



# Warehouse (MapR) Configuration Guide

for Version 11.2



Copyright © 1994-2018 Dell Inc. or its subsidiaries. All Rights Reserved.

## **Contact Information**

RSA Link at <https://community.rsa.com> contains a knowledgebase that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

## **Trademarks**

For a list of RSA trademarks, go to [www.emc.com/legal/emc-corporation-trademarks.htm#rsa](http://www.emc.com/legal/emc-corporation-trademarks.htm#rsa).

## **License Agreement**

This software and the associated documentation are proprietary and confidential to Dell, are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by Dell.

## **Third-Party Licenses**

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on RSA Link. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

## **Note on Encryption Technologies**

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

## **Distribution**

Dell believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

September 2018

# Contents

---

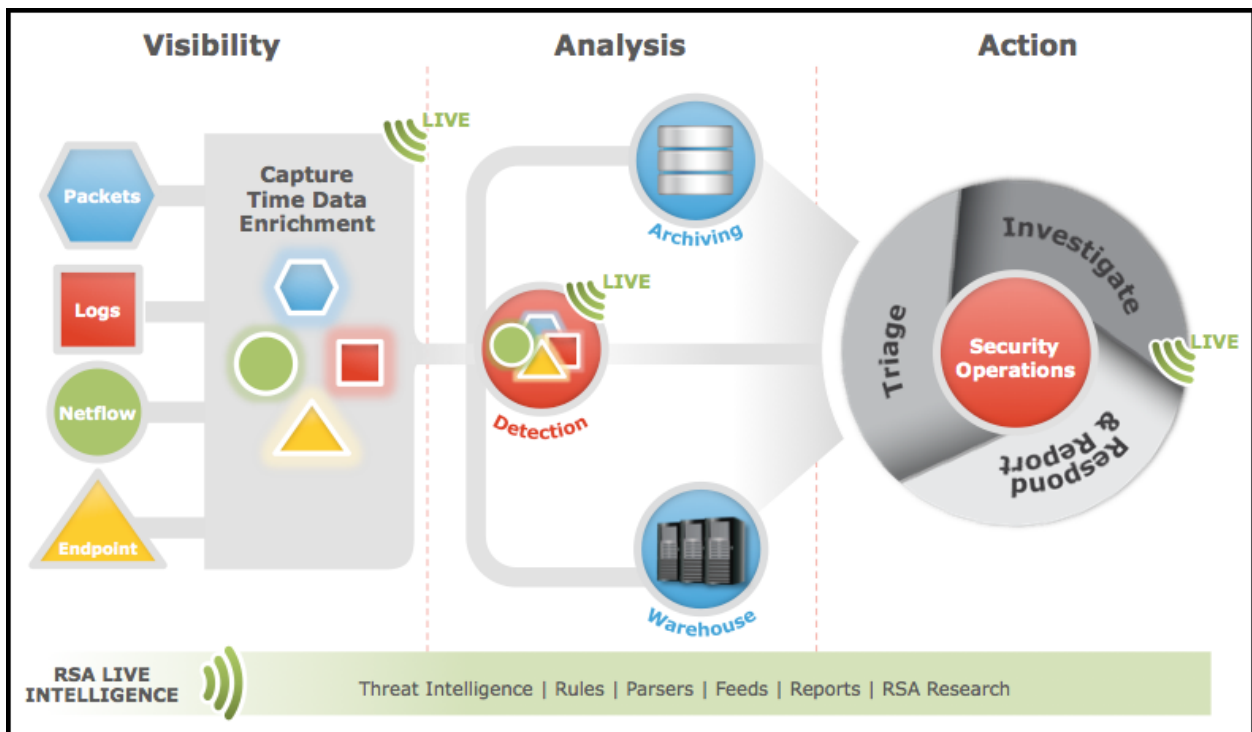
<b>RSA NetWitness Warehouse Overview</b> .....	<b>4</b>
<b>How Warehouse (MapR) Works</b> .....	<b>5</b>
<b>Configure MapR</b> .....	<b>7</b>
Generate and Update the Default UUID in Appliances .....	7
Update the Configuration Template File .....	7
Upgrade the Warehouse Cluster .....	9
Install the Warehouse License File .....	10
Generate the Virtual IP Address for Primary Appliance .....	10
Configure other NetWitness Platform Services .....	11
Stop the Hbase Services Using the Command Line .....	12
Stop the Hbase Services Using the MapR Control System .....	13
<b>Configure Warehouse Connector to Write to NetWitness Warehouse</b> .....	<b>17</b>
Verify the Network File System (NFS) Services Status .....	17
Install the Network File System Packages .....	17
Mount the Warehouse on the Warehouse Connector .....	19
<b>Manage MapR Cluster</b> .....	<b>21</b>
Access MapR Control System UI for Cluster Administration .....	21
Enable MapR Metrics on RSA NetWitnessWarehouse Cluster .....	22
Edit and Remove Virtual IP Addresses (Command Line) .....	23
Add and Remove a Virtual IP Address (MapR UI) .....	24
Add a Virtual IP Address with Multiple Nodes (MapR UI) .....	28
Optimal VIP Configuration .....	28
Optimal Configuration with the Warehouse Connector .....	29

## RSA NetWitness Warehouse Overview

RSA NetWitness Warehouse provides the capacity to process large amounts of current and long term data through a Hadoop-based distributed computing system that collects, manages, and enables advanced analytics and reporting on NetWitness Platform data. RSA NetWitness Warehouse requires a service called Warehouse Connector to collect metadata and events from Decoder and Log Decoder and write them in Avro format into a Hadoop-based distributed computing system. For more information on the Warehouse Connector, see "How Warehouse Connector Works" topic in the *Warehouse Connector Configuration Guide*.

The Warehouse is made up of three or more nodes depending on the organization's analytic, archiving, and resiliency requirements.

The following figure depicts the architecture of a NetWitness Platform network that implements the RSA NetWitness Warehouse component.



## How Warehouse (MapR) Works

---

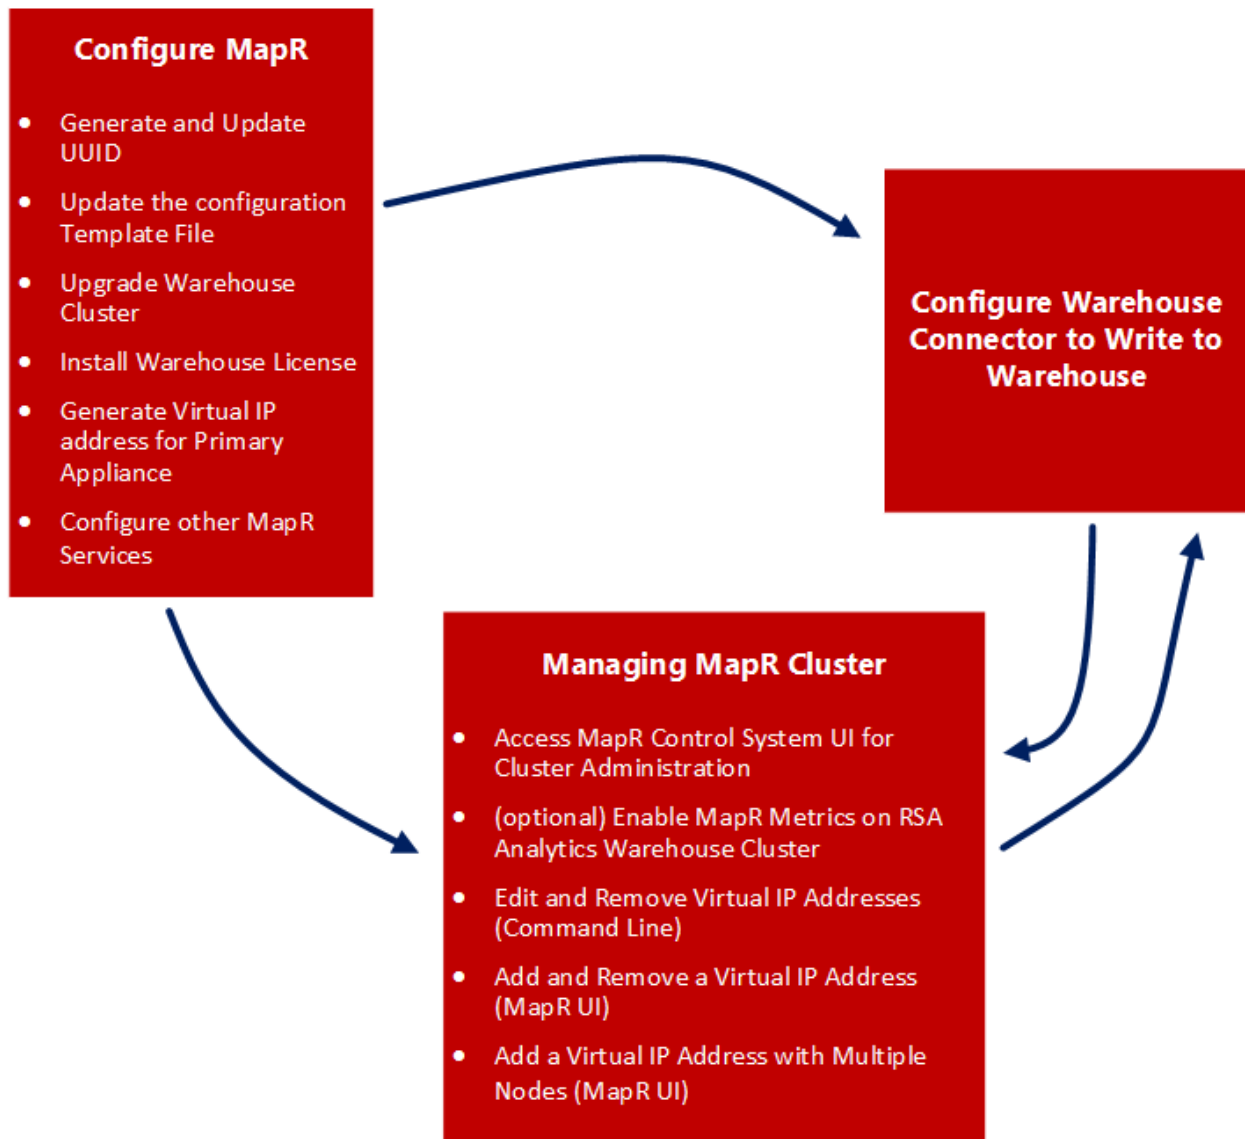
You must configure the nodes for the RSA NetWitness Warehouse (MapR). It only applies to RSA NetWitness Warehouse instances running MapR.

### Prerequisites

Make sure that you have:

- Installed the RSA NetWitness Warehouse appliance in your network environment. For more information, see "RSA Analytics Warehouse (MapR) Setup Guide" in the *Hardware Setup guides*.
- Configured the network interface of the Warehouse appliance.

This figure is an overview of the entire process of configuring Warehouse appliance in your network.



To configure the nodes for the RSA NetWitness Warehouse (MapR), perform the following:

1. [Configure MapR](#)
2. [Configure Warehouse Connector to Write to NetWitness Warehouse](#)
3. [Manage MapR Cluster](#)

**Note:** If you are planning to have a cluster of Warehouse appliances, make sure you perform the following tasks on all the appliances in the cluster.

**Caution:** Prerequisites are mandatory. Your installation will fail if you have not set the network configuration as described in the *RSA Analytics Warehouse (MapR) Setup Guide* or *Virtual Host Setup Guide* depending on your deployment.

## Configure MapR

---

You can configure MapR using the following procedure.

### Generate and Update the Default UUID in Appliances

You need to manually generate and update the default Universally Unique Identifier (UUID) on the Appliances in the cluster. The UUID must be unique to the Appliance in the cluster.

#### To generate and update the default UUID in the Appliance:

1. Log on to the Appliance as root user.
2. Generate the UUID and copy it in the correct files, using the following commands:

- `/opt/mapr/server/mruuidgen > /opt/mapr/hostid`
- `cp /opt/mapr/hostid /opt/mapr/server/hostid.xxxxx`

Where, xxxxx refers to the 5-digit number randomly assigned to the existing file.

**Note:** Review `/opt/mapr/server` for the full name of this file.

3. Restart the appliance, using the following command:

```
reboot
```

### Update the Configuration Template File

You must update the configuration template file in the RSA NetWitness Warehouse Appliance. The configuration template file in the RSA NetWitness Warehouse appliance must include the following parameters:

- nodes
- Internalnetworks
- clustername
- disks

By default, a configuration template is provided with the RSA NetWitness Warehouse appliance and is located on the RSA NetWitness Warehouse appliance at `/opt/rsa/saw/install`.

### Prerequisites

Make sure that you validated the volume in the server to identify available drive space for Warehouse to store data. The total drive space of the additional volume is considered as a single drive by the HDFS. In Warehouse, the AVRO files are stored in the drive space.

**Note:** The server contains additional volumes of identical size other than the operation system volume.

To check free space, enter the command `fdisk -l | grep /dev/s |sort` in the Warehouse node. You will get a list of disks that are not partitioned for usage. You need to list the identified disks in the configuration template file so that Warehouse utilizes this space for the Hadoop Cluster.

### To update the configuration template file in the RSA NetWitness Warehouse Appliance:

1. Log on to the appliance as the root user.
2. Navigate to `/opt/rsa/saw/install`, enter the following command:  

```
cd /opt/rsa/saw/install
```
3. Create a copy of the configuration template, enter the following command:

```
cp conf.template conf.template-<name>
```

where `<name>` is the custom name of the configuration template file.

4. Edit the configuration template file, enter the following command:

```
vi conf.template-<name>
```

Parameter	Description
Nodes	List the IP addresses of the appliances in the cluster separated by spaces. All the appliances in the cluster must be listed in the same order in every configuration file for each RSA NetWitness Warehouse appliance.
Internalnetworks	List the network addresses in CIDR format separated by spaces. This Warehouse appliance cluster communication is limited to the provided network addresses.  <b>Note:</b> RSA recommends that you do not leave this parameter blank.
Clustername	Name of the cluster. The cluster name is used to identify the Network File System (NFS) share.
Disks	Displays the list of disks recognized by the operation system, and these disks will be formatted in HDFS for the Warehouse when this configuration script is executed.

The following figure displays a sample configuration template file:



```
[root@saw-node2 install]# vi conf.template-test
[global]

# nodes: List of the first 5 node IP addresses in the cluster, separated by
#       spaces. Use addresses on internal network if restricting network traffic
nodes=xxx.108.x.25 xxx.108.x.27 xxx.108.x.33

# internalnetworks: List of network addresses, in CIDR format separated by
#                  spaces, that cluster communication will be limited to.
#                  Leave blank to allow communication over any network
internalnetworks=xxx.108.0/24

# clustername: Name of cluster. NFS share will be /mapr/<clustername>
clustername=saw

# Internal settings - changing these may result in unsupported behavior
[internal]

disks=/dev/sdb /dev/sdc /dev/sdd /dev/sde /dev/sdf /dev/sdg /dev/sdh /dev/sdi /dev/sdj
```

- Execute the configuration template file, using the following command:

```
./configure.py conf.template-<name>
```

- Restart the appliance, using the following command:

```
reboot
```

## Upgrade the Warehouse Cluster

You must upgrade the warehouse cluster after updating the configuration template file and reboot the RSA NetWitness Warehouse appliance.

### To Upgrade the Warehouse Cluster

You must manually open Hiveserver port 10000, which is not opened by default:

- Get the line number where the REJECT statement appears in the Iptable.
- Make sure that the Iptables service is running, using the following command:

```
NUM=$(iptables -L INPUT -n --line-numbers |grep 'reject-with' |awk
' {print $1}')
```

**Note:** The ACCEPT statements that follow the REJECT statement in the Iptables will not take effect. You can incorporate the line number of the REJECT statement in the command to ensure that the ACCEPT statements proceed the REJECT statement.

- Add the firewall exception for port 10000 to the Iptables. Enter the following command:

```
iptables -I INPUT $NUM -m state --state NEW -p tcp --dport 10000 -j
ACCEPT
```

- Save the Iptables. Enter the following command:

```
/etc/init.d/iptables save
```

- Restart the Iptables. Enter the following command:

```
/etc/init.d/iptables restart
```

- Verify if the firewall exceptions for the ports are added. Enter the following command:

```
Service iptables status | grep 10000
```

The following output should be displayed:

```
ACCEPT tcp -- 0.0.0.0/0 0.0.0.0/0 state NEW tcp dpt:10000
```

## Install the Warehouse License File

You need to manually install the Warehouse license file on the Warehouse appliance. If you have a cluster of Warehouse appliances, you need to install the license file on the first Warehouse appliance in the cluster.

### Prerequisites

Make sure that you have:

- Obtained the Warehouse license file.
- Copied the license file to `/root/` on the first Warehouse appliance in the cluster using a USB drive or through SCP.

### To install the Warehouse license file:

- Log on to the appliance as a root user.
- Install the license file, using the following command:

```
maprcli license add -is_file true -license <license_filename>
```

where `<license_filename>` is filename of the RSA NetWitness Warehouse license file.

The license file is installed without any output messages. If you included a network range in the `internalnetworks` parameter in the configuration template file, a warning message appears suggesting that the Warehouse is configured only to communicate with the network entered in the configuration template file. You can ignore this warning as this does not have any functional issue.

- Confirm the license file installation, using the following command:

```
maprcli license list
```

The output messages appears on the console screen. The last two lines of the output message should be similar to the following sample:

```
hash: "b8x01f1W8EMNSqq7zztn8D2BXnQ="
  3 May 14, 2013
```

- Retrieve a list of directories, run the following command:

```
hadoop fs -ls /
```

## Generate the Virtual IP Address for Primary Appliance

Generate a virtual IP address for the primary RSA NetWitness Warehouse (Warehouse) appliance.

## Prerequisites

Make sure you note down the MAC addresses of all the Warehouse appliances in the cluster. Use the following command on the appliance to view the MAC address of appliance:

```
ifconfig <interface> | grep HWaddr
```

where <interface> is the network interface.

### To generate a virtual IP address for the primary Warehouse appliance:

1. Log on to the primary appliance as root user.
2. Create the virtual IP address. Enter the following command:  

```
maprcli virtualip add -virtualip <VIP_address> -netmask <netmask> -  
macs <mac_node1> <mac_node2> <mac_node3> .....< mac_node n>
```

where:
  - <VIP\_address> is the virtual IP address for the primary Warehouse appliance.
  - <netmask> is the netmask address of the primary Warehouse appliance.
  - <mac\_node1> is the MAC address of the first node in the Warehouse cluster.
  - <mac\_node2> is the MAC address of the second node in the Warehouse cluster.For example, if the MAC address for node 1 is 01:Z1:1X:00:20:Y1 and node 2 is 32:Y2:4Z:40:10:X3, and the IP address is 192.168.100.10, then enter the command as following:  

```
maprcli virtualip add -virtualip 192.168.100.10 -netmask <netmask>  
-macs 01:Z1:1X:00:20:Y1 32:Y2:4Z:40:10:X3
```
3. Verify the virtual IP address, using the following command:  

```
maprcli virtualip list
```
4. To add or remove virtual IP addresses, you can use the command line or the MapR Control System. For more information, see "Edit and Remove Virtual IP Addresses (Command Line)" and "Add and Remove a Virtual IP Address (MapR UI)" sections in [Manage MapR Cluster](#).

## Configure other NetWitness Platform Services

Configure other NetWitness Platform services for the RSA NetWitness Warehouse (MapR).

1. If you are not using Vulnerability Response Management (VRM), disable the Hbase services to return the configured memory so that it is available for use elsewhere in the cluster. To stop the Hbase services, you can use the command line or the MapR Control System. For more information, see [Stop the Hbase Services Using the Command Line](#) and [Stop the Hbase Services Using the MapR Control System](#).
2. Add Warehouse data sources to the Reporting Engine. For the detailed procedure, see "Add Warehouse as Data Source to Reporting Engine" topic in the *Reporting Engine Configuration Guide*.

## Stop the Hbase Services Using the Command Line

This section provides the steps to stop the Hbase services using the command line. If you are not using Vulnerability Response Management (VRM), stop the Hbase services to return the configured memory so that it is available for use elsewhere in the cluster.

### To stop the Hbase services using the command line:

1. Stop the **Hbase RegionServer** service on *all of the appliances*, using the following command:  

```
maprcli node services -hbregionserver stop -filter "[hn==*]"
```
2. Stop the **Hbase RegionServer** service on *a specific node*, using the following command:  

```
maprcli node services -hbregionserver stop -filter "[hn==<Hostname>]"
```

Where <Hostname> is the specific node hostname.
3. Stop the **Hbase Master** service on *all of the appliances*, using the following command:  

```
maprcli node services -hbmaster stop -filter "[hn==*]"
```
4. Stop the **Hbase Master** service on *a specific node*, using the following command:  

```
maprcli node services -hbmaster stop -filter "[hn==<Hostname>]"
```

Where <Hostname> is the specific node hostname.

### Hbase Services Stop and Start Commands Summary

The following tables summarize the commands used to stop and start the Hbase services for the **HBase RegionServer** and **HBase Master** services.

HBase RegionServer	Command to run using the Command Line
Stop on All the Appliances	<pre>maprcli node services -hbregionserver stop -filter "[hn==*]"</pre>
Start on All the Appliances	<pre>maprcli node services -hbregionserver start -filter "[hn==*]"</pre>
Stop on Specific node	<pre>maprcli node services -hbregionserver stop -filter "[hn==&lt;Hostname&gt;]"</pre>
Start on Specific node	<pre>maprcli node services -hbregionserver start -filter "[hn==&lt;Hostname&gt;]"</pre>

HBase Master	Command to run using the Command Line
Stop on All the Appliances	<pre>maprcli node services -hbmaster stop -filter "[hn==*]"</pre>
Start on All the Appliances	<pre>maprcli node services -hbmaster start -filter "[hn==*]"</pre>


HBase Master	Command to run using the Command Line
Stop on Specific node	<code>maprcli node services -hbmaster stop -filter "[hn==&lt;Hostname&gt;]"</code>
Start on Specific node	<code>maprcli node services -hbmaster start -filter "[hn==&lt;Hostname&gt;]"</code>

Where <Hostname> is the specific node hostname.

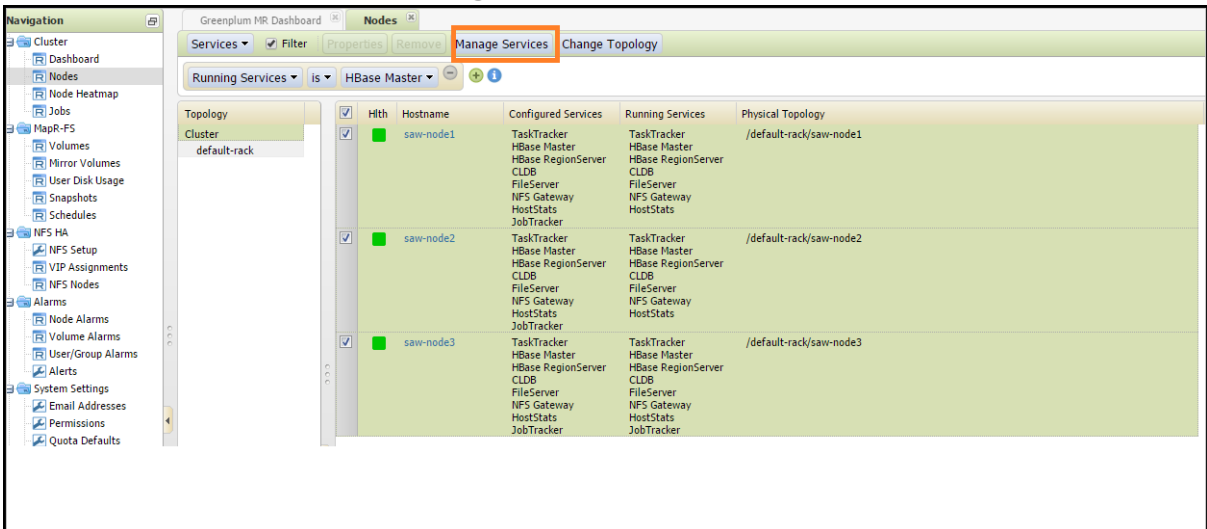
## Stop the Hbase Services Using the MapR Control System

This section provides the steps to stop the Hbase services using the MapR Control System. If you are not using Vulnerability Response Management (VRM), stop the Hbase services to return the configured memory so that it is available for use elsewhere in the cluster.

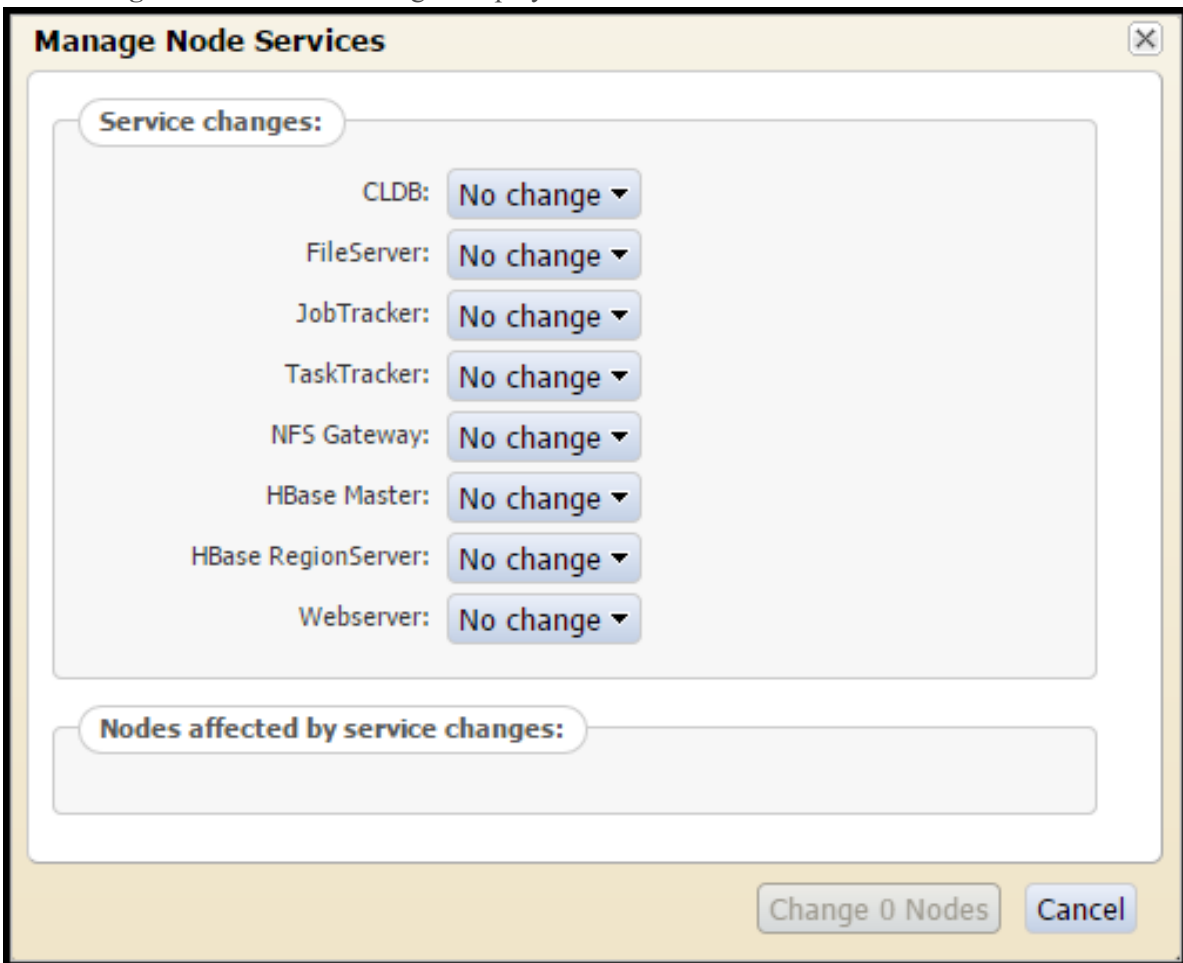
1. Log on to the MapR Control System user interface. For more information see "Access MapR Control System UI for Cluster Administration" section in [Manage MapR Cluster](#).
2. Stop the **HBase Master** services, in the **Services** section of the dashboard, click the number in the **Actv** column for the **HBase Master** service. This is the number of active services for the **HBase Master** service.



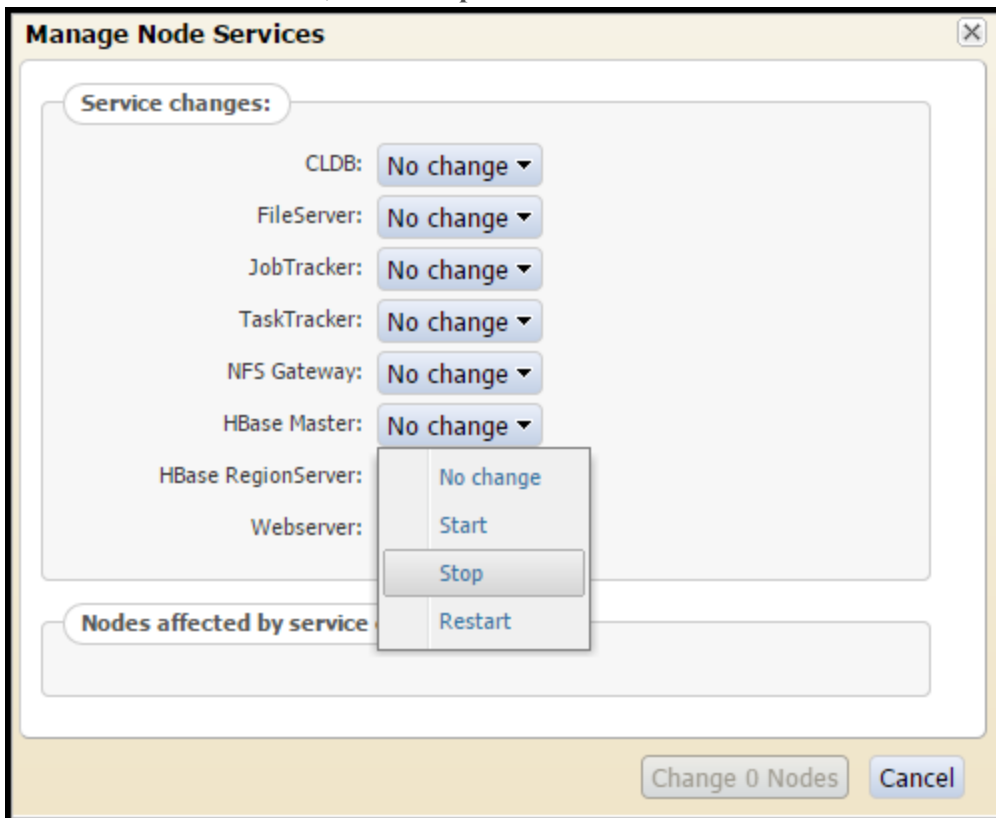
3. On the **Cluster Nodes** tab, click **Manage Services**.



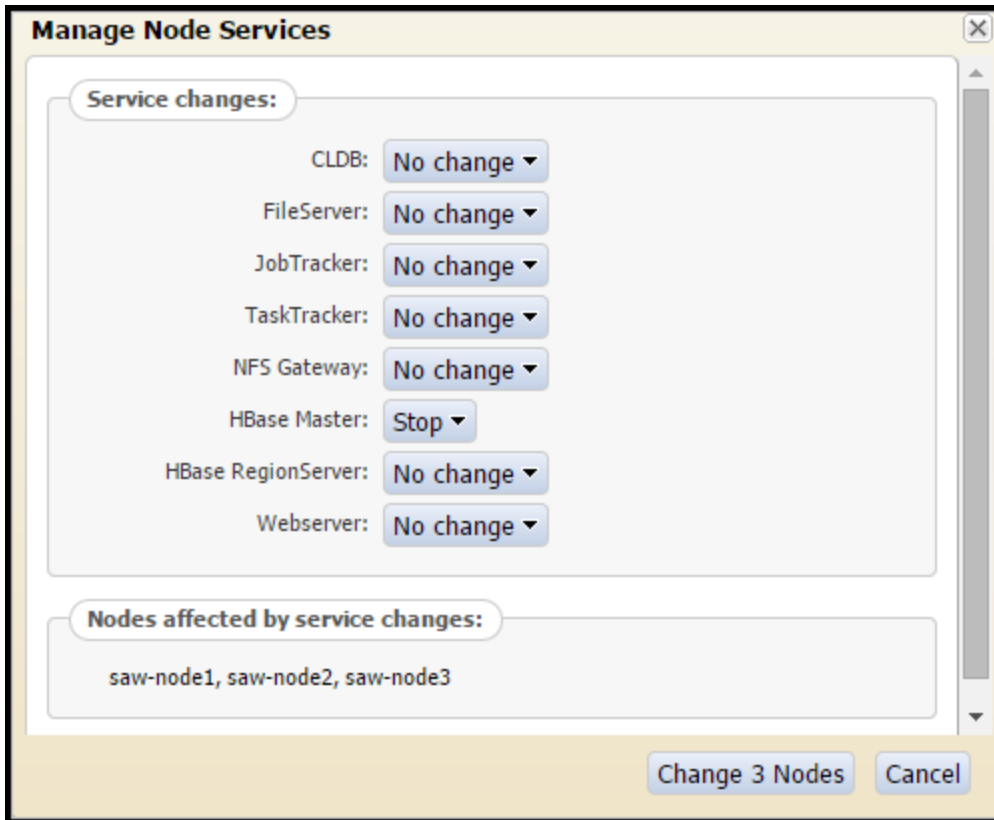
The **Manage Node Services** dialog is displayed.



4. In the **HBase Master** field, select **Stop**.



5. Click **Change <number\_of\_nodes> Nodes**.  
Where <number\_of\_nodes> is the number of active nodes selected.  
For example, click **Change 3 Nodes**.



The **Hbase Master** service on the selected nodes must be in a stopped state.

6. Stop the **Hbase RegionServer** services, repeat steps 2 to 5 for the **Hbase RegionServer** services.



## Configure Warehouse Connector to Write to NetWitness Warehouse

---

You must enable the Warehouse Connector services to write to RSA NetWitness Warehouse.

To configure Warehouse Connector to write to the NetWitness Warehouse, perform the following tasks on the Log Decoders and Decoders where the Warehouse Connectors are installed:

**Note:** If you are configuring on a virtual environment, perform these tasks on a standalone Warehouse Connector server.

### Verify the Network File System (NFS) Services Status

**To verify the NFS services status:**

1. Log on to the Warehouse Connector appliance where you have installed the Warehouse Connector service.
2. Enter the following command:  

```
rpm -qa |grep nfs
```

The NFS package names appear in the output message. For example:  

```
nfs-utils-lib-1.1.5-6.el6.x86_64  
nfs-utils-1.2.3-36.el6.x86_64
```
3. If the output message is empty, install the NFS packages.

### Install the Network File System Packages

#### Prerequisites

If the NFS packages are already downloaded on the appliances manually, install the packages and mount RSA NetWitness Warehouse. You need to have internet access to complete this task. If internet access is not available, you must download the RPM packages offline and copy them to this machine for installation.

**Note:** Install the NFS packages only if the NFS packages are not displayed when you verify the status of NFS in the Warehouse Connector appliance or on the appliance where you have installed the Warehouse Connector service.

**To install NFS packages:**

1. Log on to the Warehouse Connector appliance or on the appliance where you have installed the Warehouse Connector service.
2. Verify the NFS status, using the following command:

```
rpm -qa |grep nfs
```

The NFS package names appear in the output message. For example:

```
nfs-utils-lib-1.1.5-6.el6.x86_64  
nfs-utils-1.2.3-36.el6.x86_64
```

If the `nfs-utils` and `nfs-utils-lib` are properly identified, you can skip the remaining steps in this procedure (*Install the Network File System Packages*).

3. Search for NFS package, using the following command:

```
yum search nfs-utils
```

The output ends with the following message:

```
"name and summary matches only, use "search all" for everything."
```

**Note:** Contact RSA Customer Support if the output ends with the following message:  
"no matches found"

4. Install the NFS programs, using the following command:

```
yum install nfs-utils nfs-utils-lib
```

The output prompts for **y** or **n**. Type **y** and press **ENTER**.

The NFS packages are successfully installed.

## Mount the Warehouse on the Warehouse Connector

### To mount RSA NetWitness Warehouse on the appliance:

1. Create a new directory named `/saw`, using the following command:

```
mkdir /saw
```

2. Enter the following command:

```
ll /
```

The new directory is displayed.

3. Mount the Warehouse, using the following command:

```
mount -t nfs -o nolock,tcp,hard,intr <IP_Address_for_
SAW>:/mapr/<cluster-name> /saw
```

Where `<IP_Address_for_SAW>` is the IP address of the primary Warehouse appliance in the cluster and `<cluster-name>` is the name provided in the template file.

**Note:** If a virtual IP address is configured for the Warehouse, you have to use it as the IP address in `<IP_Address_for_SAW>`.

4. Verify if the Warehouse is mounted successfully, using the following command

```
mount
```

The IP address of the primary Warehouse appliance and other details you have provided in **step 3** appear in the last line of the output message.

5. List the content in the newly created directory, `/saw`, using the following command:

```
ll /saw
```

The following directories are displayed:

```
hbase
```

```
index-scratch
```

```
jars
```

```
logs
```

```
user
```

```
var
```

6. To add NFS to the Auto-mount options. Do the following:

- a. To check if the IP address of the primary Warehouse appliance and other details you have provided while mounting Warehouse appears in `/etc/fstab`, enter the following command:

```
cat /etc/fstab
```

If the detail does not appear in the `/etc/fstab` file, perform the following steps.

- b. Enter the following command:

```
tail -n 1 /etc/mtab
```

The IP address of the primary Warehouse appliance and other details you provided while mounting Warehouse appear in the last line of the output message.

- c. Enter the following command:

```
tail -n 1 /etc/mtab >> /etc/fstab
```

- d. Edit the `/etc/fstab` file to add the word 'auto' at the end of the file. Enter the following command:

```
vi /etc/fstab
```

For example, `10.11.111.11:/mapr/saw /saw nfs`

```
rw,nolock,tcp,auto,addr=10.11.111.11 0 0
```

## Manage MapR Cluster

---

You can manage the MapR cluster using the following procedures.

### Access MapR Control System UI for Cluster Administration

You can access the MapR Control System user interface for RSA NetWitnessWarehouse cluster administration. MapR Control System user interface enables you to administer the RSA NetWitnessWarehouse cluster. The MapR Control System user interface provides details of the following:

- Nodes
- Node Heatmap
- Jobs
- MapR Tables
- Volumes
- Mirrors
- User Disk Usage
- Snapshots
- Schedules
- NFS Setup
- Virtual IP Assignments
- NFS Nodes
- Node Alarms
- Volume Alarms
- User/Group Alarms
- HBase
- JobTracker
- CLDB

#### To access the MapR Control System user interface:

1. Log on to one of the appliances in the RSA NetWitnessWarehouse cluster.
2. Start the webserver. Enter the following command:

```
/opt/mapr/adminuiapp/webserver start
```

**Note:** The default port used by the webserver is **8443**.

**Note:** If you receive the error `/opt/mapr/conf/ssl_keystore` (No such file or directory) in the `/opt/mapr/logs/adminuiapp.log` after executing the command `/opt/mapr/adminuiapp/webserver start`, enter the following commands:

```
./configure.sh -R -genkeys
service mapr-warden restart
```

- Using a web browser to access the MapR Control System, type the following url:

`https://<NODE-IP-OR-HOSTNAME>:8443`

The MapR Control System user interface is displayed.

The screenshot shows the MapR Control System dashboard for a cluster named 'SAW'. The dashboard includes a navigation menu on the left, a cluster heatmap showing 6 nodes on 1 rack, a cluster utilization table, a table of alarms, a table of MapReduce metrics, a table of services, and a table of volumes.

Cluster Utilization				Savings: 13%
	%	Utilized	Total	
CPU	2%	1 Cores	48 Cores	
Memory	64%	55.8GB	87.6GB	
Disk Space	8%	90GB	1.1TB	

MapReduce			
Running Jobs	0		
Queued Jobs	0		
Running Tasks	0		
Running Map Tasks	0		
Running Reduce Tasks	0		
Map Task Capacity	30		
Reduce Task Capacity	24		
Map Task Prefetch Capacity	18		
Backlisted Nodes	0		

Services					
	Actv	Stby	Stop	Fail	Total
Oozie	1	-	-	1	2
FileServer	6	-	-	0	6
HiveMeta	0	-	-	0	0
NFS Gateway	3	-	-	0	3
Webserver	1	-	-	0	1
GLDB	2	-	-	0	2
TaskTracker	6	-	-	0	6
JobTracker	1	1	-	0	2
HostState	6	-	-	0	6
HiveServer 2	0	-	-	0	0

Volumes			
	#	%	Total
Mounted	25	95%	28.4GB
Unmounted	1	4%	none
<b>Total</b>	<b>26</b>	<b>100%</b>	<b>28.4GB</b>

## Enable MapR Metrics on RSA NetWitnessWarehouse Cluster

You can enable MapR Metrics on the RSA NetWitnessWarehouse cluster. This optional procedure enables Administrators to see job details in the MapR Control System UI rather than going to the JobTracker for details.

### Prerequisites

Make sure that you have the following MapR Metrics dependencies installed in your environment:

- MySQL Server installed and configured.
- Libraries hosted on the EPEL Repository.
- Libraries hosted on the CentOS base repositories.

### To Enable MapR Metrics

To enable MapR Metrics on the RSA NetWitnessWarehouse cluster, follow the instructions at the following links:

- <http://doc.mapr.com/display/MapR/Setting+up+the+MapR+Metrics+Database>
- <http://doc.mapr.com/display/MapR/MapR+Metrics+and+Job+Performance>

**Note:** Make sure you install MapR Metrics on the nodes in your RSA NetWitness Warehouse Cluster where Job Tracker or Web Server is running.

## Edit and Remove Virtual IP Addresses (Command Line)

You can edit and remove virtual IP addresses in the Warehouse cluster using the command line. This procedure is optional and used when you want to change the virtual IP addresses in the Warehouse cluster.

Adding and removing Warehouse appliances to and from a virtual IP group is accomplished by executing an **edit** command. This is the same as the add command, except that ALL of the MAC addresses are replaced with ONLY the MAC addresses that you enter.

### Prerequisites

Make sure you note down the MAC addresses of all the Warehouse appliances in the cluster. Use the following command on the appliance to view the MAC address of the appliance:

```
ifconfig <interface> | grep HWaddr
```

where <interface> is the network interface.

Also note the MAC addresses of the Warehouse appliances that you want to add.

### To add or remove a virtual IP address in the primary Warehouse appliance:

1. Log on to the primary appliance as root user.
2. Edit the virtual IP address. Enter the following command:

```
maprcli virtualip edit -virtualip <VIP_address> -netmask <netmask>  
-macs <mac_node1> <mac_node2> <mac_node3> .....< mac_node n>
```

where:

- <VIP\_address> is the virtual IP address for the primary Warehouse appliance.
- <netmask> is the netmask address of the primary Warehouse appliance.
- <mac\_node1> is the MAC address of the first node in the Warehouse cluster.
- <mac\_node2> is the MAC address of the second node in the Warehouse cluster.

For example, if the IP address of the primary warehouse is 192.168.100.10 and the MAC address for node 1 is 01:Z1:1X:00:20:Y1, node 2 is 32:Y2:4Z:40:10:X3, and you want to add node 3, which is 20:Y2:4Z:20:10:X3, then enter the following:

```
maprcli virtualip edit -virtualip 192.168.100.10 -netmask <netmask>  
-macs 01:Z1:1X:00:20:Y1 32:Y2:4Z:40:10:X3 20:Y2:4Z:20:10:X3
```

3. Verify the virtual IP addresses. Enter the following command:

```
maprcli virtualip list
```

To remove the virtual IP address of the primary Warehouse appliance group entirely:  
Enter the following command:

```
maprcli virtualip remove -virtualip 192.168.100.10
```

## Add and Remove a Virtual IP Address (MapR UI)

You can add a virtual IP address in the Warehouse cluster using the MapR Control System. This procedure is optional and used when you want to add a virtual IP address (VIP) in the Warehouse cluster.

### Prerequisites

Follow the instructions in [Access MapR Control System UI for Cluster Administration](#) before completing this procedure.

1. Log on to the MapR Control System.

The screenshot displays the MapR Control System interface for a cluster named 'SAW'. The left navigation pane is expanded to 'NFS HA > NFS Setup'. The main content area shows a 'Cluster Heatmap' with 3 nodes on 1 rack, a 'Health' indicator, and a 'Rack' dropdown. Below the heatmap is an 'Alarms' section with one alarm: 'Cluster License Near Expiration Alarm' raised 13h 55.7m ago, with a summary 'One or more licenses is about to expire within -62 days'. On the right, there are several summary tables:

	%	Utilized	Total
CPU	N/A	0 Cores	24 Cores
Memory	66%	23.3GB	34.9GB
Disk Space	10%	70GB	721GB

Running Jobs	0
Queued Jobs	0
Running Tasks	0
Running Map Tasks	0
Running Reduce Tasks	0
Map Task Capacity	18
Reduce Task Capacity	12
Map Task Prefetch Capacity	9
Blacklisted Nodes	0

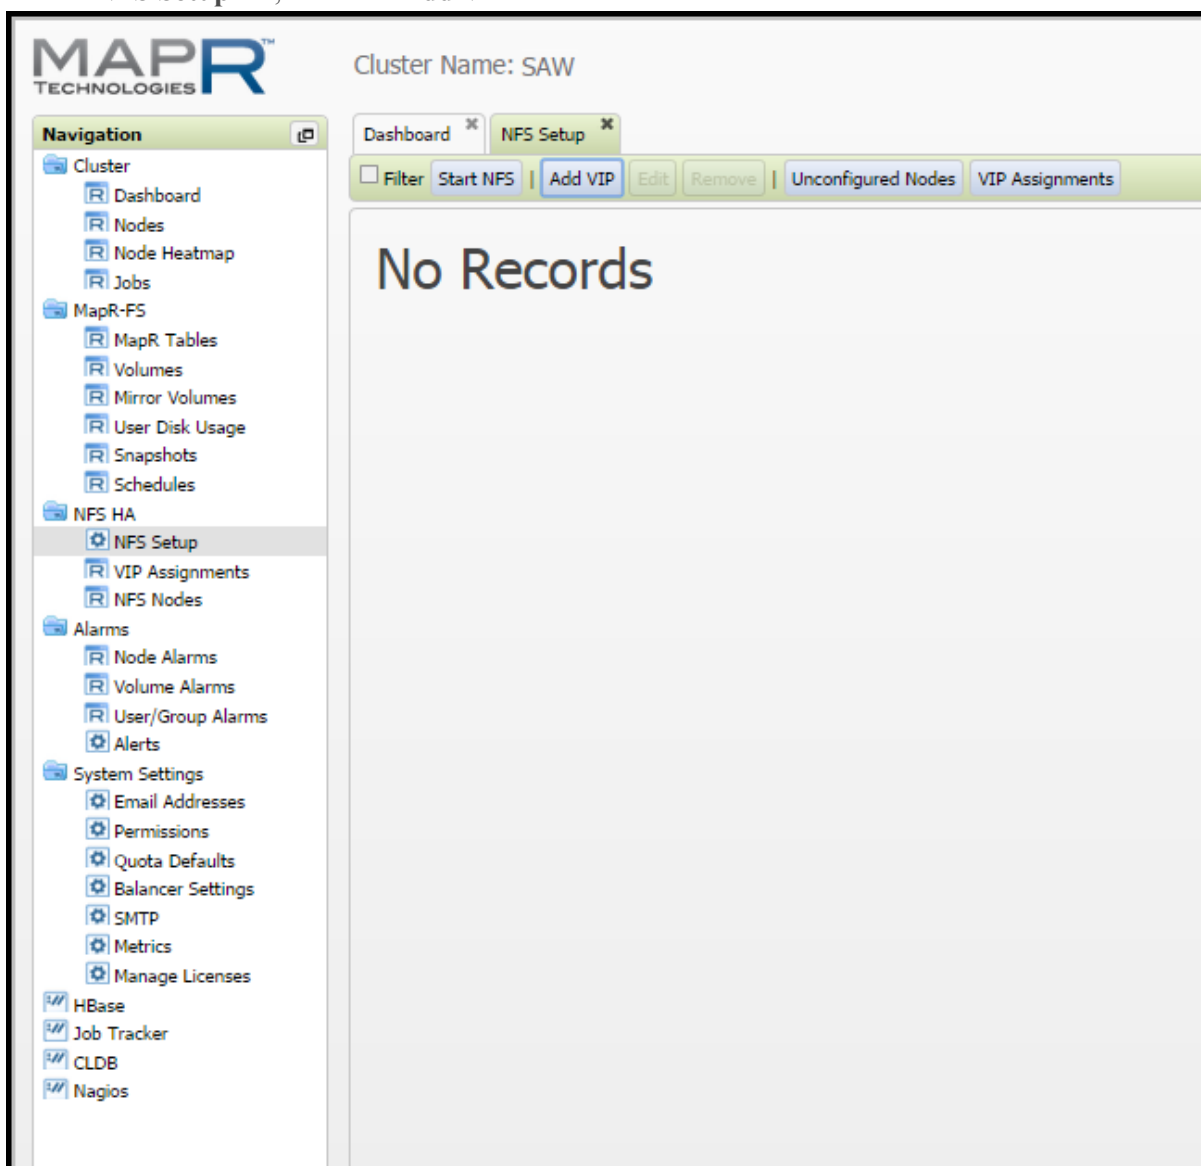
	Actv	Sbty	Stop	Fail	Total
FileServer	3	-	-	0	3
HBase RegionServer	1	-	-	0	1
NFS Gateway	1	-	-	0	1
HDFS	1	-	-	0	1
CLDB	1	-	-	0	1
TaskTracker	3	-	-	0	3
JobTracker	1	-	-	0	1
HostStats	3	-	-	0	3
HBase Master	1	-	-	0	1

	#	%	Total
Mounted	16	94%	23.3GB
Unmounted	1	6%	none
<b>Total</b>	<b>17</b>	<b>100%</b>	<b>23.2GB</b>

2. In the Navigation panel, select **NFS HA > NFS Setup**.  
The NFS Setup tab is displayed. The NFS Setup tab enables you to edit, remove or add VIPs in the Warehouse cluster.



3. On the NFS Setup tab, click the **Add VIP** button.



The **Add Virtual IP** dialog is displayed.

✕

### Add Virtual IP

▼ Virtual IP Range

\* Starting VIP:  ?

Ending VIP:  ?

\* Netmask:  ?

Preferred MAC address   ?

▼ Virtual IP Range

Use all network interfaces on all nodes that are running the NFS Gateway service.<br/>If additional NFS Gateway services are started, the network interfaces on their nodes will automatically become candidates for the VIPs in this range

Select the desired network interfaces:

Filter

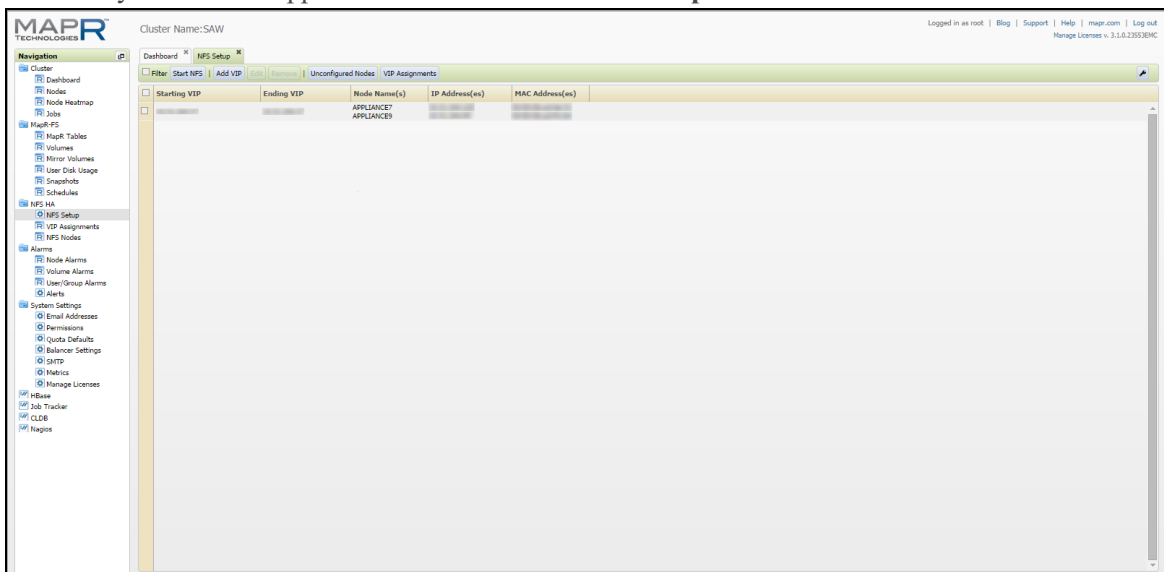
Node Name	IP Address	MAC Address	+
APPLIANCE7			Selected
APPLIANCE7	0.0.0.0		+
APPLIANCE9			Selected

<< < Showing 1-3 of 3 > >> 🔍

Node Name	IP Address	MAC Address	-
APPLIANCE7			-
APPLIANCE9			-


4. In the **Add Virtual IP** dialog, do the following:
  - a. In the **Starting VIP** field, type the starting IP Address for VIP.
  - b. In the **Ending VIP** field, type the ending IP Address for VIP. If this field is left blank, only one IP address is used for VIP allocation.
  - c. In the **Netmask** field, type the Netmask for the deployment.
  - d. Select **Select Desired Network Interfaces** to choose the available Network Interfaces that need to be used for VIP assignment. Select all of the external Interfaces from the list of available nodes by clicking the plus button next to the interface entry. Selected Interfaces will appear in the bottom list.
  - e. Click **OK** to add the VIP.

The newly added VIP appears in the list on the **NFS Setup** tab.



**Note:** VIP allocation can also be removed or edited from the **NFS HA > NFS Setup** tab by selecting a VIP and clicking the **Edit** or **Remove** button.

5. In the Navigation panel, select **NFS HA > VIP Assignment** to view the node that is assigned to the newly added VIP.



## Add a Virtual IP Address with Multiple Nodes (MapR UI)

You can add a virtual IP address (VIP) with multiple nodes. Virtual IP (VIP) is a technique used to load balance data access into HDFS by using a floating IP Address among the cluster nodes. This technique is mostly used by the MapR Hadoop Distribution along with the MapR-NFS Service. VIP can provide High Availability and Load Balancing by dynamically allocating the Floating IP among the nodes.

### Optimal VIP Configuration

We recommend using one VIP for every three Nodes, because the replication factor for HDFS is 3 by default. This also helps in optimizing the performance of the cluster.

In the case of High Data Load (>20K EPS), a single NFS might overload while replicating the file into the cluster. If the NFS Server crashes before the data is replicated, you may lose data.

Multiple NFS Servers also allow more distributed data locality which helps in High Availability and Fault Tolerance.

### Prerequisites

Calculate how many VIPs you can afford.

- We suggest **One VIP per 3 Nodes**.
- In case the number of nodes that you have is not a multiple of three, you can allocate multiple VIPs to more than three nodes. For example, two VIPs among five Nodes.

The steps to add the VIP are the same as adding any other VIP, but instead of choosing “all nodes” for VIP, you choose a subset of nodes to participate in the VIP.


- A node can participate in Multiple VIPs.
- For more information, see <http://doc.mapr.com/display/MapR/Setting+Up+VIPs+for+NFS>

## Optimal Configuration with the Warehouse Connector

The recommended configuration is to have one VIP per Warehouse Connector. In cases where Warehouse Connector numbers are higher than VIPs, configure multiple Warehouse Connectors to write to a VIP in a way so that traffic on VIPs can be normalized.

### Add a Virtual IP Address that has Multiple Nodes

1. Log on to the MapR Control System.
2. In the Navigation panel, select **NFS-HA > VIP Assignments**.
3. On the **NFS Setup** tab, click the **Add VIP** button.



4. In the **Add Virtual IP** dialog, do the following:

- a. Specify the Starting and Ending VIP as the same IP address.

**Add Virtual IP**

▼ Virtual IP Range

\* Starting VIP:  ?

Ending VIP:  ?

\* Netmask:  ?

Preferred MAC address   ?

Specify IP Address for VIP

▼ Virtual IP Range

Use all network interfaces on all nodes that are running the NFS Gateway service. <br/>If additional NFS Gateway services are started, the network interfaces on their nodes will automatically become candidates for the VIPs in this range

Select the desired network interfaces:

Filter

Node Name	IP Address	MAC Address	
saw-node1			+ Selected
saw-node1	0.0.0.0		+ Selected
saw-node1			
saw-node1	0.0.0.0		+ Selected

Click on + to add a NIC for VIP

Showing 1-4 of 4

Participating VIP will appear here

Node Name	IP Address	MAC Address	-
saw-node1			-

- b. Select **Select the Desired Network Interfaces** to choose the available Network Interfaces that need to be used for the VIP assignment. Select the NIC Cards that you want to participate in the VIP. A node can have multiple NICs, so depending on the Network Configuration you can select them.
- c. Click **OK** to add the VIP.

## Example VIP Configurations

The following table shows example configurations of virtual IP addresses (VIPs) with different numbers of nodes in the cluster.

Number of Nodes in Cluster	Number of VIPs
3 Nodes	1 VIP
5 Nodes	2 VIPs (3 Nodes each, 1 Common Node)
7 Nodes	2 VIPs (3 Nodes each, 1 Free Node)
8 Nodes	3 VIPs (3 Nodes each, 1 Common Node among 2 VIPs)
11 Nodes	4 VIPs (3 Nodes each, 1 Common Node among 2 VIPs)
11 Nodes	3 VIPs (3 Nodes each, 2 Free Nodes)

