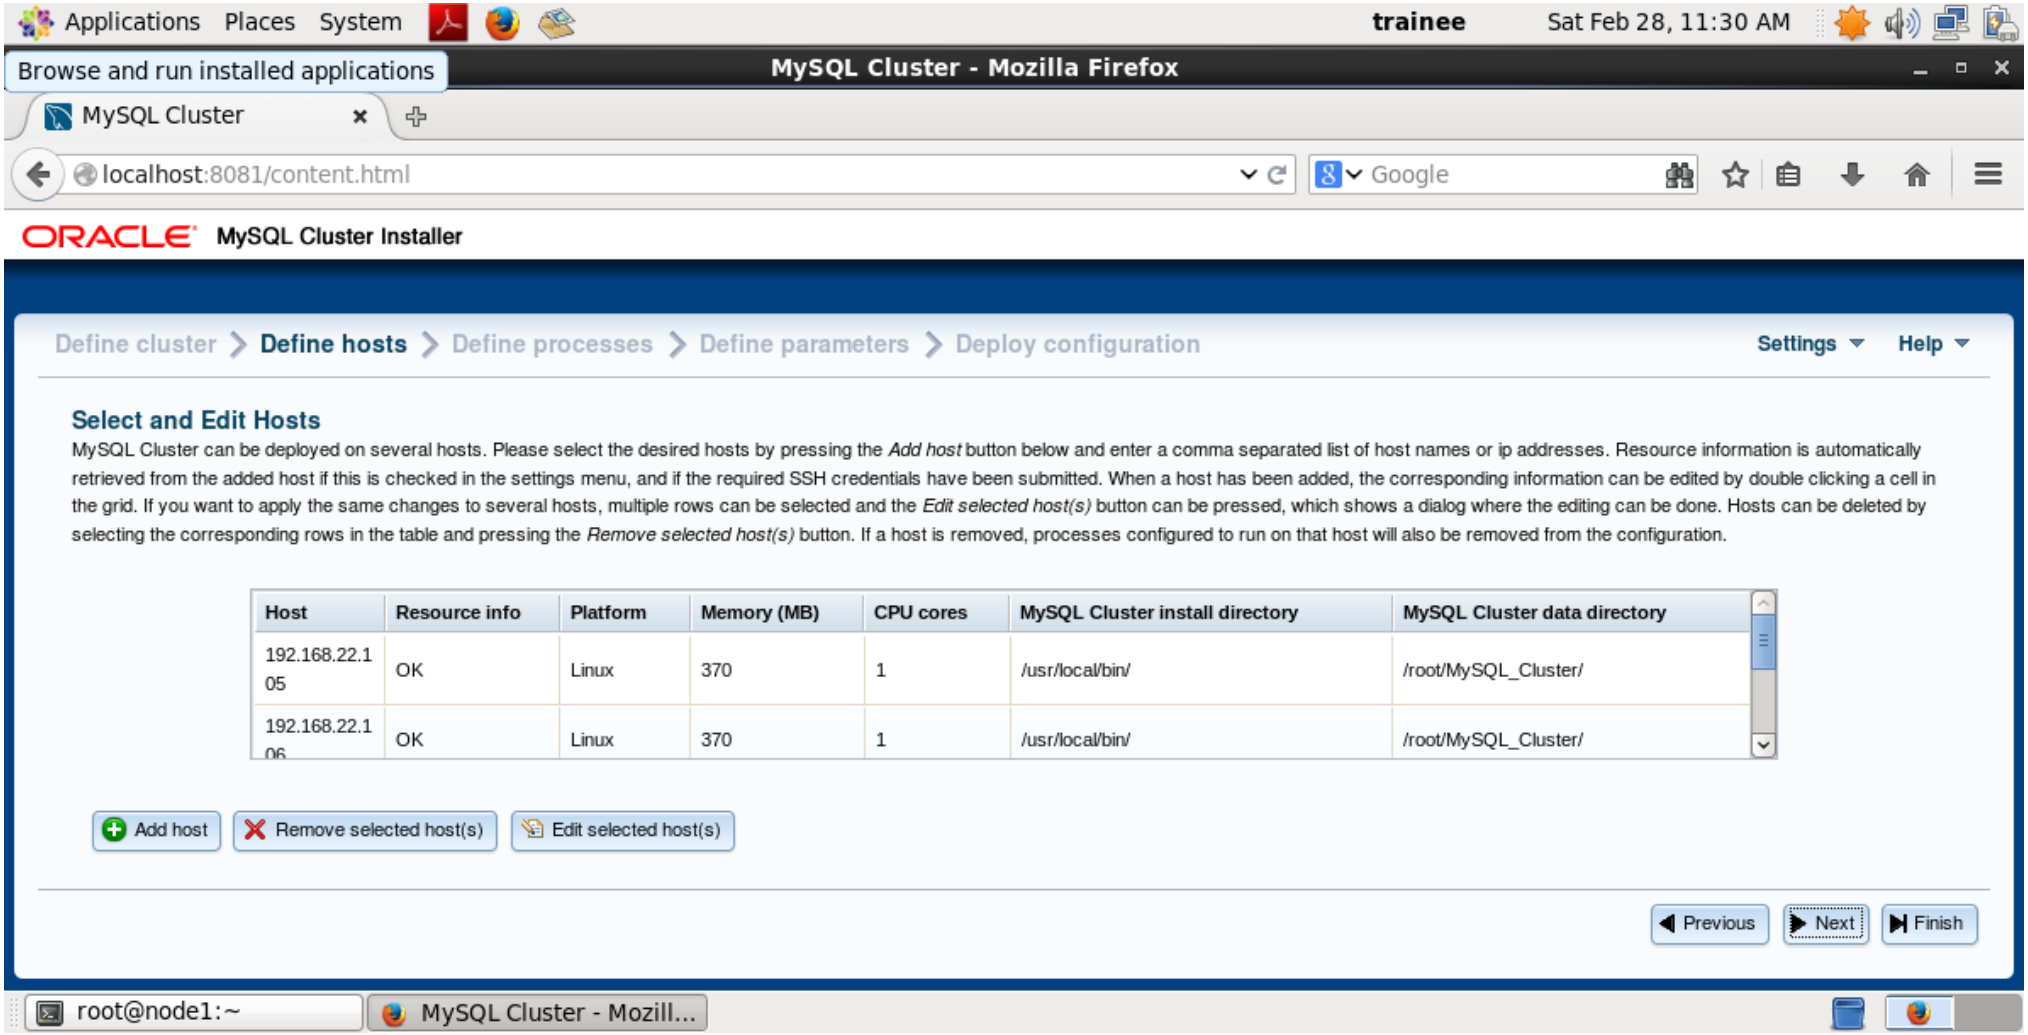

Change the directories as shown below for **all** the nodes and click on the **Next** button:

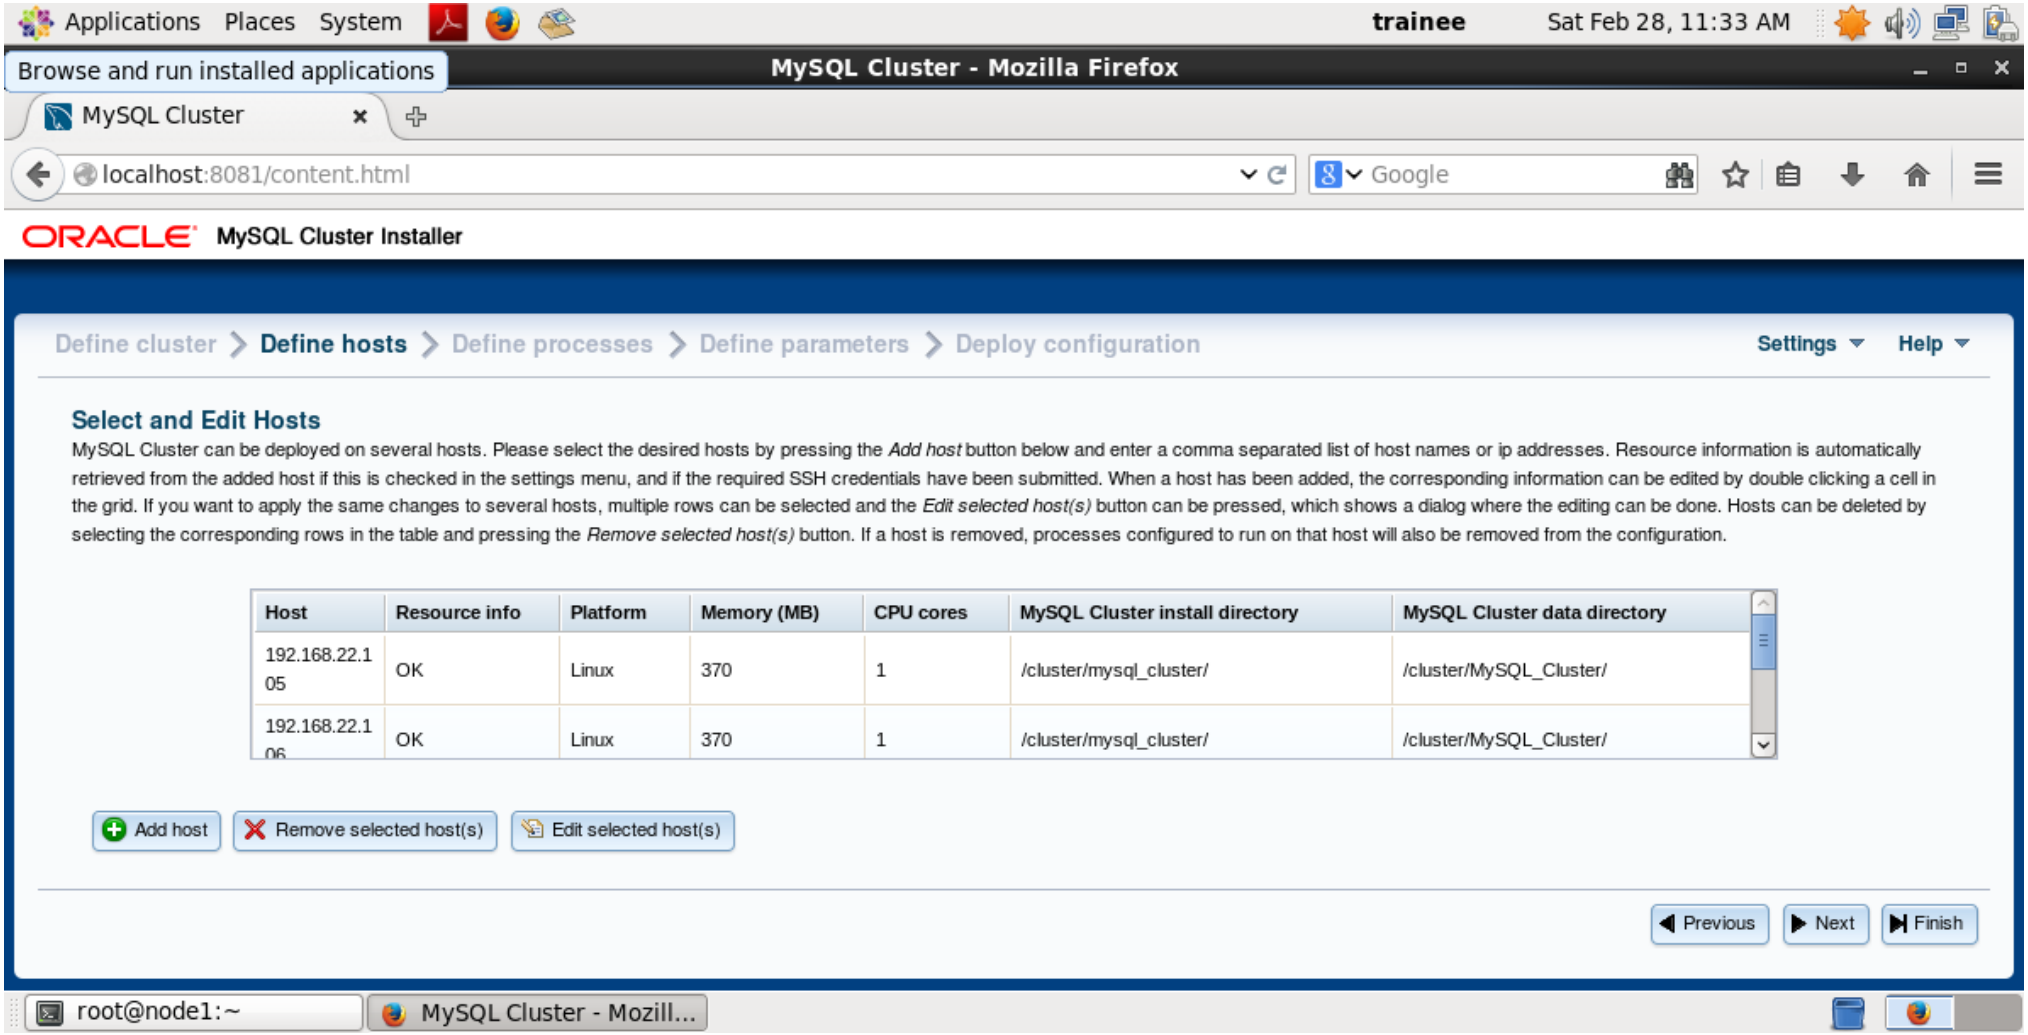

Examine the physical topology of the cluster. Note how the installer has set up the high availability. Since we do not need them for this LAB, delete the API services and click on **Next** :

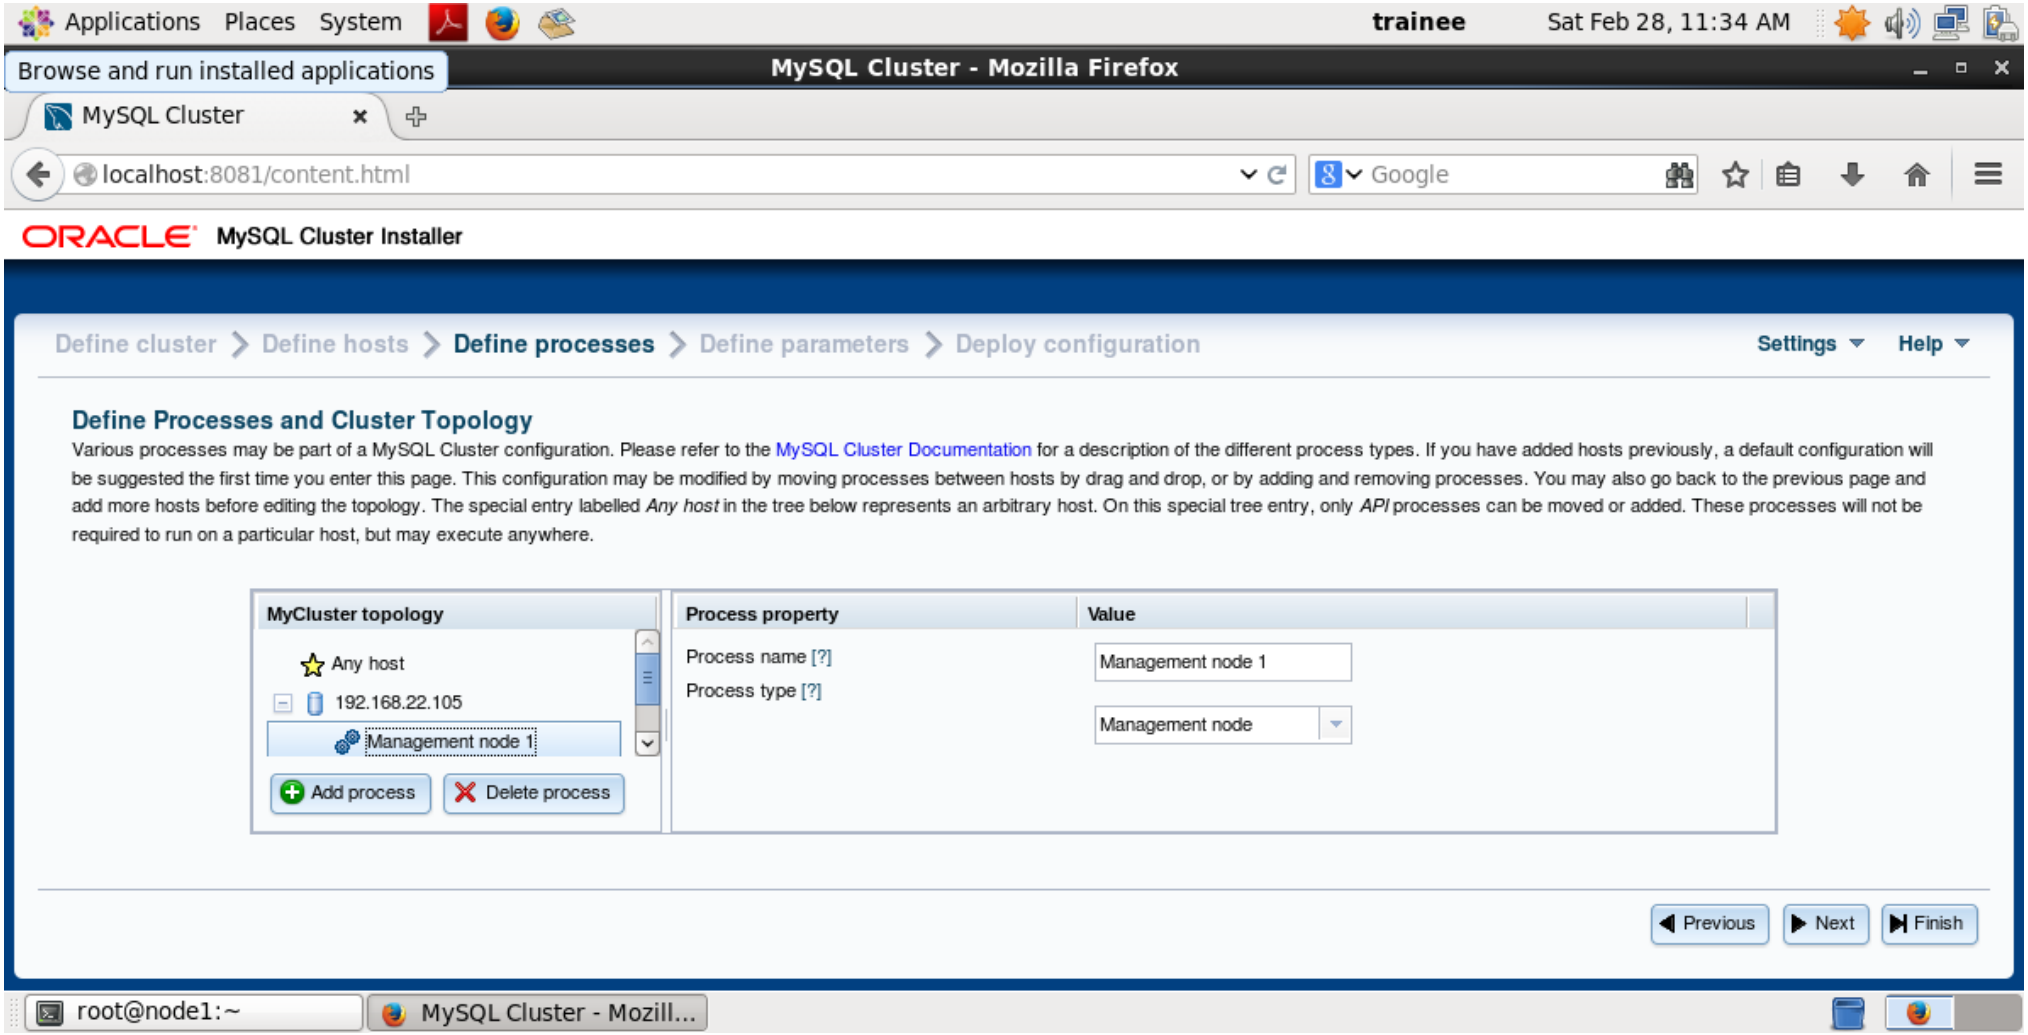

Note that each node or service is associated with a node number. Examine the services and click on **Next** :

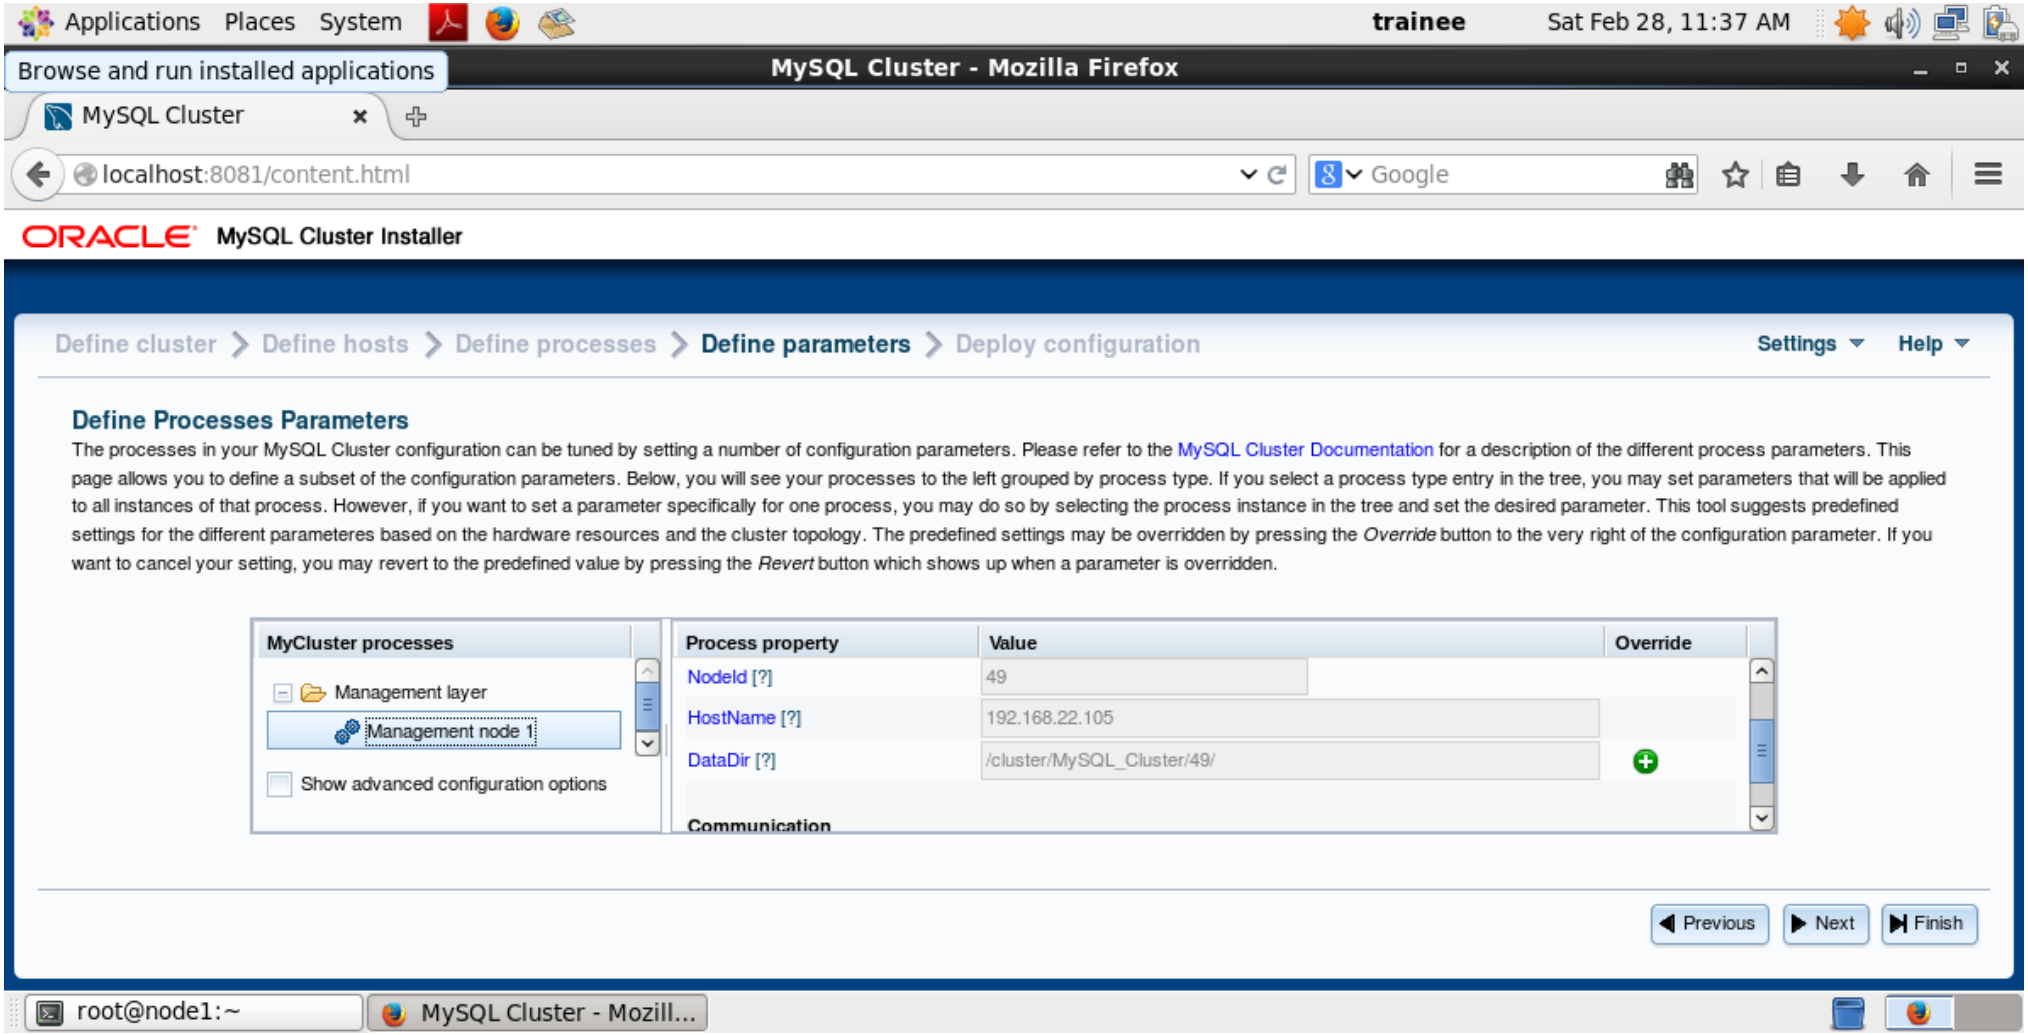

Now click on **Deploy cluster**:

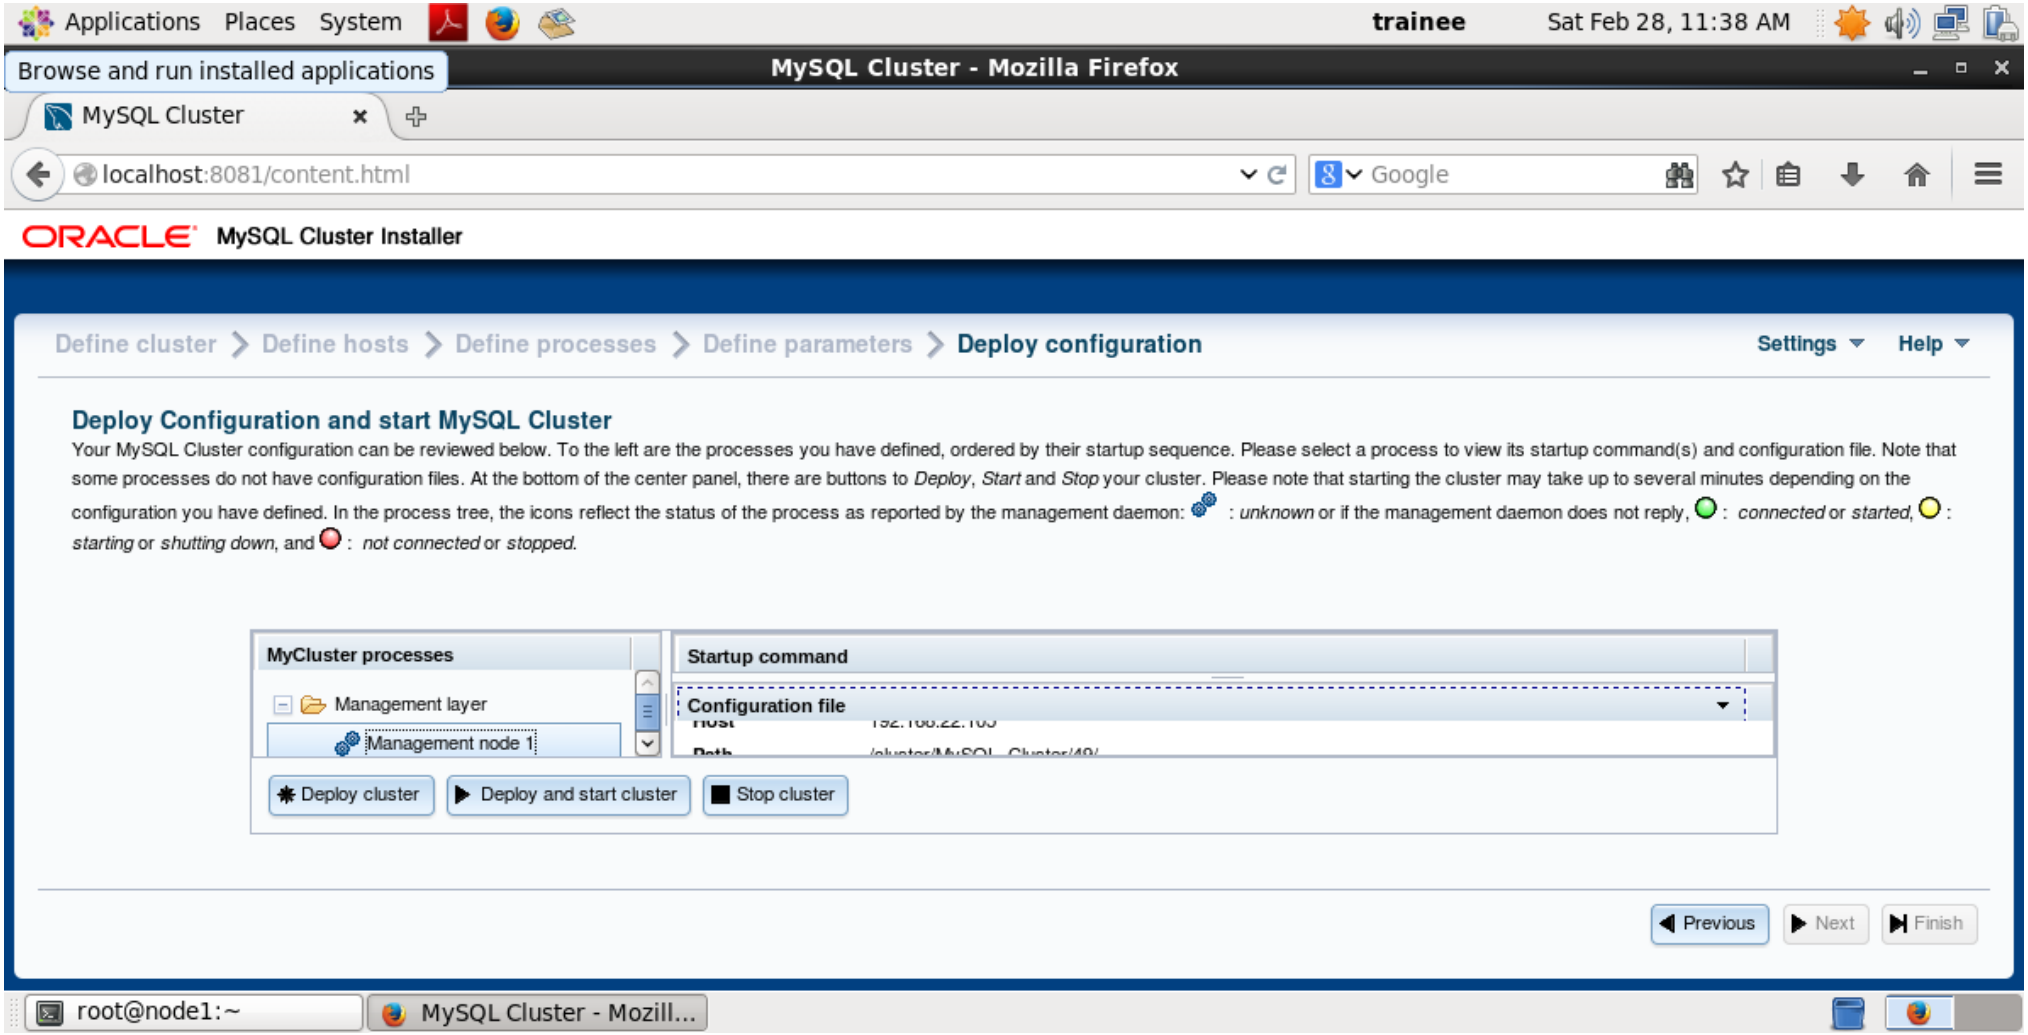

Deploying the cluster can take upto 30 minutes:

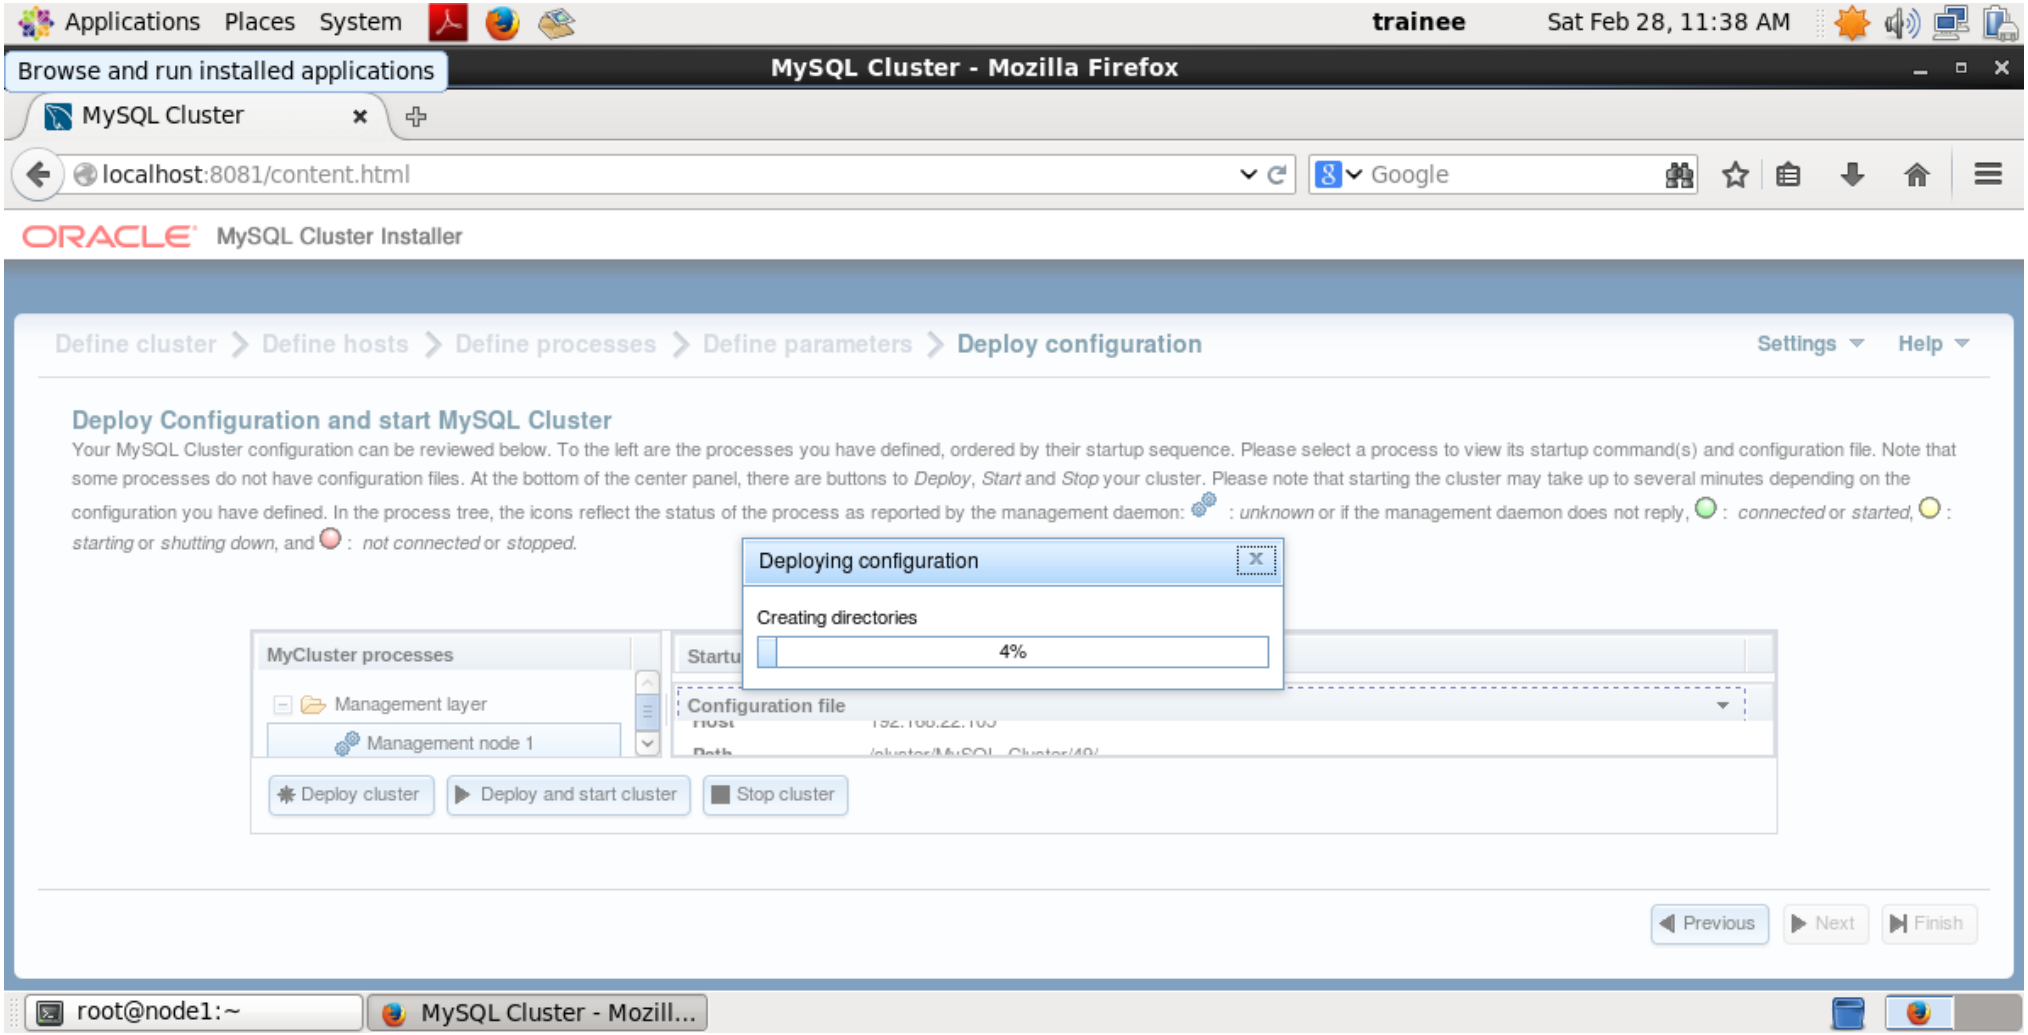

Once the deployment has finished, you should see the following message:

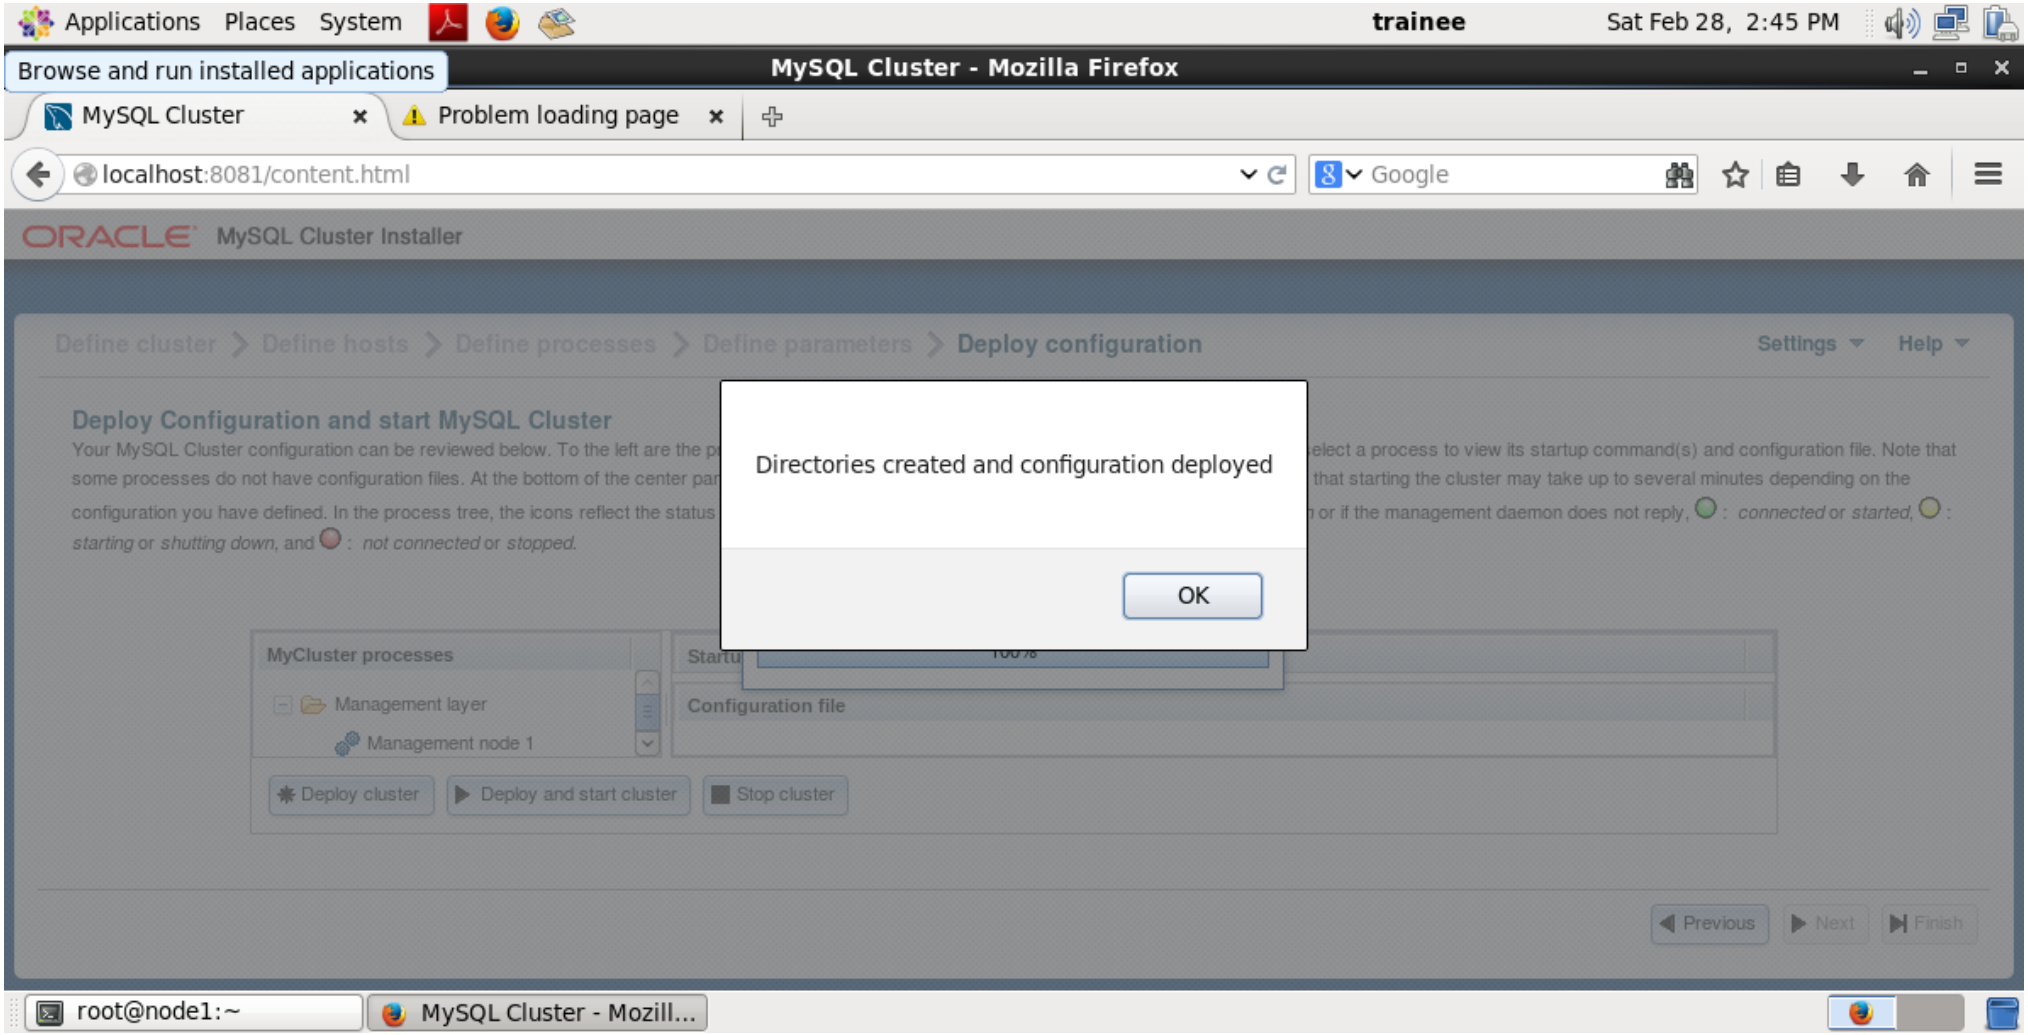

Cluster Configuration

The cluster configuration is held in the **/cluster/MySQL_Cluster/** directory. The following shows the contents of the node1 host's configuration directory:

```
[root@node1 ~]# ls -lR /cluster/MySQL_Cluster/
/cluster/MySQL_Cluster/:
total 12
drwxr-xr-x. 2 root root 4096 Feb 28 14:42 49
drwxr-xr-x. 5 root root 4096 Feb 28 14:42 55
drwxr-xr-x. 5 root root 4096 Feb 28 14:42 56

/cluster/MySQL_Cluster/49:
total 4
-rw-r--r--. 1 root root 1426 Feb 28 14:42 config.ini

/cluster/MySQL_Cluster/55:
total 16
-rw-r--r--. 1 root root  346 Feb 28 14:42 my.cnf
drwxr-xr-x. 2 root root 4096 Feb 28 14:20 mysql
drwxr-xr-x. 2 root root 4096 Feb 28 14:19 test
drwxr-xr-x. 2 root root 4096 Feb 28 14:22 tmp

/cluster/MySQL_Cluster/55/mysql:
total 0

/cluster/MySQL_Cluster/55/test:
total 0

/cluster/MySQL_Cluster/55/tmp:
total 0

/cluster/MySQL_Cluster/56:
total 16
-rw-r--r--. 1 root root  346 Feb 28 14:42 my.cnf
drwxr-xr-x. 2 root root 4096 Feb 28 14:26 mysql
drwxr-xr-x. 2 root root 4096 Feb 28 14:24 test
drwxr-xr-x. 2 root root 4096 Feb 28 14:27 tmp
```

```
/cluster/MySQL_Cluster/56/mysql:
total 0

/cluster/MySQL_Cluster/56/test:
total 0

/cluster/MySQL_Cluster/56/tmp:
total 0
```

<note important> In the above example, we can see the configuration directories of the Cluster Management node (**49**) and the two Cluster SQL node (**55** and **56**). </note>

Node2's configuration is identical to that of node1, with the exception of the directory numbers:

```
[root@node2 ~]# ls -lR /cluster/MySQL_Cluster/
/cluster/MySQL_Cluster/:
total 12
drwxr-xr-x. 2 root root 4096 Feb 28 14:42 52
drwxr-xr-x. 5 root root 4096 Feb 28 14:42 57
drwxr-xr-x. 5 root root 4096 Feb 28 14:42 58

/cluster/MySQL_Cluster/52:
total 4
-rw-r--r--. 1 root root 1426 Feb 28 14:42 config.ini

/cluster/MySQL_Cluster/57:
total 16
-rw-r--r--. 1 root root  346 Feb 28 14:42 my.cnf
drwxr-xr-x. 2 root root 4096 Feb 28 14:31 mysql
drwxr-xr-x. 2 root root 4096 Feb 28 14:30 test
drwxr-xr-x. 2 root root 4096 Feb 28 14:32 tmp

/cluster/MySQL_Cluster/57/mysql:
total 0
```

```
/cluster/MySQL_Cluster/57/test:
total 0

/cluster/MySQL_Cluster/57/tmp:
total 0

/cluster/MySQL_Cluster/58:
total 16
-rw-r--r--. 1 root root  346 Feb 28 14:42 my.cnf
drwxr-xr-x. 2 root root 4096 Feb 28 14:37 mysql
drwxr-xr-x. 2 root root 4096 Feb 28 14:35 test
drwxr-xr-x. 2 root root 4096 Feb 28 14:38 tmp

/cluster/MySQL_Cluster/58/mysql:
total 0

/cluster/MySQL_Cluster/58/test:
total 0

/cluster/MySQL_Cluster/58/tmp:
total 0
```

Node3 and Node4's configuration are different in so much as they represent the services installed on them, Multi threaded data node 1 and Multi threaded data node 2 respectively:

```
[root@node3 ~]# ls -lR /cluster/MySQL_Cluster/
/cluster/MySQL_Cluster/:
total 4
drwxr-xr-x. 2 root root 4096 Feb 28 14:39 1

/cluster/MySQL_Cluster/1:
total 0
```

```
[root@node1 ~]# ls -lR /cluster/MySQL_Cluster/
```

```
/cluster/MySQL_Cluster/:
total 4
drwxr-xr-x. 2 root root 4096 Feb 28 14:41 2

/cluster/MySQL_Cluster/2:
total 0
```

The MySQL Cluster global configuration file is named **config.ini**, but only by convention. If the file does not exist, the configuration is read from the **configuration cache**. The location of this file has no default value and it can be placed in any location that can be read by ndb_mgmd. The location and name of the configuration file are specified using –config-file=path_name. The config.ini file uses the INI format where the content is split into [SECTIONS] containing parameter=value declarations which can either use the **=** or **:** characters as seperators. One addition to the classic INI format is the use of a unique ID specified as a parameter within the section to specifically identify several nodes of the same type. Each node has a section in the config.ini file. In our case the cluster has two management nodes, so the configuration file contains two [NDB_MGMD] sections.

<note important> All Cluster parameter names are case-insensitive. </note>

The sections within this file are:

Section	Description
[NDB_MGMD DEFAULT]	Contains parameters that apply to all cluster management server nodes (MGM nodes)
[MGM] or [NDB_MGMD]	Defines a cluster management server node
[TCP] or [TCP DEFAULT]	Defines a TCP/IP connection between cluster nodes
[NDBD DEFAULT]	Contains parameters that apply to all data nodes (ndbd processes)
[NDBD]	Defines a cluster data node
[MYSQLD DEFAULT]	Contains parameters that apply to all MySQL server nodes (SQL or API nodes)
[MYSQLD]	Defines the cluster's MySQL server nodes

```
[root@node1 ~]# cat /cluster/MySQL_Cluster/49/config.ini
#
# Configuration file for MyCluster
#

[NDB_MGMD DEFAULT]
Portnumber=1186
```

```
[NDB_MGMD]
NodeId=49
HostName=192.168.22.105
DataDir=/cluster/MySQL_Cluster/49/
Portnumber=1186

[NDB_MGMD]
NodeId=52
HostName=192.168.22.106
DataDir=/cluster/MySQL_Cluster/52/
Portnumber=1186

[TCP DEFAULT]
SendBufferMemory=4M
ReceiveBufferMemory=4M

[NDBD DEFAULT]
BackupMaxWriteSize=1M
BackupDataBufferSize=16M
BackupLogBufferSize=4M
BackupMemory=20M
BackupReportFrequency=10
MemReportFrequency=30
LogLevelStartup=15
LogLevelShutdown=15
LogLevelCheckpoint=8
LogLevelNodeRestart=15
DataMemory=1M
IndexMemory=1M
MaxNoOfTables=4096
MaxNoOfTriggers=3500
NoOfReplicas=2
StringMemory=25
DiskPageBufferMemory=64M
```

```
SharedGlobalMemory=20M
LongMessageBuffer=32M
MaxNoOfConcurrentTransactions=16384
BatchSizePerLocalScan=512
FragmentLogFileSize=64M
NoOfFragmentLogFiles=16
RedoBuffer=32M
MaxNoOfExecutionThreads=2
StopOnError=false
LockPagesInMainMemory=1
TimeBetweenEpochsTimeout=32000
TimeBetweenWatchdogCheckInitial=60000
TransactionInactiveTimeout=60000
HeartbeatIntervalDbDb=15000
HeartbeatIntervalDbApi=15000

[NDBD]
NodeId=1
HostName=192.168.22.107
DataDir=/cluster/MySQL_Cluster/1/

[NDBD]
NodeId=2
HostName=192.168.22.108
DataDir=/cluster/MySQL_Cluster/2/

[MYSQLD DEFAULT]

[MYSQLD]
NodeId=55
HostName=192.168.22.105

[MYSQLD]
NodeId=56
```

```
HostName=192.168.22.105

[MYSQLD]
NodeId=57
HostName=192.168.22.106

[MYSQLD]
NodeId=58
HostName=192.168.22.106
```

```
[root@node2 ~]# cat /cluster/MySQL_Cluster/52/config.ini
#
# Configuration file for MyCluster
#

[NDB_MGMD DEFAULT]
Portnumber=1186

[NDB_MGMD]
NodeId=49
HostName=192.168.22.105
DataDir=/cluster/MySQL_Cluster/49/
Portnumber=1186

[NDB_MGMD]
NodeId=52
HostName=192.168.22.106
DataDir=/cluster/MySQL_Cluster/52/
Portnumber=1186

[TCP DEFAULT]
SendBufferMemory=4M
ReceiveBufferMemory=4M
```

```
[NDBD DEFAULT]
BackupMaxWriteSize=1M
BackupDataBufferSize=16M
BackupLogBufferSize=4M
BackupMemory=20M
BackupReportFrequency=10
MemReportFrequency=30
LogLevelStartup=15
LogLevelShutdown=15
LogLevelCheckpoint=8
LogLevelNodeRestart=15
DataMemory=1M
IndexMemory=1M
MaxNoOfTables=4096
MaxNoOfTriggers=3500
NoOfReplicas=2
StringMemory=25
DiskPageBufferMemory=64M
SharedGlobalMemory=20M
LongMessageBuffer=32M
MaxNoOfConcurrentTransactions=16384
BatchSizePerLocalScan=512
FragmentLogFileSize=64M
NoOfFragmentLogFiles=16
RedoBuffer=32M
MaxNoOfExecutionThreads=2
StopOnError=false
LockPagesInMainMemory=1
TimeBetweenEpochsTimeout=32000
TimeBetweenWatchdogCheckInitial=60000
TransactionInactiveTimeout=60000
HeartbeatIntervalDbDb=15000
HeartbeatIntervalDbApi=15000
```

```
[NDBD]
NodeId=1
HostName=192.168.22.107
DataDir=/cluster/MySQL_Cluster/1/

[NDBD]
NodeId=2
HostName=192.168.22.108
DataDir=/cluster/MySQL_Cluster/2/

[MYSQLD DEFAULT]

[MYSQLD]
NodeId=55
HostName=192.168.22.105

[MYSQLD]
NodeId=56
HostName=192.168.22.105

[MYSQLD]
NodeId=57
HostName=192.168.22.106

[MYSQLD]
NodeId=58
HostName=192.168.22.106
```

<note important> These two files are and must remain identical in order for the cluster to continue to function correctly. </note>

The definitions of the parameters in these files are as follows:

Section	Default Value	Parameter	Definition

Section	Default Value	Parameter	Definition
[NDB_MGMD DEFAULT]	1186	Portnumber=	http://dev.mysql.com/doc/refman/5.6/en/mysql-cluster-mgm-definition.html
[NDB_MGMD]	N/A	NodeId=	http://dev.mysql.com/doc/refman/5.6/en/mysql-cluster-mgm-definition.html
[NDB_MGMD]	N/A	HostName=	http://dev.mysql.com/doc/refman/5.6/en/mysql-cluster-mgm-definition.html
[NDB_MGMD]	Process CWD	DataDir=	http://dev.mysql.com/doc/refman/5.6/en/mysql-cluster-mgm-definition.html
[NDB_MGMD]	1186	Portnumber=	http://dev.mysql.com/doc/refman/5.6/en/mysql-cluster-mgm-definition.html
[TCP DEFAULT]	2GB	SendBufferMemory=	http://dev.mysql.com/doc/refman/5.6/en/mysql-cluster-tcp-definition.html#ndbparam-tcp-sendbuffermemory
[TCP DEFAULT]	2GB	ReceiveBufferMemory=	http://dev.mysql.com/doc/refman/5.6/en/mysql-cluster-tcp-definition.html#ndbparam-tcp-receivebuffermemory
[NDBD DEFAULT]	1MB	BackupMaxWriteSize=	http://dev.mysql.com/doc/refman/5.6/en/mysql-cluster-ndbd-definition.html#ndbparam-ndbd-backupmaxwritesize
[NDBD DEFAULT]	16MB	BackupDataBufferSize=	http://dev.mysql.com/doc/refman/5.6/en/mysql-cluster-ndbd-definition.html#ndbparam-ndbd-backupdatabuffersize
[NDBD DEFAULT]	16MB	BackupLogBufferSize=	http://dev.mysql.com/doc/refman/5.6/en/mysql-cluster-ndbd-definition.html#ndbparam-ndbd-backuplogbuffersize
[NDBD DEFAULT]	N/A	BackupMemory=	http://dev.mysql.com/doc/refman/5.6/en/mysql-cluster-ndbd-definition.html#ndbparam-ndbd-backupmemory
[NDBD DEFAULT]	0s	BackupReportFrequency=	http://dev.mysql.com/doc/refman/5.6/en/mysql-cluster-ndbd-definition.html#ndbparam-ndbd-backupreportfrequency
[NDBD DEFAULT]	0s	MemReportFrequency=	http://dev.mysql.com/doc/refman/5.6/en/mysql-cluster-ndbd-definition.html#ndbparam-ndbd-memreportfrequency
[NDBD DEFAULT]	1s	LogLevelStartup=	http://dev.mysql.com/doc/refman/5.6/en/mysql-cluster-ndbd-definition.html#ndbparam-ndbd-loglevelstartup
[NDBD DEFAULT]	0s	LogLevelShutdown=	http://dev.mysql.com/doc/refman/5.6/en/mysql-cluster-ndbd-definition.html#ndbparam-ndbd-loglevelshutdown
[NDBD DEFAULT]	0s	LogLevelCheckpoint=	http://dev.mysql.com/doc/refman/5.6/en/mysql-cluster-ndbd-definition.html#ndbparam-ndbd-loglevelcheckpoint
[NDBD DEFAULT]	0s	LogLevelNodeRestart=	http://dev.mysql.com/doc/refman/5.6/en/mysql-cluster-ndbd-definition.html#ndbparam-ndbd-loglevelnoderestart
[NDBD DEFAULT]	80MB	DataMemory=	http://dev.mysql.com/doc/refman/5.6/en/mysql-cluster-ndbd-definition.html#ndbparam-ndbd-datamemory
[NDBD DEFAULT]	18M	IndexMemory=	http://dev.mysql.com/doc/refman/5.6/en/mysql-cluster-ndbd-definition.html#ndbparam-ndbd-indexmemory
[NDBD DEFAULT]	128	MaxNoOfTables=	http://dev.mysql.com/doc/refman/5.6/en/mysql-cluster-ndbd-definition.html#ndbparam-ndbd-maxnooftables
[NDBD DEFAULT]	768	MaxNoOfTriggers=	http://dev.mysql.com/doc/refman/5.6/en/mysql-cluster-ndbd-definition.html#ndbparam-ndbd-maxnooftriggers

Section	Default Value	Parameter	Definition
[NDBD DEFAULT]	2	NoOfReplicas=	http://dev.mysql.com/doc/refman/5.6/en/mysql-cluster-ndbd-definition.html#ndbparam-ndbd-noofreplicas
[NDBD DEFAULT]	25	StringMemory=	http://dev.mysql.com/doc/refman/5.6/en/mysql-cluster-ndbd-definition.html#ndbparam-ndbd-stringmemory
[NDBD DEFAULT]	64M	DiskPageBufferMemory=	http://dev.mysql.com/doc/refman/5.6/en/mysql-cluster-ndbd-definition.html#ndbparam-ndbd-diskpagebuffermemory
[NDBD DEFAULT]	128M	SharedGlobalMemory=	http://dev.mysql.com/doc/refman/5.6/en/mysql-cluster-ndbd-definition.html#ndbparam-ndbd-sharedglobalmemory
[NDBD DEFAULT]	64M	LongMessageBuffer=	http://dev.mysql.com/doc/refman/5.6/en/mysql-cluster-ndbd-definition.html#ndbparam-ndbd-longmessagebuffer
[NDBD DEFAULT]	4096	MaxNoOfConcurrentTransactions=	http://dev.mysql.com/doc/refman/5.6/en/mysql-cluster-ndbd-definition.html#ndbparam-ndbd-maxnoofconcurrenttransactions
[NDBD DEFAULT]	256	BatchSizePerLocalScan=	http://dev.mysql.com/doc/refman/5.6/en/mysql-cluster-ndbd-definition.html#ndbparam-ndbd-batchsizeperlocalscan
[NDBD DEFAULT]	16M	FragmentLogFileSize=	http://dev.mysql.com/doc/refman/5.6/en/mysql-cluster-ndbd-definition.html#ndbparam-ndbd-fragmentlogfilesize
[NDBD DEFAULT]	16	NoOfFragmentLogFiles=	http://dev.mysql.com/doc/refman/5.6/en/mysql-cluster-ndbd-definition.html#ndbparam-ndbd-nooffragmentlogfiles
[NDBD DEFAULT]	32M	RedoBuffer=	http://dev.mysql.com/doc/refman/5.6/en/mysql-cluster-ndbd-definition.html#ndbparam-ndbd-redobuffer
[NDBD DEFAULT]	2	MaxNoOfExecutionThreads=	http://dev.mysql.com/doc/refman/5.6/en/mysql-cluster-ndbd-definition.html#ndbparam-ndbmtd-maxnoofexecutionthreads
[NDBD DEFAULT]	1	StopOnError=	http://dev.mysql.com/doc/refman/5.6/en/mysql-cluster-ndbd-definition.html#ndbparam-ndbd-stoponerror
[NDBD DEFAULT]	0	LockPagesInMainMemory=	http://dev.mysql.com/doc/refman/5.6/en/mysql-cluster-ndbd-definition.html#ndbparam-ndbd-lockpagesinmainmemory
[NDBD DEFAULT]	0	TimeBetweenEpochsTimeout=	http://dev.mysql.com/doc/refman/5.6/en/mysql-cluster-ndbd-definition.html#ndbparam-ndbd-timebetweenepochstimeout
[NDBD DEFAULT]	6000ms	TimeBetweenWatchdogCheckInitial=	http://dev.mysql.com/doc/refman/5.6/en/mysql-cluster-ndbd-definition.html#ndbparam-ndbd-timebetweenwatchdogcheckinitial
[NDBD DEFAULT]	4GB	TransactionInactiveTimeout=	http://dev.mysql.com/doc/refman/5.6/en/mysql-cluster-ndbd-definition.html#ndbparam-ndbd-transactioninactivetimeout
[NDBD DEFAULT]	5000ms	HeartbeatIntervalDbDb=	http://dev.mysql.com/doc/refman/5.6/en/mysql-cluster-ndbd-definition.html#ndbparam-ndbd-heartbeatintervaldbdb
[NDBD DEFAULT]	1500ms	HeartbeatIntervalDbApi	http://dev.mysql.com/doc/refman/5.6/en/mysql-cluster-ndbd-definition.html#ndbparam-ndbd-heartbeatintervaldbapi
[NDBD]	N/A	NodeId=	http://dev.mysql.com/doc/refman/5.6/en/mysql-cluster-ndbd-definition.html
[NDBD]	N/A	HostName=	http://dev.mysql.com/doc/refman/5.6/en/mysql-cluster-ndbd-definition.html
[NDBD]	N/A	DataDir=	http://dev.mysql.com/doc/refman/5.6/en/mysql-cluster-ndbd-definition.html
[MYSQLD]	N/A	NodeId=	http://dev.mysql.com/doc/refman/5.6/en/mysql-cluster-api-definition.html#ndbparam-api-nodeid

Section	Default Value	Parameter	Definition
[MYSQLD]	N/A	HostName=	http://dev.mysql.com/doc/refman/5.6/en/mysql-cluster-api-definition.html

MySQL Configuration

The four Cluster SQL nodes contain a **my.cnf** file:

```
[root@node1 ~]# cat /cluster/MySQL_Cluster/55/my.cnf
#
# Configuration file for MyCluster
# Generated by mcc
#
[mysqld]
log-error=mysqld.55.err
datadir="/cluster/MySQL_Cluster/55/"
tmpdir="/cluster/MySQL_Cluster/55/tmp"
basedir="/cluster/mysql_cluster/"
port=3306
ndbcluster=on
ndb-nodeid=55
ndb-connectstring=192.168.22.105:1186,192.168.22.106:1186,
socket="/cluster/MySQL_Cluster/55/mysql.socket"
[root@node1 ~]# cat /cluster/MySQL_Cluster/56/my.cnf
#
# Configuration file for MyCluster
# Generated by mcc
#
[mysqld]
log-error=mysqld.56.err
datadir="/cluster/MySQL_Cluster/56/"
tmpdir="/cluster/MySQL_Cluster/56/tmp"
basedir="/cluster/mysql_cluster/"
```

```
port=3307
ndbcluster=on
ndb-nodeid=56
ndb-connectstring=192.168.22.105:1186,192.168.22.106:1186,
socket="/cluster/MySQL_Cluster/56/mysql.socket"
```

```
[root@node2 ~]# cat /cluster/MySQL_Cluster/57/my.cnf
#
# Configuration file for MyCluster
# Generated by mcc
#
[mysqld]
log-error=mysqld.57.err
datadir="/cluster/MySQL_Cluster/57/"
tmpdir="/cluster/MySQL_Cluster/57/tmp"
basedir="/cluster/mysql_cluster/"
port=3306
ndbcluster=on
ndb-nodeid=57
ndb-connectstring=192.168.22.105:1186,192.168.22.106:1186,
socket="/cluster/MySQL_Cluster/57/mysql.socket"
[root@node2 ~]# cat /cluster/MySQL_Cluster/58/my.cnf
#
# Configuration file for MyCluster
# Generated by mcc
#
[mysqld]
log-error=mysqld.58.err
datadir="/cluster/MySQL_Cluster/58/"
tmpdir="/cluster/MySQL_Cluster/58/tmp"
basedir="/cluster/mysql_cluster/"
port=3307
ndbcluster=on
ndb-nodeid=58
```

```
ndb-connectstring=192.168.22.105:1186,192.168.22.106:1186,
socket="/cluster/MySQL_Cluster/58/mysql.socket"
```

<note important> All **my.cnf** parameter names are **case-sensitive**. </note>

As with a traditional my.cnf file, these parameters can be split out into different sections so that they apply only to specific cluster nodes, for example:

Section	Description	Example
[mysql_cluster]	Contains settings for all nodes	ndb-connectstring=192.168.22.105:1186,192.168.22.106:1186
[ndbd]	Contains settings specific to data nodes	ndb-connectstring=192.168.22.105:1186,192.168.22.106:1186
[ndb_mgm]	Contains settings specific to management nodes	ndb-connectstring=192.168.22.105:1186,192.168.22.106:1186
[mysqld]	Contains settings specific to SQL nodes	ndbcluster=on

The definitions of the non-standard MySQL parameters in these files are as follows:

Section	Parameter	Description
[mysqld]	ndbcluster=	Enables/disables the ndbcluster storage engine
[mysqld]	ndb-nodeid=	Sets the node's MySQL Cluster node ID.
[mysqld]	ndb-connectstring=	Stipulates the connection string for the management server hosts (default port = 1186)

Connection Strings

Each MySQL or API node must use a Connection String to connect to the CLuster Management node. The Connection String is defined either by the **bind-address** parameter or the **ndb-connectstring** parameter (used for backward compatibility) in **my.cnf** or in a file called **Ndb.cfg**, placed in the mysqld's startup directory.

An example of the most complete form of a Connection String is:

```
nodeid=55,mymanagementhost:1186,192.168.22.105:1186,192.168.22.106:1186, ...
```

In the above example the MySQL Cluster Data node or API node will poll each host until a connection to a Cluster Management node is successful. As a result there is a Cluster Management node redundancy. In the example shown, the line can be split as follows:

```
NodeId, host-definition, host_definition, host-definition
```

<note important> If the Connection String is not defined, the default value of **localhost:1186** is used. </note>

The bind-address Parameter

If a bind address is defined before the host definitions, the same bind address is used for all of the host definitions:

```
bind-address=192.168.22.105, server1:1186, server2:1186
```

In order to poll two servers, the format of the line becomes:

```
server1:1186;bind-address=192.168.22.105, server2:1186;bind-address=192.168.22.106
```

Starting the Cluster

Starting the cluster consists firstly of starting the Cluster Management service on node1 and node2. The Cluster Management deamon **ndb_mgmd** can be invoked with many command line switches as shown below:

```
[root@node1 ~]# /cluster/mysql_cluster/bin/ndb_mgmd --help
MySQL Cluster Management Server mysql-5.6.23 ndb-7.4.4
Usage: /cluster/mysql_cluster/bin/ndb_mgmd [OPTIONS]
MySQL distrib mysql-5.6.23 ndb-7.4.4, for linux-glibc2.5 (i686)

Default options are read from the following files in the given order:
/etc/my.cnf /etc/mysql/my.cnf /usr/local/mysql/etc/my.cnf ~/.my.cnf
The following groups are read: mysql_cluster ndb_mgmd
The following options may be given as the first argument:
--print-defaults        Print the program argument list and exit.
--no-defaults           Don't read default options from any option file,
```

```
                        except for login file.
--defaults-file=#       Only read default options from the given file #.
--defaults-extra-file=# Read this file after the global files are read.
--defaults-group-suffix=#
                        Also read groups with concat(group, suffix)
--login-path=#          Read this path from the login file.

  -?, --usage         Display this help and exit.
  -?, --help          Display this help and exit.
  -V, --version       Output version information and exit.
  -c, --ndb-connectstring=name
                      Set connect string for connecting to ndb_mgmd. Syntax:
                      "[nodeid=<id>;][host=]<hostname>[:<port>]". Overrides
                      specifying entries in NDB_CONNECTSTRING and my.cnf
  --ndb-mgmd-host=name
                      same as --ndb-connectstring
  --ndb-nodeid=#      Set node id for this node. Overrides node id specified in
                      --ndb-connectstring.
  --ndb-optimized-node-selection
                      Select nodes for transactions in a more optimal way
                      (Defaults to on; use --skip-ndb-optimized-node-selection to disable.)
  -c, --connect-string=name
                      same as --ndb-connectstring
  --core-file         Write core on errors.
  --character-sets-dir=name
                      Directory where character sets are.
  -f, --config-file=name
                      Specify cluster configuration file
  -P, --print-full-config
                      Print full config and exit
  -d, --daemon        Run ndb_mgmd in daemon mode (default)
                      (Defaults to on; use --skip-daemon to disable.)
  --interactive       Run interactive. Not supported but provided for testing
                      purposes
```

```
  --no-nodeid-checks  Do not provide any node id checks
  --nodaemon          Don't run as daemon, but don't read from stdin
  --mycnf             Read cluster config from my.cnf
  --bind-address=name Local bind address
  --configdir=name    Directory for the binary configuration files (alias for
                      --config-dir)
  --config-dir=name   Directory for the binary configuration files
  --config-cache      Enable configuration cache and change management
                      (Defaults to on; use --skip-config-cache to disable.)
  -v, --verbose       Write more log messages
  --reload            Reload config from config.ini or my.cnf if it has changed
                      on startup
  --initial           Delete all binary config files and start from config.ini
                      or my.cnf
  --log-name=name     Name to use when logging messages for this node
  --nowait-nodes=name Nodes that will not be waited for during start

Variables (--variable-name=value)
and boolean options {FALSE|TRUE}  Value (after reading options)
--------------------------------- ----------------------------------------
ndb-connectstring                 (No default value)
ndb-mgmd-host                     (No default value)
ndb-nodeid                        0
ndb-optimized-node-selection      TRUE
connect-string                    (No default value)
core-file                         FALSE
character-sets-dir                (No default value)
config-file                       (No default value)
print-full-config                 FALSE
daemon                            TRUE
interactive                       FALSE
no-nodeid-checks                  FALSE
nodaemon                          FALSE
mycnf                             FALSE
```

```
bind-address                    (No default value)
configdir                       /usr/local/mysql/mysql-cluster
config-dir                      /usr/local/mysql/mysql-cluster
config-cache                    TRUE
verbose                         FALSE
reload                          FALSE
initial                         FALSE
log-name                        MgmtSrvr
nowait-nodes                    (No default value)
```

<note important> At the end of the above output we can see the default values for some of the switches. </note>

Starting the Cluster Management Nodes

Start the Management service on node1 by stipulating the –initial switch such that the config.ini and my.cnf files are read:

```
[root@node1 ~]# /cluster/mysql_cluster/bin/ndb_mgmd --initial --ndb-nodeid=49 --config-
dir=/cluster/MySQL_Cluster/49/ --config-file=/cluster/MySQL_Cluster/49/config.ini
MySQL Cluster Management Server mysql-5.6.23 ndb-7.4.4
```

Start the Management service on node2 by stipulating the –initial switch such that the config.ini and my.cnf files are read:

```
[root@node2 ~]# /cluster/mysql_cluster/bin/ndb_mgmd --initial --ndb-nodeid=52 --config-
dir=/cluster/MySQL_Cluster/52/ --config-file=/cluster/MySQL_Cluster/52/config.ini
MySQL Cluster Management Server mysql-5.6.23 ndb-7.4.4
```

Starting the SQL Nodes

Start the SQL node on node3:

```
[root@node3 ~]# /cluster/mysql_cluster/bin/ndbmtd --ndb-nodeid=1 --ndb-
```

```
connectstring=192.168.22.105:1186,192.168.22.106:1186
2015-03-02 09:26:14 [ndbd] INFO     -- Angel connected to '192.168.22.105:1186'
2015-03-02 09:26:14 [ndbd] INFO     -- Angel allocated nodeid: 1
```

Start the SQL node on node4:

```
[root@node4 ~]# /cluster/mysql_cluster/bin/ndbmtd --ndb-nodeid=2 --ndb-
connectstring=192.168.22.105:1186,192.168.22.106:1186
2015-03-02 09:26:49 [ndbd] INFO     -- Angel connected to '192.168.22.105:1186'
2015-03-02 09:26:49 [ndbd] INFO     -- Angel allocated nodeid: 2
```

<note important> Note that **ndbmtd** is the multi-threaded version of ndbd, used on host computers having multiple CPU cores. For further information about the differences between **ndbm** and **ndbmtd**, consult this
http://dev.mysql.com/doc/refman/5.6/en/mysql-cluster-programs-ndbmtd.html[link </note>

We now need to create the MySQL Server system tables on node1 and node2. In order to achieve this, Oracle supplies a bash script called **mysql_install_db** which can be found in **/cluster/mysql_cluster/scripts/** :

```
[root@node1 ~]# /cluster/mysql_cluster/scripts/mysql_install_db --no-defaults --basedir=/cluster/mysql_cluster/ -
-datadir=/cluster/MySQL_Cluster/55/ --user=root --explicit_defaults_for_timestamp
WARNING: The host 'node1.fenestros.loc' could not be looked up with /cluster/mysql_cluster//bin/resolveip.
This probably means that your libc libraries are not 100 % compatible
with this binary MySQL version. The MySQL daemon, mysqld, should work
normally with the exception that host name resolving will not work.
This means that you should use IP addresses instead of hostnames
when specifying MySQL privileges !

Installing MySQL system tables...2015-03-02 09:30:29 30822 [Note] InnoDB: Using mutexes to ref count buffer pool
pages
2015-03-02 09:30:29 30822 [Note] InnoDB: The InnoDB memory heap is disabled
2015-03-02 09:30:29 30822 [Note] InnoDB: Mutexes and rw_locks use InnoDB's own implementation
2015-03-02 09:30:29 30822 [Note] InnoDB: Memory barrier is not used
2015-03-02 09:30:29 30822 [Note] InnoDB: Compressed tables use zlib 1.2.3
2015-03-02 09:30:29 30822 [Note] InnoDB: Using Linux native AIO
```

```
2015-03-02 09:30:29 30822 [Note] InnoDB: Not using CPU crc32 instructions
2015-03-02 09:30:29 30822 [Note] InnoDB: Initializing buffer pool, size = 128.0M
2015-03-02 09:30:29 30822 [Note] InnoDB: Completed initialization of buffer pool
2015-03-02 09:30:29 30822 [Note] InnoDB: The first specified data file ./ibdata1 did not exist: a new database to
be created!
2015-03-02 09:30:29 30822 [Note] InnoDB: Setting file ./ibdata1 size to 12 MB
2015-03-02 09:30:29 30822 [Note] InnoDB: Database physically writes the file full: wait...
2015-03-02 09:30:30 30822 [Note] InnoDB: Setting log file ./ib_logfile101 size to 48 MB
2015-03-02 09:30:31 30822 [Note] InnoDB: Setting log file ./ib_logfile1 size to 48 MB
2015-03-02 09:30:32 30822 [Note] InnoDB: Renaming log file ./ib_logfile101 to ./ib_logfile0
2015-03-02 09:30:32 30822 [Warning] InnoDB: New log files created, LSN=45781
2015-03-02 09:30:32 30822 [Note] InnoDB: Doublewrite buffer not found: creating new
2015-03-02 09:30:32 30822 [Note] InnoDB: Doublewrite buffer created
2015-03-02 09:30:32 30822 [Note] InnoDB: 128 rollback segment(s) are active.
2015-03-02 09:30:32 30822 [Warning] InnoDB: Creating foreign key constraint system tables.
2015-03-02 09:30:32 30822 [Note] InnoDB: Foreign key constraint system tables created
2015-03-02 09:30:32 30822 [Note] InnoDB: Creating tablespace and datafile system tables.
2015-03-02 09:30:32 30822 [Note] InnoDB: Tablespace and datafile system tables created.
2015-03-02 09:30:32 30822 [Note] InnoDB: Waiting for purge to start
2015-03-02 09:30:32 30822 [Note] InnoDB: 5.6.23 started; log sequence number 0
2015-03-02 09:30:33 30822 [Note] Binlog end
2015-03-02 09:30:33 30822 [Note] InnoDB: FTS optimize thread exiting.
2015-03-02 09:30:33 30822 [Note] InnoDB: Starting shutdown...
2015-03-02 09:30:35 30822 [Note] InnoDB: Shutdown completed; log sequence number 1625977
OK

Filling help tables...2015-03-02 09:30:35 30845 [Note] InnoDB: Using mutexes to ref count buffer pool pages
2015-03-02 09:30:35 30845 [Note] InnoDB: The InnoDB memory heap is disabled
2015-03-02 09:30:35 30845 [Note] InnoDB: Mutexes and rw_locks use InnoDB's own implementation
2015-03-02 09:30:35 30845 [Note] InnoDB: Memory barrier is not used
2015-03-02 09:30:35 30845 [Note] InnoDB: Compressed tables use zlib 1.2.3
2015-03-02 09:30:35 30845 [Note] InnoDB: Using Linux native AIO
2015-03-02 09:30:35 30845 [Note] InnoDB: Not using CPU crc32 instructions
2015-03-02 09:30:35 30845 [Note] InnoDB: Initializing buffer pool, size = 128.0M
```

```
2015-03-02 09:30:35 30845 [Note] InnoDB: Completed initialization of buffer pool
2015-03-02 09:30:35 30845 [Note] InnoDB: Highest supported file format is Barracuda.
2015-03-02 09:30:35 30845 [Note] InnoDB: 128 rollback segment(s) are active.
2015-03-02 09:30:35 30845 [Note] InnoDB: Waiting for purge to start
2015-03-02 09:30:35 30845 [Note] InnoDB: 5.6.23 started; log sequence number 1625977
2015-03-02 09:30:36 30845 [Note] Binlog end
2015-03-02 09:30:36 30845 [Note] InnoDB: FTS optimize thread exiting.
2015-03-02 09:30:36 30845 [Note] InnoDB: Starting shutdown...
2015-03-02 09:30:37 30845 [Note] InnoDB: Shutdown completed; log sequence number 1625987
OK

To start mysqld at boot time you have to copy
support-files/mysql.server to the right place for your system

PLEASE REMEMBER TO SET A PASSWORD FOR THE MySQL root USER !
To do so, start the server, then issue the following commands:

  /cluster/mysql_cluster//bin/mysqladmin -u root password 'new-password'
  /cluster/mysql_cluster//bin/mysqladmin -u root -h node1.fenestros.loc password 'new-password'

Alternatively you can run:

  /cluster/mysql_cluster//bin/mysql_secure_installation

which will also give you the option of removing the test
databases and anonymous user created by default.  This is
strongly recommended for production servers.

See the manual for more instructions.

You can start the MySQL daemon with:

  cd . ; /cluster/mysql_cluster//bin/mysqld_safe &
```

```
You can test the MySQL daemon with mysql-test-run.pl

  cd mysql-test ; perl mysql-test-run.pl

Please report any problems at http://bugs.mysql.com/

The latest information about MySQL is available on the web at

  http://www.mysql.com

Support MySQL by buying support/licenses at http://shop.mysql.com

WARNING: Found existing config file /cluster/mysql_cluster//my.cnf on the system.
Because this file might be in use, it was not replaced,
but was used in bootstrap (unless you used --defaults-file)
and when you later start the server.
The new default config file was created as /cluster/mysql_cluster//my-new.cnf,
please compare it with your file and take the changes you need.

WARNING: Default config file /etc/my.cnf exists on the system
This file will be read by default by the MySQL server
If you do not want to use this, either remove it, or use the
--defaults-file argument to mysqld_safe when starting the server

[root@node1 ~]# /cluster/mysql_cluster/scripts/mysql_install_db --no-defaults --basedir=/cluster/mysql_cluster/ -
-datadir=/cluster/MySQL_Cluster/56/ --user=root --explicit_defaults_for_timestamp
WARNING: The host 'node1.fenestros.loc' could not be looked up with /cluster/mysql_cluster//bin/resolveip.
This probably means that your libc libraries are not 100 % compatible
with this binary MySQL version. The MySQL daemon, mysqld, should work
normally with the exception that host name resolving will not work.
This means that you should use IP addresses instead of hostnames
when specifying MySQL privileges !

Installing MySQL system tables...2015-03-02 09:31:54 30876 [Note] InnoDB: Using mutexes to ref count buffer pool
```

```
pages
2015-03-02 09:31:54 30876 [Note] InnoDB: The InnoDB memory heap is disabled
2015-03-02 09:31:54 30876 [Note] InnoDB: Mutexes and rw_locks use InnoDB's own implementation
2015-03-02 09:31:54 30876 [Note] InnoDB: Memory barrier is not used
2015-03-02 09:31:54 30876 [Note] InnoDB: Compressed tables use zlib 1.2.3
2015-03-02 09:31:54 30876 [Note] InnoDB: Using Linux native AIO
2015-03-02 09:31:54 30876 [Note] InnoDB: Not using CPU crc32 instructions
2015-03-02 09:31:54 30876 [Note] InnoDB: Initializing buffer pool, size = 128.0M
2015-03-02 09:31:55 30876 [Note] InnoDB: Completed initialization of buffer pool
2015-03-02 09:31:55 30876 [Note] InnoDB: The first specified data file ./ibdata1 did not exist: a new database to
be created!
2015-03-02 09:31:55 30876 [Note] InnoDB: Setting file ./ibdata1 size to 12 MB
2015-03-02 09:31:55 30876 [Note] InnoDB: Database physically writes the file full: wait...
2015-03-02 09:31:55 30876 [Note] InnoDB: Setting log file ./ib_logfile101 size to 48 MB
2015-03-02 09:31:57 30876 [Note] InnoDB: Setting log file ./ib_logfile1 size to 48 MB
2015-03-02 09:31:59 30876 [Note] InnoDB: Renaming log file ./ib_logfile101 to ./ib_logfile0
2015-03-02 09:31:59 30876 [Warning] InnoDB: New log files created, LSN=45781
2015-03-02 09:31:59 30876 [Note] InnoDB: Doublewrite buffer not found: creating new
2015-03-02 09:31:59 30876 [Note] InnoDB: Doublewrite buffer created
2015-03-02 09:31:59 30876 [Note] InnoDB: 128 rollback segment(s) are active.
2015-03-02 09:31:59 30876 [Warning] InnoDB: Creating foreign key constraint system tables.
2015-03-02 09:31:59 30876 [Note] InnoDB: Foreign key constraint system tables created
2015-03-02 09:31:59 30876 [Note] InnoDB: Creating tablespace and datafile system tables.
2015-03-02 09:31:59 30876 [Note] InnoDB: Tablespace and datafile system tables created.
2015-03-02 09:31:59 30876 [Note] InnoDB: Waiting for purge to start
2015-03-02 09:31:59 30876 [Note] InnoDB: 5.6.23 started; log sequence number 0
2015-03-02 09:31:59 30876 [Note] Binlog end
2015-03-02 09:31:59 30876 [Note] InnoDB: FTS optimize thread exiting.
2015-03-02 09:31:59 30876 [Note] InnoDB: Starting shutdown...
2015-03-02 09:32:01 30876 [Note] InnoDB: Shutdown completed; log sequence number 1625977
OK

Filling help tables...2015-03-02 09:32:01 30899 [Note] InnoDB: Using mutexes to ref count buffer pool pages
2015-03-02 09:32:01 30899 [Note] InnoDB: The InnoDB memory heap is disabled
```

```
2015-03-02 09:32:01 30899 [Note] InnoDB: Mutexes and rw_locks use InnoDB's own implementation
2015-03-02 09:32:01 30899 [Note] InnoDB: Memory barrier is not used
2015-03-02 09:32:01 30899 [Note] InnoDB: Compressed tables use zlib 1.2.3
2015-03-02 09:32:01 30899 [Note] InnoDB: Using Linux native AIO
2015-03-02 09:32:01 30899 [Note] InnoDB: Not using CPU crc32 instructions
2015-03-02 09:32:01 30899 [Note] InnoDB: Initializing buffer pool, size = 128.0M
2015-03-02 09:32:01 30899 [Note] InnoDB: Completed initialization of buffer pool
2015-03-02 09:32:01 30899 [Note] InnoDB: Highest supported file format is Barracuda.
2015-03-02 09:32:01 30899 [Note] InnoDB: 128 rollback segment(s) are active.
2015-03-02 09:32:01 30899 [Note] InnoDB: Waiting for purge to start
2015-03-02 09:32:01 30899 [Note] InnoDB: 5.6.23 started; log sequence number 1625977
2015-03-02 09:32:01 30899 [Note] Binlog end
2015-03-02 09:32:01 30899 [Note] InnoDB: FTS optimize thread exiting.
2015-03-02 09:32:01 30899 [Note] InnoDB: Starting shutdown...
2015-03-02 09:32:03 30899 [Note] InnoDB: Shutdown completed; log sequence number 1625987
OK

To start mysqld at boot time you have to copy
support-files/mysql.server to the right place for your system

PLEASE REMEMBER TO SET A PASSWORD FOR THE MySQL root USER !
To do so, start the server, then issue the following commands:

  /cluster/mysql_cluster//bin/mysqladmin -u root password 'new-password'
  /cluster/mysql_cluster//bin/mysqladmin -u root -h node1.fenestros.loc password 'new-password'

Alternatively you can run:

  /cluster/mysql_cluster//bin/mysql_secure_installation

which will also give you the option of removing the test
databases and anonymous user created by default.  This is
strongly recommended for production servers.
```

```
See the manual for more instructions.

You can start the MySQL daemon with:

  cd . ; /cluster/mysql_cluster//bin/mysqld_safe &

You can test the MySQL daemon with mysql-test-run.pl

  cd mysql-test ; perl mysql-test-run.pl

Please report any problems at http://bugs.mysql.com/

The latest information about MySQL is available on the web at

  http://www.mysql.com

Support MySQL by buying support/licenses at http://shop.mysql.com

WARNING: Found existing config file /cluster/mysql_cluster//my.cnf on the system.
Because this file might be in use, it was not replaced,
but was used in bootstrap (unless you used --defaults-file)
and when you later start the server.
The new default config file was created as /cluster/mysql_cluster//my-new.cnf,
please compare it with your file and take the changes you need.

WARNING: Default config file /etc/my.cnf exists on the system
This file will be read by default by the MySQL server
If you do not want to use this, either remove it, or use the
--defaults-file argument to mysqld_safe when starting the server
```

```
[root@node2 ~]# /cluster/mysql_cluster/scripts/mysql_install_db --no-defaults --basedir=/cluster/mysql_cluster/ --datadir=/cluster/MySQL_Cluster/57/ --user=root --explicit_defaults_for_timestamp
[root@node2 ~]# /cluster/mysql_cluster/scripts/mysql_install_db --no-defaults --basedir=/cluster/mysql_cluster/ --datadir=/cluster/MySQL_Cluster/58/ --user=root --explicit_defaults_for_timestamp
```

Now we need to start the SQL nodes on node1:

```
[root@node1 ~]# /cluster/mysql_cluster/bin/mysqld --defaults-file=/cluster/MySQL_Cluster/55/my.cnf --
basedir=/cluster/mysql_cluster/ --datadir=/cluster/MySQL_Cluster/55/ --user=root --
explicit_defaults_for_timestamp &
[1] 31287
[root@node1 ~]#  /cluster/mysql_cluster/bin/mysqld --defaults-file=/cluster/MySQL_Cluster/56/my.cnf --
basedir=/cluster/mysql_cluster/ --datadir=/cluster/MySQL_Cluster/56/ --user=root --
explicit_defaults_for_timestamp &
[2] 31319
[root@node1 ~]# ps aux | grep mysqld
root     31287  6.2  0.4 811932  1728 pts/1    Sl   10:22   0:01 /cluster/mysql_cluster/bin/mysqld --defaults-
file=/cluster/MySQL_Cluster/55/my.cnf --basedir=/cluster/mysql_cluster/ --datadir=/cluster/MySQL_Cluster/55/ --
user=root --explicit_defaults_for_timestamp
root     31319 12.4 73.9 812040 280812 pts/1   Sl   10:22   0:00 /cluster/mysql_cluster/bin/mysqld --defaults-
file=/cluster/MySQL_Cluster/56/my.cnf --basedir=/cluster/mysql_cluster/ --datadir=/cluster/MySQL_Cluster/56/ --
user=root --explicit_defaults_for_timestamp
root     31352  0.0  0.1   4352   732 pts/1    S+   10:22   0:00 grep mysqld
```

Equally we need to do the same on node2:

```
[root@node2 ~]# /cluster/mysql_cluster/bin/mysqld --defaults-file=/cluster/MySQL_Cluster/57/my.cnf --
basedir=/cluster/mysql_cluster/ --datadir=/cluster/MySQL_Cluster/57/ --user=root --
explicit_defaults_for_timestamp &
[1] 24817
[root@node2 ~]# /cluster/mysql_cluster/bin/mysqld --defaults-file=/cluster/MySQL_Cluster/58/my.cnf --
basedir=/cluster/mysql_cluster/ --datadir=/cluster/MySQL_Cluster/58/ --user=root --
explicit_defaults_for_timestamp &
[2] 24847
[root@node2 ~]# ps aux | grep mysqld
root     24817  4.1  7.0 811728 26688 pts/0    Sl   10:25   0:01 /cluster/mysql_cluster/bin/mysqld --defaults-
file=/cluster/MySQL_Cluster/57/my.cnf --basedir=/cluster/mysql_cluster/ --datadir=/cluster/MySQL_Cluster/57/ --
user=root --explicit_defaults_for_timestamp
root     24847  3.0 56.4 288124 214500 pts/0   D    10:25   0:00 /cluster/mysql_cluster/bin/mysqld --defaults-
```

```
file=/cluster/MySQL_Cluster/58/my.cnf --basedir=/cluster/mysql_cluster/ --datadir=/cluster/MySQL_Cluster/58/ --
user=root --explicit_defaults_for_timestamp
root     24851  0.3  0.1   4352   504 pts/0    S+   10:26   0:00 grep mysqld
```

<note important> You can find more information about the **–explicit_defaults_for_timestamp** switch **here**. </note>

Check the Cluster Status

You can use the NDB Cluster Management Client, **ndb_mgm**, from either node1 or node2 to check the cluster status:

```
[root@node1 ~]# /cluster/mysql_cluster/bin/ndb_mgm
-- NDB Cluster -- Management Client --
ndb_mgm> SHOW
Connected to Management Server at: localhost:1186
Cluster Configuration
---------------------
[ndbd(NDB)] 2 node(s)
id=1    @192.168.22.107  (mysql-5.6.23 ndb-7.4.4, Nodegroup: 0, *)
id=2    @192.168.22.108  (mysql-5.6.23 ndb-7.4.4, Nodegroup: 0)

[ndb_mgmd(MGM)] 2 node(s)
id=49   @192.168.22.105  (mysql-5.6.23 ndb-7.4.4)
id=52   @192.168.22.106  (mysql-5.6.23 ndb-7.4.4)

[mysqld(API)]   4 node(s)
id=55   @192.168.22.105  (mysql-5.6.23 ndb-7.4.4)
id=56   @192.168.22.105  (mysql-5.6.23 ndb-7.4.4)
id=57   @192.168.22.106  (mysql-5.6.23 ndb-7.4.4)
id=58   @192.168.22.106  (mysql-5.6.23 ndb-7.4.4)

ndb_mgm> exit
[root@node1 ~]#
```

```
[root@node2 ~]# /cluster/mysql_cluster/bin/ndb_mgm
-- NDB Cluster -- Management Client --
ndb_mgm> SHOW
Connected to Management Server at: localhost:1186
Cluster Configuration
---------------------
[ndbd(NDB)] 2 node(s)
id=1    @192.168.22.107  (mysql-5.6.23 ndb-7.4.4, Nodegroup: 0, *)
id=2    @192.168.22.108  (mysql-5.6.23 ndb-7.4.4, Nodegroup: 0)

[ndb_mgmd(MGM)] 2 node(s)
id=49   @192.168.22.105  (mysql-5.6.23 ndb-7.4.4)
id=52   @192.168.22.106  (mysql-5.6.23 ndb-7.4.4)

[mysqld(API)]   4 node(s)
id=55   @192.168.22.105  (mysql-5.6.23 ndb-7.4.4)
id=56   @192.168.22.105  (mysql-5.6.23 ndb-7.4.4)
id=57   @192.168.22.106  (mysql-5.6.23 ndb-7.4.4)
id=58   @192.168.22.106  (mysql-5.6.23 ndb-7.4.4)

ndb_mgm> exit
[root@node2 ~]#
```

The ndb_mgm command can also be used as follows:

```
[root@node1 ~]# ndb_mgm -e "SHOW"
-bash: ndb_mgm: command not found
[root@node1 ~]# /cluster/mysql_cluster/bin/ndb_mgm -e "SHOW"
Connected to Management Server at: localhost:1186
Cluster Configuration
---------------------
[ndbd(NDB)] 2 node(s)
id=1    @192.168.22.107  (mysql-5.6.23 ndb-7.4.4, Nodegroup: 0, *)
id=2    @192.168.22.108  (mysql-5.6.23 ndb-7.4.4, Nodegroup: 0)
```

```
[ndb_mgmd(MGM)] 2 node(s)
id=49   @192.168.22.105  (mysql-5.6.23 ndb-7.4.4)
id=52   @192.168.22.106  (mysql-5.6.23 ndb-7.4.4)

[mysqld(API)]   4 node(s)
id=55   @192.168.22.105  (mysql-5.6.23 ndb-7.4.4)
id=56   @192.168.22.105  (mysql-5.6.23 ndb-7.4.4)
id=57   @192.168.22.106  (mysql-5.6.23 ndb-7.4.4)
id=58   @192.168.22.106  (mysql-5.6.23 ndb-7.4.4)

[root@node1 ~]# /cluster/mysql_cluster/bin/ndb_mgm --execute "SHOW"
Connected to Management Server at: localhost:1186
Cluster Configuration
---------------------
[ndbd(NDB)] 2 node(s)
id=1    @192.168.22.107  (mysql-5.6.23 ndb-7.4.4, Nodegroup: 0, *)
id=2    @192.168.22.108  (mysql-5.6.23 ndb-7.4.4, Nodegroup: 0)

[ndb_mgmd(MGM)] 2 node(s)
id=49   @192.168.22.105  (mysql-5.6.23 ndb-7.4.4)
id=52   @192.168.22.106  (mysql-5.6.23 ndb-7.4.4)

[mysqld(API)]   4 node(s)
id=55   @192.168.22.105  (mysql-5.6.23 ndb-7.4.4)
id=56   @192.168.22.105  (mysql-5.6.23 ndb-7.4.4)
id=57   @192.168.22.106  (mysql-5.6.23 ndb-7.4.4)
id=58   @192.168.22.106  (mysql-5.6.23 ndb-7.4.4)
```

The command line switches of the nbd_mgm command are:

```
[root@node2 ~]# /cluster/mysql_cluster/bin/ndb_mgm --help
Usage: /cluster/mysql_cluster/bin/ndb_mgm [OPTIONS] [hostname [port]]
MySQL distrib mysql-5.6.23 ndb-7.4.4, for linux-glibc2.5 (i686)
```

```
Default options are read from the following files in the given order:
/etc/my.cnf /etc/mysql/my.cnf /usr/local/mysql/etc/my.cnf ~/.my.cnf
The following groups are read: mysql_cluster ndb_mgm
The following options may be given as the first argument:
--print-defaults        Print the program argument list and exit.
--no-defaults           Don't read default options from any option file,
                        except for login file.
--defaults-file=#       Only read default options from the given file #.
--defaults-extra-file=# Read this file after the global files are read.
--defaults-group-suffix=#
                        Also read groups with concat(group, suffix)
--login-path=#          Read this path from the login file.

  -?, --usage         Display this help and exit.
  -?, --help          Display this help and exit.
  -V, --version       Output version information and exit.
  -c, --ndb-connectstring=name
                      Set connect string for connecting to ndb_mgmd. Syntax:
                      "[nodeid=<id>;][host=]<hostname>[:<port>]". Overrides
                      specifying entries in NDB_CONNECTSTRING and my.cnf
  --ndb-mgmd-host=name
                      same as --ndb-connectstring
  --ndb-nodeid=#      Set node id for this node. Overrides node id specified in
                      --ndb-connectstring.
  --ndb-optimized-node-selection
                      Select nodes for transactions in a more optimal way
                      (Defaults to on; use --skip-ndb-optimized-node-selection to disable.)
  -c, --connect-string=name
                      same as --ndb-connectstring
  --core-file         Write core on errors.
  --character-sets-dir=name
                      Directory where character sets are.
  -e, --execute=name  execute command and exit
  -t, --try-reconnect=#
```

```
                      Specify number of tries for connecting to ndb_mgmd (0 =
                      infinite)
  -v, --verbose=#     Control the amount of printout

Variables (--variable-name=value)
and boolean options {FALSE|TRUE}  Value (after reading options)
--------------------------------- ----------------------------------------
ndb-connectstring                 (No default value)
ndb-mgmd-host                     (No default value)
ndb-nodeid                        0
ndb-optimized-node-selection      TRUE
connect-string                    (No default value)
core-file                         FALSE
character-sets-dir                (No default value)
execute                           (No default value)
try-reconnect                     3
verbose                           1
```

Stopping and Starting the Cluster

Graceful Shutdown

To perform a graceful shutdown of the cluster, we need to use the **ndb-adm** command on a host running a management node. This command causes the ndb_mgm, ndb_mgmd, and any ndbd or ndbmtd processes to terminate gracefully as shown below:

```
[root@node1 ~]# /cluster/mysql_cluster/bin/ndb_mgm -e shutdown
Connected to Management Server at: localhost:1186
4 NDB Cluster node(s) have shutdown.
Disconnecting to allow management server to shutdown.
```

Note however that the mysqld servers are not shutdown:

```
[root@node1 ~]# ps aux | grep ndb_mgmd
root      3410  0.0  0.1   4352   732 pts/0    S+   13:37   0:00 grep ndb_mgmd
[root@node1 ~]# ps aux | grep mysqld
root      2983  1.7  0.7 812036  2908 pts/0    Sl   12:41   1:01 /cluster/mysql_cluster/bin/mysqld --defaults-
file=/cluster/MySQL_Cluster/55/my.cnf --basedir=/cluster/mysql_cluster/ --datadir=/cluster/MySQL_Cluster/55/ --
user=root --explicit_defaults_for_timestamp
root      3016  1.7 37.8 811548 143712 pts/0   Sl   12:41   1:00 /cluster/mysql_cluster/bin/mysqld --defaults-
file=/cluster/MySQL_Cluster/56/my.cnf --basedir=/cluster/mysql_cluster/ --datadir=/cluster/MySQL_Cluster/56/ --
user=root --explicit_defaults_for_timestamp
root      3412  0.0  0.1   4352   732 pts/0    S+   13:38   0:00 grep mysqld
```

```
[root@node2 ~]# ps aux | grep ndb_mgmd
root      2726  0.0  0.1   4352   724 pts/1    S+   13:39   0:00 grep ndb_mgmd
[root@node2 ~]# ps aux | grep mysqld
root      2348  1.8  0.8 832032  3128 ?        Sl   12:41   1:04 /cluster/mysql_cluster/bin/mysqld --defaults-
file=/cluster/MySQL_Cluster/57/my.cnf --basedir=/cluster/mysql_cluster/ --datadir=/cluster/MySQL_Cluster/57/ --
user=root --explicit_defaults_for_timestamp
root      2382  1.8  1.1 811524  4456 ?        Sl   12:42   1:03 /cluster/mysql_cluster/bin/mysqld --defaults-
file=/cluster/MySQL_Cluster/58/my.cnf --basedir=/cluster/mysql_cluster/ --datadir=/cluster/MySQL_Cluster/58/ --
user=root --explicit_defaults_for_timestamp
root      2729  0.0  0.1   4356   736 pts/1    S+   13:39   0:00 grep mysqld
```

As opposed to the ndbd nodes which have been shutdown:

```
[root@node3 ~]# ps aux | grep ndbdtd
root      2870  0.0  0.0   4352   724 pts/0    S+   13:42   0:00 grep ndbdtd
```

```
[root@node4 ~]# ps aux | grep ndbdtd
root      2887  0.0  0.0   4352   724 pts/0    S+   13:43   0:00 grep ndbdtd
```

It therefore necessary to manually shutdown the mysqld nodes as shown below:

```
[root@node1 ~]# /cluster/mysql_cluster/bin/mysqladmin --socket='/cluster/MySQL_Cluster/55/mysql.socket' shutdown
```

```
[1]-  Done                    /cluster/mysql_cluster/bin/mysqld --defaults-file=/cluster/MySQL_Cluster/55/my.cnf
--basedir=/cluster/mysql_cluster/ --datadir=/cluster/MySQL_Cluster/55/ --user=root --
explicit_defaults_for_timestamp
[root@node1 ~]# /cluster/mysql_cluster/bin/mysqladmin --socket='/cluster/MySQL_Cluster/56/mysql.socket' shutdown
[2]+  Done                    /cluster/mysql_cluster/bin/mysqld --defaults-file=/cluster/MySQL_Cluster/56/my.cnf
--basedir=/cluster/mysql_cluster/ --datadir=/cluster/MySQL_Cluster/56/ --user=root --
explicit_defaults_for_timestamp
[root@node1 ~]# ps aux | grep mysqld
root      3454  0.0  0.1   4352   728 pts/0    S+   13:49   0:00 grep mysqld
```

```
[root@node2 ~]# /cluster/mysql_cluster/bin/mysqladmin --socket='/cluster/MySQL_Cluster/57/mysql.socket' shutdown
[root@node2 ~]# ps aux | grep mysqld
root      2382  2.0  1.1 811524  4456 ?        Sl   12:42   1:28 /cluster/mysql_cluster/bin/mysqld --defaults-
file=/cluster/MySQL_Cluster/58/my.cnf --basedir=/cluster/mysql_cluster/ --datadir=/cluster/MySQL_Cluster/58/ --
user=root --explicit_defaults_for_timestamp
root      2779  0.0  0.1   4356   736 pts/1    S+   13:53   0:00 grep mysqld
[root@node2 ~]# /cluster/mysql_cluster/bin/mysqladmin --socket='/cluster/MySQL_Cluster/58/mysql.socket' shutdown
[root@node2 ~]# ps aux | grep mysqld
root      2784  0.0  0.1   4352   728 pts/1    S+   13:53   0:00 grep mysqld
```

Starting the Cluster

Starting the cluster consists of first starting the Management nodes:

```
[root@node1 ~]# /cluster/mysql_cluster/bin/ndb_mgmd -f /cluster/MySQL_Cluster/49/config.ini --config-
dir=/cluster/MySQL_Cluster/49/
MySQL Cluster Management Server mysql-5.6.23 ndb-7.4.4
[root@node1 ~]# /cluster/mysql_cluster/bin/ndb_mgm -e SHOW
Connected to Management Server at: localhost:1186
Cluster Configuration
---------------------
[ndbd(NDB)] 2 node(s)
id=1 (not connected, accepting connect from 192.168.22.107)
```

```
id=2 (not connected, accepting connect from 192.168.22.108)

[ndb_mgmd(MGM)] 2 node(s)
id=49   @192.168.22.105  (mysql-5.6.23 ndb-7.4.4)
id=52 (not connected, accepting connect from 192.168.22.106)

[mysqld(API)]   4 node(s)
id=55 (not connected, accepting connect from 192.168.22.105)
id=56 (not connected, accepting connect from 192.168.22.105)
id=57 (not connected, accepting connect from 192.168.22.106)
id=58 (not connected, accepting connect from 192.168.22.106)
```

```
[root@node2 ~]# /cluster/mysql_cluster/bin/ndb_mgmd -f /cluster/MySQL_Cluster/52/config.ini --config-
dir=/cluster/MySQL_Cluster/52/
MySQL Cluster Management Server mysql-5.6.23 ndb-7.4.4
[root@node2 ~]# /cluster/mysql_cluster/bin/ndb_mgm -e SHOW
Connected to Management Server at: localhost:1186
Cluster Configuration
---------------------
[ndbd(NDB)] 2 node(s)
id=1 (not connected, accepting connect from 192.168.22.107)
id=2 (not connected, accepting connect from 192.168.22.108)

[ndb_mgmd(MGM)] 2 node(s)
id=49   @192.168.22.105  (mysql-5.6.23 ndb-7.4.4)
id=52   @192.168.22.106  (mysql-5.6.23 ndb-7.4.4)

[mysqld(API)]   4 node(s)
id=55 (not connected, accepting connect from 192.168.22.105)
id=56 (not connected, accepting connect from 192.168.22.105)
id=57 (not connected, accepting connect from 192.168.22.106)
id=58 (not connected, accepting connect from 192.168.22.106)
```

We now need to start the Data nodes:

```
[root@node3 ~]# /cluster/mysql_cluster/bin/ndbmtd --ndb-connectstring=192.168.22.105:1186,192.168.22.106:1186
2015-03-02 14:41:18 [ndbd] INFO     -- Angel connected to '192.168.22.105:1186'
2015-03-02 14:41:18 [ndbd] INFO     -- Angel allocated nodeid: 1
```

```
[root@node4 ~]# /cluster/mysql_cluster/bin/ndbmtd --ndb-connectstring=192.168.22.105:1186,192.168.22.106:1186
2015-03-02 14:41:22 [ndbd] INFO     -- Angel connected to '192.168.22.105:1186'
2015-03-02 14:41:22 [ndbd] INFO     -- Angel allocated nodeid: 2
```

Now we can check the cluster status from the node4 host:

```
[root@node4 ~]# /cluster/mysql_cluster/bin/ndb_mgm --ndb-connectstring=192.168.22.105:1186,192.168.22.106:1186 -e
SHOW
Connected to Management Server at: 192.168.22.105:1186
Cluster Configuration
---------------------
[ndbd(NDB)] 2 node(s)
id=1    @192.168.22.107  (mysql-5.6.23 ndb-7.4.4, Nodegroup: 0)
id=2    @192.168.22.108  (mysql-5.6.23 ndb-7.4.4, Nodegroup: 0, *)

[ndb_mgmd(MGM)] 2 node(s)
id=49   @192.168.22.105  (mysql-5.6.23 ndb-7.4.4)
id=52   @192.168.22.106  (mysql-5.6.23 ndb-7.4.4)

[mysqld(API)]   4 node(s)
id=55 (not connected, accepting connect from 192.168.22.105)
id=56 (not connected, accepting connect from 192.168.22.105)
id=57 (not connected, accepting connect from 192.168.22.106)
id=58 (not connected, accepting connect from 192.168.22.106)
```

Finally we need to start the SQL nodes:

```
[root@node1 ~]# /cluster/mysql_cluster/bin/mysqld --defaults-file=/cluster/MySQL_Cluster/55/my.cnf --
basedir=/cluster/mysql_cluster/ --datadir=/cluster/MySQL_Cluster/55/ --user=root --
```

```
explicit_defaults_for_timestamp &
[1] 4134
[root@node1 ~]# /cluster/mysql_cluster/bin/mysqld --defaults-file=/cluster/MySQL_Cluster/56/my.cnf --
basedir=/cluster/mysql_cluster/ --datadir=/cluster/MySQL_Cluster/56/ --user=root --
explicit_defaults_for_timestamp &
[2] 4167
[root@node1 ~]# ps aux | grep mysqld
root      4134  4.3  0.5 812040  2020 pts/0    Sl   14:43   0:01 /cluster/mysql_cluster/bin/mysqld --defaults-
file=/cluster/MySQL_Cluster/55/my.cnf --basedir=/cluster/mysql_cluster/ --datadir=/cluster/MySQL_Cluster/55/ --
user=root --explicit_defaults_for_timestamp
root      4167  4.9 74.6 812040 283320 pts/0   Sl   14:43   0:00 /cluster/mysql_cluster/bin/mysqld --defaults-
file=/cluster/MySQL_Cluster/56/my.cnf --basedir=/cluster/mysql_cluster/ --datadir=/cluster/MySQL_Cluster/56/ --
user=root --explicit_defaults_for_timestamp
root      4200  0.0  0.1   4352   728 pts/0    S+   14:44   0:00 grep mysqld
[root@node1 ~]# /cluster/mysql_cluster/bin/ndb_mgm -e SHOW
Connected to Management Server at: localhost:1186
Cluster Configuration
---------------------
[ndbd(NDB)] 2 node(s)
id=1    @192.168.22.107  (mysql-5.6.23 ndb-7.4.4, Nodegroup: 0)
id=2    @192.168.22.108  (mysql-5.6.23 ndb-7.4.4, Nodegroup: 0, *)

[ndb_mgmd(MGM)] 2 node(s)
id=49   @192.168.22.105  (mysql-5.6.23 ndb-7.4.4)
id=52   @192.168.22.106  (mysql-5.6.23 ndb-7.4.4)

[mysqld(API)]   4 node(s)
id=55   @192.168.22.105  (mysql-5.6.23 ndb-7.4.4)
id=56   @192.168.22.105  (mysql-5.6.23 ndb-7.4.4)
id=57 (not connected, accepting connect from 192.168.22.106)
id=58 (not connected, accepting connect from 192.168.22.106)
```

```
[root@node2 ~]# /cluster/mysql_cluster/bin/mysqld --defaults-file=/cluster/MySQL_Cluster/57/my.cnf --
basedir=/cluster/mysql_cluster/ --datadir=/cluster/MySQL_Cluster/57/ --user=root --
```

```
explicit_defaults_for_timestamp &
[1] 2987
[root@node2 ~]# /cluster/mysql_cluster/bin/mysqld --defaults-file=/cluster/MySQL_Cluster/58/my.cnf --
basedir=/cluster/mysql_cluster/ --datadir=/cluster/MySQL_Cluster/58/ --user=root --
explicit_defaults_for_timestamp &
[2] 3019
[root@node2 ~]# ps aux | grep mysqld
root      2987  4.2  0.8 812040  3384 pts/1    Sl   14:45   0:01 /cluster/mysql_cluster/bin/mysqld --defaults-
file=/cluster/MySQL_Cluster/57/my.cnf --basedir=/cluster/mysql_cluster/ --datadir=/cluster/MySQL_Cluster/57/ --
user=root --explicit_defaults_for_timestamp
root      3019  7.0 80.5 812040 305976 pts/1   Sl   14:45   0:00 /cluster/mysql_cluster/bin/mysqld --defaults-
file=/cluster/MySQL_Cluster/58/my.cnf --basedir=/cluster/mysql_cluster/ --datadir=/cluster/MySQL_Cluster/58/ --
user=root --explicit_defaults_for_timestamp
root      3052  0.5  0.1   4352   748 pts/1    S+   14:45   0:00 grep mysqld
[root@node2 ~]# /cluster/mysql_cluster/bin/ndb_mgm -e SHOW
Connected to Management Server at: localhost:1186
Cluster Configuration
---------------------
[ndbd(NDB)] 2 node(s)
id=1    @192.168.22.107  (mysql-5.6.23 ndb-7.4.4, Nodegroup: 0)
id=2    @192.168.22.108  (mysql-5.6.23 ndb-7.4.4, Nodegroup: 0, *)

[ndb_mgmd(MGM)] 2 node(s)
id=49   @192.168.22.105  (mysql-5.6.23 ndb-7.4.4)
id=52   @192.168.22.106  (mysql-5.6.23 ndb-7.4.4)

[mysqld(API)]   4 node(s)
id=55   @192.168.22.105  (mysql-5.6.23 ndb-7.4.4)
id=56   @192.168.22.105  (mysql-5.6.23 ndb-7.4.4)
id=57   @192.168.22.106  (mysql-5.6.23 ndb-7.4.4)
id=58   @192.168.22.106  (mysql-5.6.23 ndb-7.4.4)
```

<note important> It is obviously possible to perform a **rolling restart** to ensure high availability. This process is detailed **here** </note>

Logging and Debugging

MySQL Logs

These are generated in their respective node directories. They are of a traditional MySQL format and detail both successful operations and errors:

```
[root@node1 ~]# tail /cluster/MySQL_Cluster/55/mysqld.55.err
2015-03-02 14:43:47 4134 [Note] NDB Binlog: starting log at epoch 5130/3
2015-03-02 14:43:47 4134 [Note] NDB Binlog: Got first event
2015-03-02 14:43:47 4134 [Note] NDB Binlog: ndb tables writable
2015-03-02 14:43:47 4134 [Note] NDB Binlog: Startup and setup completed
2015-03-02 14:44:09 4134 [Note] NDB Schema dist: Data node: 1 reports subscribe from node 56, subscriber bitmask
10000000
2015-03-02 14:44:09 4134 [Note] NDB Schema dist: Data node: 2 reports subscribe from node 56, subscriber bitmask
10000000
2015-03-02 14:45:32 4134 [Note] NDB Schema dist: Data node: 1 reports subscribe from node 57, subscriber bitmask
30000000
2015-03-02 14:45:32 4134 [Note] NDB Schema dist: Data node: 2 reports subscribe from node 57, subscriber bitmask
30000000
2015-03-02 14:45:50 4134 [Note] NDB Schema dist: Data node: 1 reports subscribe from node 58, subscriber bitmask
70000000
2015-03-02 14:45:50 4134 [Note] NDB Schema dist: Data node: 2 reports subscribe from node 58, subscriber bitmask
70000000
[root@node1 ~]# tail /cluster/MySQL_Cluster/56/mysqld.56.err
2015-03-02 14:44:10 4167 [Note] NDB Binlog: starting log at epoch 5140/19
2015-03-02 14:44:10 4167 [Note] NDB Binlog: Got first event
2015-03-02 14:44:10 4167 [Note] NDB Binlog: ndb tables writable
2015-03-02 14:44:10 4167 [Note] NDB Binlog: Startup and setup completed
2015-03-02 14:44:10 4167 [Note] NDB Schema dist: Data node: 2 reports subscribe from node 55, subscriber bitmask
8000000
2015-03-02 14:44:10 4167 [Note] NDB Schema dist: Data node: 1 reports subscribe from node 55, subscriber bitmask
8000000
```

```
2015-03-02 14:45:32 4167 [Note] NDB Schema dist: Data node: 1 reports subscribe from node 57, subscriber bitmask
28000000
2015-03-02 14:45:32 4167 [Note] NDB Schema dist: Data node: 2 reports subscribe from node 57, subscriber bitmask
28000000
2015-03-02 14:45:50 4167 [Note] NDB Schema dist: Data node: 1 reports subscribe from node 58, subscriber bitmask
68000000
2015-03-02 14:45:50 4167 [Note] NDB Schema dist: Data node: 2 reports subscribe from node 58, subscriber bitmask
68000000
```

```
[root@node2 ~]# tail /cluster/MySQL_Cluster/57/mysqld.57.err
2015-03-02 14:45:32 2987 [Note] NDB Binlog: starting log at epoch 5180/16
2015-03-02 14:45:32 2987 [Note] NDB Binlog: Got first event
2015-03-02 14:45:32 2987 [Note] NDB Binlog: ndb tables writable
2015-03-02 14:45:32 2987 [Note] NDB Binlog: Startup and setup completed
2015-03-02 14:45:32 2987 [Note] NDB Schema dist: Data node: 1 reports subscribe from node 56, subscriber bitmask
10000000
2015-03-02 14:45:32 2987 [Note] NDB Schema dist: Data node: 1 reports subscribe from node 55, subscriber bitmask
18000000
2015-03-02 14:45:32 2987 [Note] NDB Schema dist: Data node: 2 reports subscribe from node 56, subscriber bitmask
10000000
2015-03-02 14:45:32 2987 [Note] NDB Schema dist: Data node: 2 reports subscribe from node 55, subscriber bitmask
18000000
2015-03-02 14:45:50 2987 [Note] NDB Schema dist: Data node: 1 reports subscribe from node 58, subscriber bitmask
58000000
2015-03-02 14:45:50 2987 [Note] NDB Schema dist: Data node: 2 reports subscribe from node 58, subscriber bitmask
58000000
[root@node2 ~]# tail /cluster/MySQL_Cluster/58/mysqld.58.err
2015-03-02 14:45:51 3019 [Note] NDB Binlog: starting log at epoch 5189/15
2015-03-02 14:45:51 3019 [Note] NDB Binlog: Got first event
2015-03-02 14:45:51 3019 [Note] NDB Binlog: ndb tables writable
2015-03-02 14:45:51 3019 [Note] NDB Binlog: Startup and setup completed
2015-03-02 14:45:51 3019 [Note] NDB Schema dist: Data node: 1 reports subscribe from node 57, subscriber bitmask
20000000
2015-03-02 14:45:51 3019 [Note] NDB Schema dist: Data node: 1 reports subscribe from node 56, subscriber bitmask
```

```
30000000
2015-03-02 14:45:51 3019 [Note] NDB Schema dist: Data node: 1 reports subscribe from node 55, subscriber bitmask
38000000
2015-03-02 14:45:51 3019 [Note] NDB Schema dist: Data node: 2 reports subscribe from node 57, subscriber bitmask
20000000
2015-03-02 14:45:51 3019 [Note] NDB Schema dist: Data node: 2 reports subscribe from node 56, subscriber bitmask
30000000
2015-03-02 14:45:51 3019 [Note] NDB Schema dist: Data node: 2 reports subscribe from node 55, subscriber bitmask
38000000
```

Cluster Logs

By default the cluster log is saved to a file named ndb_**node_id**_cluster.log situated in the Management node's data directory:

```
[root@node1 ~]# locate *_cluster.log
/cluster/MySQL_Cluster/49/ndb_49_cluster.log
```

```
[root@node2 ~]# locate *_cluster.log
/cluster/MySQL_Cluster/52/ndb_52_cluster.log
```

<note important> This location can be changed by the **DataDir** parameter. The use of Syslog is made possible by the **LogDestination** parameter. The cluster log provides logging information for the entire cluster in a single location. </note>

For example:

```
[root@node1 ~]# tail /cluster/MySQL_Cluster/49/ndb_49_cluster.log
2015-03-02 15:39:37 [MgmtSrvr] INFO     -- Node 2: Data usage is 71%(23 32K pages of total 32)
2015-03-02 15:39:37 [MgmtSrvr] INFO     -- Node 2: Index usage is 12%(20 8K pages of total 160)
2015-03-02 15:40:03 [MgmtSrvr] INFO     -- Node 1: Data usage is 71%(23 32K pages of total 32)
2015-03-02 15:40:03 [MgmtSrvr] INFO     -- Node 1: Index usage is 12%(20 8K pages of total 160)
2015-03-02 15:40:08 [MgmtSrvr] INFO     -- Node 2: Data usage is 71%(23 32K pages of total 32)
2015-03-02 15:40:08 [MgmtSrvr] INFO     -- Node 2: Index usage is 12%(20 8K pages of total 160)
2015-03-02 15:40:34 [MgmtSrvr] INFO     -- Node 1: Data usage is 71%(23 32K pages of total 32)
```

```
2015-03-02 15:40:34 [MgmtSrvr] INFO     -- Node 1: Index usage is 12%(20 8K pages of total 160)
2015-03-02 15:40:38 [MgmtSrvr] INFO     -- Node 2: Data usage is 71%(23 32K pages of total 32)
2015-03-02 15:40:38 [MgmtSrvr] INFO     -- Node 2: Index usage is 12%(20 8K pages of total 160)
```

```
[root@node2 ~]# tail /cluster/MySQL_Cluster/52/ndb_52_cluster.log
2015-03-02 15:39:37 [MgmtSrvr] INFO     -- Node 2: Data usage is 71%(23 32K pages of total 32)
2015-03-02 15:39:37 [MgmtSrvr] INFO     -- Node 2: Index usage is 12%(20 8K pages of total 160)
2015-03-02 15:40:03 [MgmtSrvr] INFO     -- Node 1: Data usage is 71%(23 32K pages of total 32)
2015-03-02 15:40:03 [MgmtSrvr] INFO     -- Node 1: Index usage is 12%(20 8K pages of total 160)
2015-03-02 15:40:07 [MgmtSrvr] INFO     -- Node 2: Data usage is 71%(23 32K pages of total 32)
2015-03-02 15:40:07 [MgmtSrvr] INFO     -- Node 2: Index usage is 12%(20 8K pages of total 160)
2015-03-02 15:40:34 [MgmtSrvr] INFO     -- Node 1: Data usage is 71%(23 32K pages of total 32)
2015-03-02 15:40:34 [MgmtSrvr] INFO     -- Node 1: Index usage is 12%(20 8K pages of total 160)
2015-03-02 15:40:38 [MgmtSrvr] INFO     -- Node 2: Data usage is 71%(23 32K pages of total 32)
2015-03-02 15:40:38 [MgmtSrvr] INFO     -- Node 2: Index usage is 12%(20 8K pages of total 160)
```

Node Logs

Node logs are specific to each node and are held in the DataDir:

```
[root@node1 ~]# locate _out.log
/cluster/MySQL_Cluster/49/ndb_49_out.log
```

```
[root@node2 ~]# locate *_out.log
/cluster/MySQL_Cluster/52/ndb_52_out.log
```

```
[root@node3 ~]# locate _out.log
/cluster/MySQL_Cluster/1/ndb_1_out.log
```

```
[root@node4 ~]# locate _out.log
/cluster/MySQL_Cluster/2/ndb_2_out.log
```

For example:

```
[root@node1 ~]# tail /cluster/MySQL_Cluster/49/ndb_49_out.log
==INITIAL==
==CONFIRMED==
==INITIAL==
==CONFIRMED==
Node 52 failed
stop checker 0
==CONFIRMED==
```

```
[root@node2 ~]# tail /cluster/MySQL_Cluster/52/ndb_52_out.log
==INITIAL==
==CONFIRMED==
==INITIAL==
==CONFIRMED==
stop checker 0
==CONFIRMED==
```

```
[root@node3 ~]# tail /cluster/MySQL_Cluster/1/ndb_1_out.log
2015-03-02 14:41:26 [ndbd] INFO     -- NDB start phase 7 completed
2015-03-02 14:41:26 [ndbd] INFO     -- Start phase 8 completed
2015-03-02 14:41:26 [ndbd] INFO     -- Phase 8 enabled foreign keys and waited forall nodes to complete start up
to this point
2015-03-02 14:41:26 [ndbd] INFO     -- Start phase 9 completed
2015-03-02 14:41:26 [ndbd] INFO     -- Phase 9 enabled APIs to start connecting
2015-03-02 14:41:26 [ndbd] INFO     -- Start phase 101 completed
2015-03-02 14:41:26 [ndbd] INFO     -- Phase 101 was used by SUMA to take over responsibility for sending some of
the asynchronous change events
2015-03-02 14:41:26 [ndbd] INFO     -- Node started
2015-03-02 14:41:26 [ndbd] INFO     -- Prepare arbitrator node 49 [ticket=0c5c000100a3b073]
2015-03-02 14:43:47 [ndbd] INFO     -- Allocate event buffering page chunk in SUMA, 16 pages, first page ref =
3594
```

```
[root@node4 ~]# tail /cluster/MySQL_Cluster/2/ndb_2_out.log
2015-03-02 14:41:26 [ndbd] INFO     -- Grant nodes to start phase: 9, nodes: 0000000000000006
2015-03-02 14:41:26 [ndbd] INFO     -- Start phase 9 completed
2015-03-02 14:41:26 [ndbd] INFO     -- Phase 9 enabled APIs to start connecting
2015-03-02 14:41:26 [ndbd] INFO     -- Grant nodes to start phase: 10, nodes: 0000000000000006
2015-03-02 14:41:26 [ndbd] INFO     -- Start phase 101 completed
2015-03-02 14:41:26 [ndbd] INFO     -- Phase 101 was used by SUMA to take over responsibility for sending some of
the asynchronous change events
2015-03-02 14:41:26 [ndbd] INFO     -- Grant nodes to start phase: 102, nodes: 0000000000000006
2015-03-02 14:41:26 [ndbd] INFO     -- Node started
2015-03-02 14:41:27 [ndbd] INFO     -- Started arbitrator node 49 [ticket=0c5c000100a3b073]
2015-03-02 14:43:47 [ndbd] INFO     -- Allocate event buffering page chunk in SUMA, 16 pages, first page ref =
3720
```

The ndbinfo MySQL Cluster Information Database

It is possible to interrogate the cluster in real time by using the ndbinfo database.

For example, we can connect to the first SQL node on the node1 host as follows:

```
[root@node1 ~]# /cluster/mysql_cluster/bin/mysql --socket="/cluster/MySQL_Cluster/55/mysql.socket"
Welcome to the MySQL monitor.  Commands end with ; or \g.
Your MySQL connection id is 2
Server version: 5.6.23-ndb-7.4.4-cluster-gpl MySQL Cluster Community Server (GPL)

Copyright (c) 2000, 2015, Oracle and/or its affiliates. All rights reserved.

Oracle is a registered trademark of Oracle Corporation and/or its
affiliates. Other names may be trademarks of their respective
owners.

Type 'help;' or '\h' for help. Type '\c' to clear the current input statement.
```

```
mysql>
```

We first check that the ndbinfo support is active:

```
mysql> SHOW PLUGINS;
+----------------------------------+----------+--------------------+---------+---------+
| Name                             | Status   | Type               | Library | License |
+----------------------------------+----------+--------------------+---------+---------+
| binlog                           | ACTIVE   | STORAGE ENGINE     | NULL    | GPL     |
| mysql_native_password            | ACTIVE   | AUTHENTICATION     | NULL    | GPL     |
| mysql_old_password               | ACTIVE   | AUTHENTICATION     | NULL    | GPL     |
| sha256_password                  | ACTIVE   | AUTHENTICATION     | NULL    | GPL     |
| MEMORY                           | ACTIVE   | STORAGE ENGINE     | NULL    | GPL     |
| MRG_MYISAM                       | ACTIVE   | STORAGE ENGINE     | NULL    | GPL     |
| CSV                              | ACTIVE   | STORAGE ENGINE     | NULL    | GPL     |
| MyISAM                           | ACTIVE   | STORAGE ENGINE     | NULL    | GPL     |
| ndbcluster                       | ACTIVE   | STORAGE ENGINE     | NULL    | GPL     |
| ndbinfo                          | ACTIVE   | STORAGE ENGINE     | NULL    | GPL     |
| ndb_transid_mysql_connection_map | ACTIVE   | INFORMATION SCHEMA | NULL    | GPL     |
| FEDERATED                        | DISABLED | STORAGE ENGINE     | NULL    | GPL     |
| ARCHIVE                          | ACTIVE   | STORAGE ENGINE     | NULL    | GPL     |
| InnoDB                           | ACTIVE   | STORAGE ENGINE     | NULL    | GPL     |
| INNODB_TRX                       | ACTIVE   | INFORMATION SCHEMA | NULL    | GPL     |
| INNODB_LOCKS                     | ACTIVE   | INFORMATION SCHEMA | NULL    | GPL     |
| INNODB_LOCK_WAITS                | ACTIVE   | INFORMATION SCHEMA | NULL    | GPL     |
| INNODB_CMP                       | ACTIVE   | INFORMATION SCHEMA | NULL    | GPL     |
| INNODB_CMP_RESET                 | ACTIVE   | INFORMATION SCHEMA | NULL    | GPL     |
| INNODB_CMPMEM                    | ACTIVE   | INFORMATION SCHEMA | NULL    | GPL     |
| INNODB_CMPMEM_RESET              | ACTIVE   | INFORMATION SCHEMA | NULL    | GPL     |
| INNODB_CMP_PER_INDEX             | ACTIVE   | INFORMATION SCHEMA | NULL    | GPL     |
| INNODB_CMP_PER_INDEX_RESET       | ACTIVE   | INFORMATION SCHEMA | NULL    | GPL     |
| INNODB_BUFFER_PAGE               | ACTIVE   | INFORMATION SCHEMA | NULL    | GPL     |
| INNODB_BUFFER_PAGE_LRU           | ACTIVE   | INFORMATION SCHEMA | NULL    | GPL     |
| INNODB_BUFFER_POOL_STATS         | ACTIVE   | INFORMATION SCHEMA | NULL    | GPL     |
```

```
| INNODB_METRICS                    | ACTIVE   | INFORMATION SCHEMA | NULL    | GPL     |
| INNODB_FT_DEFAULT_STOPWORD        | ACTIVE   | INFORMATION SCHEMA | NULL    | GPL     |
| INNODB_FT_DELETED                 | ACTIVE   | INFORMATION SCHEMA | NULL    | GPL     |
| INNODB_FT_BEING_DELETED           | ACTIVE   | INFORMATION SCHEMA | NULL    | GPL     |
| INNODB_FT_CONFIG                  | ACTIVE   | INFORMATION SCHEMA | NULL    | GPL     |
| INNODB_FT_INDEX_CACHE             | ACTIVE   | INFORMATION SCHEMA | NULL    | GPL     |
| INNODB_FT_INDEX_TABLE             | ACTIVE   | INFORMATION SCHEMA | NULL    | GPL     |
| INNODB_SYS_TABLES                 | ACTIVE   | INFORMATION SCHEMA | NULL    | GPL     |
| INNODB_SYS_TABLESTATS             | ACTIVE   | INFORMATION SCHEMA | NULL    | GPL     |
| INNODB_SYS_INDEXES                | ACTIVE   | INFORMATION SCHEMA | NULL    | GPL     |
| INNODB_SYS_COLUMNS                | ACTIVE   | INFORMATION SCHEMA | NULL    | GPL     |
| INNODB_SYS_FIELDS                 | ACTIVE   | INFORMATION SCHEMA | NULL    | GPL     |
| INNODB_SYS_FOREIGN                | ACTIVE   | INFORMATION SCHEMA | NULL    | GPL     |
| INNODB_SYS_FOREIGN_COLS           | ACTIVE   | INFORMATION SCHEMA | NULL    | GPL     |
| INNODB_SYS_TABLESPACES            | ACTIVE   | INFORMATION SCHEMA | NULL    | GPL     |
| INNODB_SYS_DATAFILES              | ACTIVE   | INFORMATION SCHEMA | NULL    | GPL     |
| BLACKHOLE                         | ACTIVE   | STORAGE ENGINE     | NULL    | GPL     |
| PERFORMANCE_SCHEMA                | ACTIVE   | STORAGE ENGINE     | NULL    | GPL     |
| partition                         | ACTIVE   | STORAGE ENGINE     | NULL    | GPL     |
+-----------------------------------+----------+--------------------+---------+---------+
45 rows in set (0.11 sec)

mysql>
```

Now we use the ndbinfo database:

```
mysql> SHOW DATABASES;
+--------------------+
| Database           |
+--------------------+
| information_schema |
| mysql              |
| ndbinfo            |
| performance_schema |
```

```
| test               |
| tmp                |
+--------------------+
6 rows in set (0.06 sec)

mysql> USE ndbinfo;
Reading table information for completion of table and column names
You can turn off this feature to get a quicker startup with -A

Database changed
mysql> SHOW TABLES;
+---------------------------------+
| Tables_in_ndbinfo               |
+---------------------------------+
| arbitrator_validity_detail      |
| arbitrator_validity_summary     |
| blocks                          |
| cluster_operations              |
| cluster_transactions            |
| config_params                   |
| counters                        |
| dict_obj_types                  |
| disk_write_speed_aggregate      |
| disk_write_speed_aggregate_node |
| disk_write_speed_base           |
| diskpagebuffer                  |
| logbuffers                      |
| logspaces                       |
| membership                      |
| memory_per_fragment             |
| memoryusage                     |
| nodes                           |
| operations_per_fragment         |
| resources                       |
```

```
| restart_info                    |
| server_operations               |
| server_transactions             |
| threadblocks                    |
| threadstat                      |
| transporters                    |
+---------------------------------+
26 rows in set (0.01 sec)

mysql>
```

We are now able to query this database for real time statistics such as the memory usage:

```
mysql> SELECT * FROM memoryusage;
+---------+---------------------+--------+------------+----------+-------------+
| node_id | memory_type         | used   | used_pages | total    | total_pages |
+---------+---------------------+--------+------------+----------+-------------+
|       1 | Data memory         | 753664 |         23 |  1048576 |          32 |
|       1 | Index memory        | 163840 |         20 |  1310720 |         160 |
|       1 | Long message buffer | 262144 |       1024 | 33554432 |      131072 |
|       2 | Data memory         | 753664 |         23 |  1048576 |          32 |
|       2 | Index memory        | 163840 |         20 |  1310720 |         160 |
|       2 | Long message buffer | 393216 |       1536 | 33554432 |      131072 |
+---------+---------------------+--------+------------+----------+-------------+
6 rows in set (0.02 sec)

mysql>
```

Backing Up MySQL Cluster

A backup of MySQL Cluster is a snapshot of the database. The snapshot consists of three main parts:

- **Metadata** - The names and definitions of all database tables
  - File name format: BACKUP-backup_id.node_id.ctl
- **Table records** - The data stored in the database tables
  - File name format: BACKUP-backup_id-0.node_id.data
- **Transaction log** - A sequential record of how and when data was stored in the database
  - File name format: BACKUP-backup_id.node_id.log

The backup is started by using the **ndb_mgm** command with one of the following options:

- **NOWAIT** - ndb_mgm immediately returns a prompt
- **STARTED** - ndb_mgm returns a prompt after the backup has started
- **COMPLETED** - ndb_mgm returns a prompt after the backup has completed

A snapshot option can also be stipulated:

- **SNAPSHOTSTART** - the snapshot reflects the status at the start of the backup process
- **SNAPSHOTEND** - the snapshot reflects the status at the end of the backup process

For example:

```
[root@node1 ~]# /cluster/mysql_cluster/bin/ndb_mgm
-- NDB Cluster -- Management Client --
ndb_mgm> START BACKUP SNAPSHOTEND WAIT COMPLETED
Connected to Management Server at: localhost:1186
Waiting for completed, this may take several minutes
Node 1: Backup 2 started from node 49
Node 1: Backup 2 started from node 49 completed
 StartGCP: 10336 StopGCP: 10339
 #Records: 2054 #LogRecords: 0
 Data: 50696 bytes Log: 0 bytes
ndb_mgm>

[root@node2 ~]# /cluster/mysql_cluster/bin/ndb_mgm
-- NDB Cluster -- Management Client --
ndb_mgm> START BACKUP SNAPSHOTEND WAIT COMPLETED
```

```
Connected to Management Server at: localhost:1186
Waiting for completed, this may take several minutes
Node 1: Backup 3 started from node 52
Node 1: Backup 3 started from node 52 completed
 StartGCP: 12310 StopGCP: 12313
 #Records: 2054 #LogRecords: 0
 Data: 50696 bytes Log: 0 bytes
ndb_mgm>
```

Backups are saved to the directory/BACKUP stipulated by the value of the **BackupDataDir** parameter on the data node hosts which defaults to DataDir/BACKUP:

```
[root@node3 ~]# ls /cluster/MySQL_Cluster/1/BACKUP
BACKUP-1  BACKUP-2
```

```
[root@node4 BACKUP-1]# ls /cluster/MySQL_Cluster/2/BACKUP
BACKUP-1
```

~~DISCUSSION:off~~
