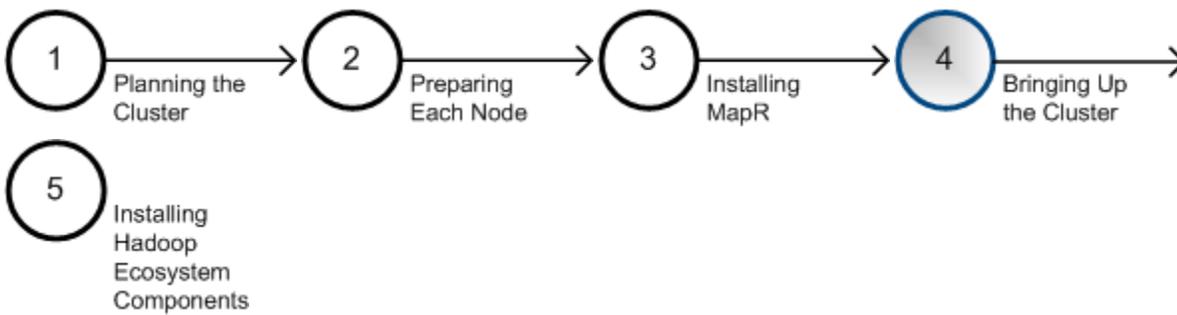


Bringing Up the Cluster



The installation of software across a cluster of nodes will go more smoothly if the services have been pre-planned and each node has been validated. Referring to the cluster design developed in [Planning the Cluster](#), ensure that each node has been prepared and meets the minimum requirements described in [Preparing Each Node](#), and that the MapR packages have been [installed on each node](#) in accordance with the plan.

- [Initialization Sequence](#)
- [Troubleshooting Initialization](#)
- [Installing the Cluster License](#)
- [Verifying Cluster Status](#)
- [Adding Volumes](#)
- [Next Step](#)

Bringing up the cluster involves setting up the administrative user, and installing a MapR license. Once these initial steps are done, the cluster is functional. You can use the MapR Control System Dashboard, or the MapR Command Line Interface, to examine nodes and activity on the cluster.

Initialization Sequence

i On versions 3.1 and later of the MapR distribution for Hadoop, the `configure.sh` script initializes the cluster automatically after a successful setup, and you can skip this process.

First, start the ZooKeeper service. It is important that all ZooKeeper instances start up, because the rest of the system cannot start unless a majority (or *quorum*) of ZooKeeper instances are up and running. Next, start the **warden** service on each node, or at least on the nodes that host the CLDB and webserver services. The warden service manages all MapR services on the node (except ZooKeeper) and helps coordinate communications. Starting the warden automatically starts the CLDB.

To bring up the cluster

1. Start **ZooKeeper** on all nodes where it is installed, by issuing the following command:

```
service mapr-zookeeper start
```
2. Verify that the quorum has been successfully established. Issue the following command and make sure that one Zookeeper is the Leader and the rest are Followers before starting the **warden**:

```
service mapr-zookeeper qstatus
```

i When the service `mapr-zookeeper qstatus` finds that a running ZooKeeper service is neither a leader nor a follower, the service issues an alarm that you can view with the `maprccli alarm list` command.

3. Start the **warden** on all nodes where CLDB is installed by issuing the following command:

```
service mapr-warden start
```

i Before continuing, wait 30 to 60 seconds for the warden to start the CLDB service. Calls to `maprccli` commands may fail if executed before the CLDB has started successfully.

4. Verify that a CLDB master is running by issuing the `maprccli node cldbmaster` command. For example:

```
# maprccli node cldbmaster
cldbmaster
ServerID: 4553404820491236337 HostName: node-36.boston
```

Do not proceed until a CLDB master is active.

5. Start the **warden** on all remaining nodes using the following command:

```
service mapr-warden start
```

6. Issue the following command to give full permission to the chosen administrative user:
`/opt/mapr/bin/maprcli acl edit -type cluster -user <user>:fc`

Troubleshooting Initialization

Difficulty bringing up the cluster seems daunting, but most cluster problems are easily resolved. For the latest support tips, visit <http://answers.mapr.com>.

- Can each node connect with the others? For a list of ports that must be open, see [Ports Used by MapR](#).
- Is the warden running on each node? On the node, run the following command as root:

```
$ service mapr-warden status
WARDEN running as process 18732
```

If the warden service is not running, check the warden log file, `/opt/mapr/logs/warden.log`, for clues
To restart the warden service:

```
$ service mapr-warden start
```

- The ZooKeeper service is not running on one or more nodes
 - Check the warden log file for errors related to resources, such as low memory
 - Check the warden log file for errors related to user permissions
 - Check for DNS and other connectivity issues between ZooKeeper nodes
- The MapR CLI program `/opt/mapr/bin/maprcli` won't run
 - Did you configure this node? See [Installing MapR Software](#).
- Permission errors appear in the log
 - Check that MapR changes to the following files have not been overwritten by automated configuration management tools:

<code>/etc/sudoers</code>	Allows the <code>mapr</code> user to invoke commands as root
<code>/etc/security/limits.conf</code>	Allows MapR services to increase limits on resources such as memory, file handles, threads and processes, and maximum priority level
<code>/etc/udev/rules.d/99-mapr-disk.rules</code>	Covers permissions and ownership of raw disk devices

Before contacting MapR Support, you can collect your cluster's logs using the [mapr-support-collect](#) script.

Installing the Cluster License

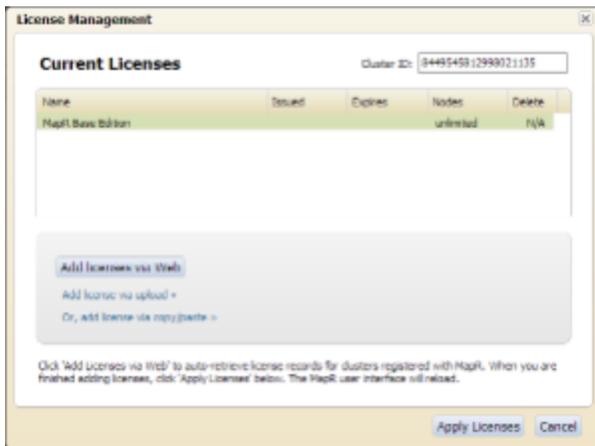
MapR Hadoop requires a valid license file, even for the free M3 Community Edition.



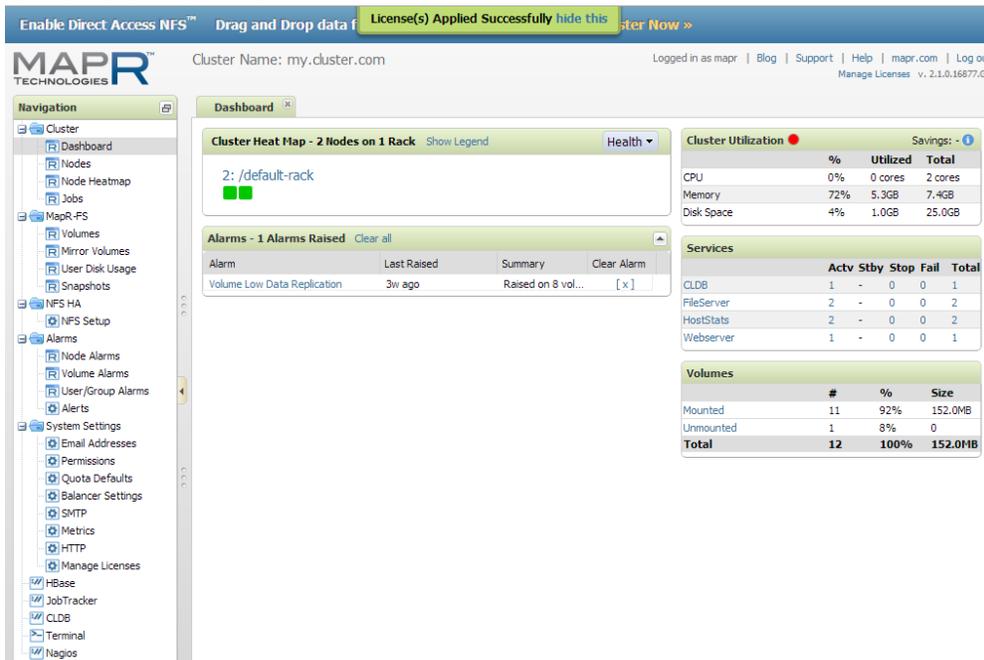
Click on the thumbnail images to view them full-size.

Using the web-based MCS to install the license

1. On a machine that is connected to the cluster and to the Internet, perform the following steps to open the MapR Control System and install the license:
 - a. In a browser, view the MapR Control System by navigating to the node that is running the MapR Control System: `https://<web server>/:8443`
Your computer won't have an HTTPS certificate yet, so the browser will warn you that the connection is not trustworthy. You can ignore the warning this time.
 - b. The first time MapR starts, you must accept the Terms of Use and choose whether to enable the MapR [Dial Home](#) service.
 - c. Log in to the MapR Control System as the administrative user you designated earlier.
Until a license is applied, the MapR Control System dashboard might show some nodes in the amber "degraded" state. Don't worry if not all nodes are green and "healthy" at this stage.
2. In the navigation pane of the MapR Control System, expand the [System Settings Views](#) group and click [Manage Licenses](#) to display the MapR License Management dialog.
3. Click **Add Licenses via Web**.



- If the cluster is already registered, the license is applied automatically. Otherwise, click **OK** to register the cluster on MapR.com and follow the instructions there.



Installing a license from the command line

Use the following steps if it is not possible to connect to the cluster and the Internet at the same time.

- Obtain a valid license file from MapR
- Copy the license file to a cluster node
- Run the following command to add the license:

```
maprcli license add [ -cluster <name> ] -license <filename> -is_file true
```

Verifying Cluster Status

To view cluster status using the web interface

- Log in to the MapR Control System.
- Under the **Cluster** group in the left pane, click **Dashboard**.
- Check the **Services** pane and make sure each service is running the correct number of instances, according to your cluster plan.

To view cluster status using the command line interface

- Log in to a cluster node

2. Use the following command to list MapR services:

```
$ maprcli service list
name      state  logpath                               displayname
fileserver 0      /opt/mapr/logs/mfs.log                 FileServer
webserver  0      /opt/mapr/logs/adminuiapp.log          WebServer
cldb      0      /opt/mapr/logs/cldb.log                 CLDB
hoststats 0      /opt/mapr/logs/hoststats.log           HostStats

$ maprcli license list
$ maprcli disk list -host <name or IP address>
```

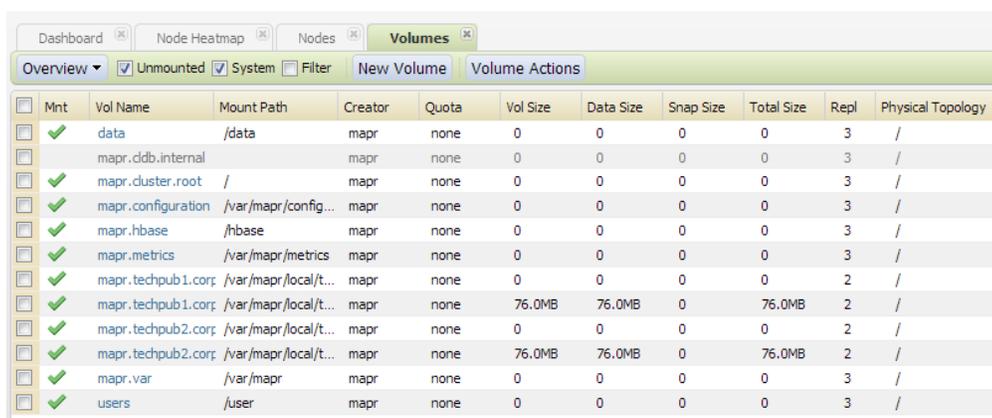
Next, start the warden on all remaining nodes using one of the following commands:

```
service mapr-warden start
/etc/init.d/mapr-warden start
```

Adding Volumes

Referring to the volume plan created in [Planning the Cluster](#), use the MapR Control System or the `maprcli` command to create and mount distinct volumes to allow more granularity in specifying policy for subsets of data.

 If you do not set up volumes, and instead store all data in the single volume mounted at `/`, it creates problems in administering data policy later as data size grows.



Mnt	Vol Name	Mount Path	Creator	Quota	Vol Size	Data Size	Snap Size	Total Size	Repl	Physical Topology
<input checked="" type="checkbox"/>	data	/data	mapr	none	0	0	0	0	3	/
<input type="checkbox"/>	mapr.cldb.internal		mapr	none	0	0	0	0	3	/
<input checked="" type="checkbox"/>	mapr.cluster.root	/	mapr	none	0	0	0	0	3	/
<input checked="" type="checkbox"/>	mapr.configuration	/var/mapr/config...	mapr	none	0	0	0	0	3	/
<input checked="" type="checkbox"/>	mapr.hbase	/hbase	mapr	none	0	0	0	0	3	/
<input checked="" type="checkbox"/>	mapr.metrics	/var/mapr/metrics	mapr	none	0	0	0	0	3	/
<input checked="" type="checkbox"/>	mapr.techpub1.corp	/var/mapr/local/t...	mapr	none	0	0	0	0	2	/
<input checked="" type="checkbox"/>	mapr.techpub1.corp	/var/mapr/local/t...	mapr	none	76.0MB	76.0MB	0	76.0MB	2	/
<input checked="" type="checkbox"/>	mapr.techpub2.corp	/var/mapr/local/t...	mapr	none	0	0	0	0	2	/
<input checked="" type="checkbox"/>	mapr.techpub2.corp	/var/mapr/local/t...	mapr	none	76.0MB	76.0MB	0	76.0MB	2	/
<input checked="" type="checkbox"/>	mapr.var	/var/mapr	mapr	none	0	0	0	0	3	/
<input checked="" type="checkbox"/>	users	/user	mapr	none	0	0	0	0	3	/

Next Step

Now that the MapR Hadoop cluster is up and running, the final installation step is to [install Hadoop Ecosystem Components](#). If you will not install any Hadoop components, see [Next Steps After Installation](#) for a list of post-install considerations.

