

RSA[®] Security Analytics


MapR 4.1 Upgrade Instructions 4.1

Contents

| | |
|--|-----------|
| <u>Upgrade Security Analytics Warehouse Cluster (MapR 3.1 to MapR 4.1)</u> | <u>2</u> |
| <u>Upload MapR packages to SA Repository</u> | <u>3</u> |
| <u>Prerequisites</u> | <u>5</u> |
| <u>Upgrade the SAW Cluster</u> | <u>7</u> |
| <u>Configure MapReduce V2 Framework (yarn)</u> | <u>18</u> |

Upgrade Security Analytics Warehouse Cluster (MapR 3.1 to MapR 4.1)

Security Analytics Warehouse upgrade process includes upgrading MapR packages on the SAW Cluster Appliances. The upgrade package consisting of *sa-10.6-Mapr-4.1-Update-1-of-2-EL6.zip* and *sa-10.6-Mapr-4.1-Update-2-of-2-EL6.zip* are available on [SCOL](#).

| | |
|--|---|
| <p>NOTE</p>  | <p>Make sure you are using SA 10.6 or later versions to upgrade MapR 4.1.</p> |
|--|---|

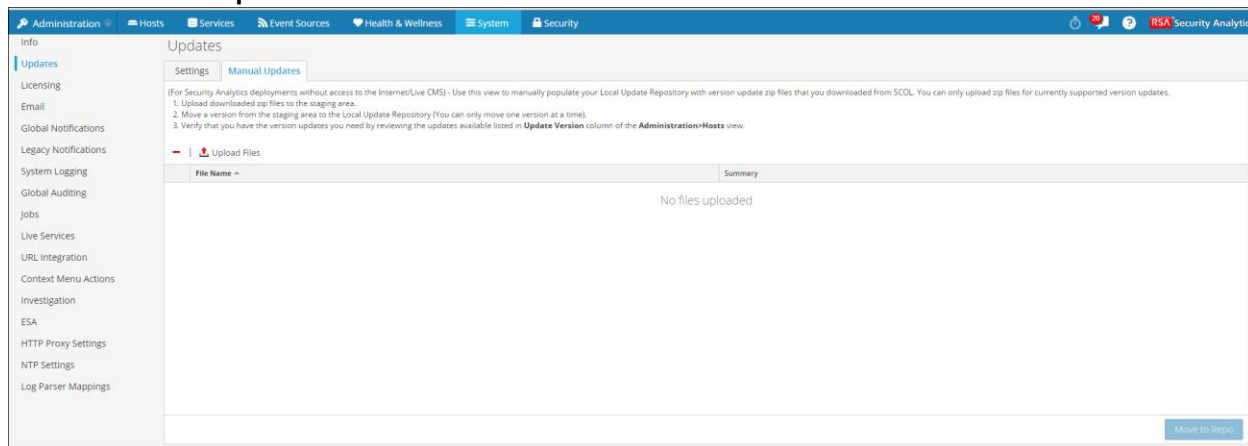
MapR 4.1 Upgrade Instructions



Upload MapR packages to SA Repository


Before you upgrade to MapR 4.1, you must upload the MapR packages to Security Analytics repository.

To upload the MapR packages to SA repository:

1. Log in to Security Analytics.
2. In the Security Analytics menu, select **Administration > System**.
3. In the left navigation panel, select **Updates**.
4. Click the **Manual Updates** tab.




5. Click  **Upload Files**
6. Click .
7. Select the .zip files to be uploaded.

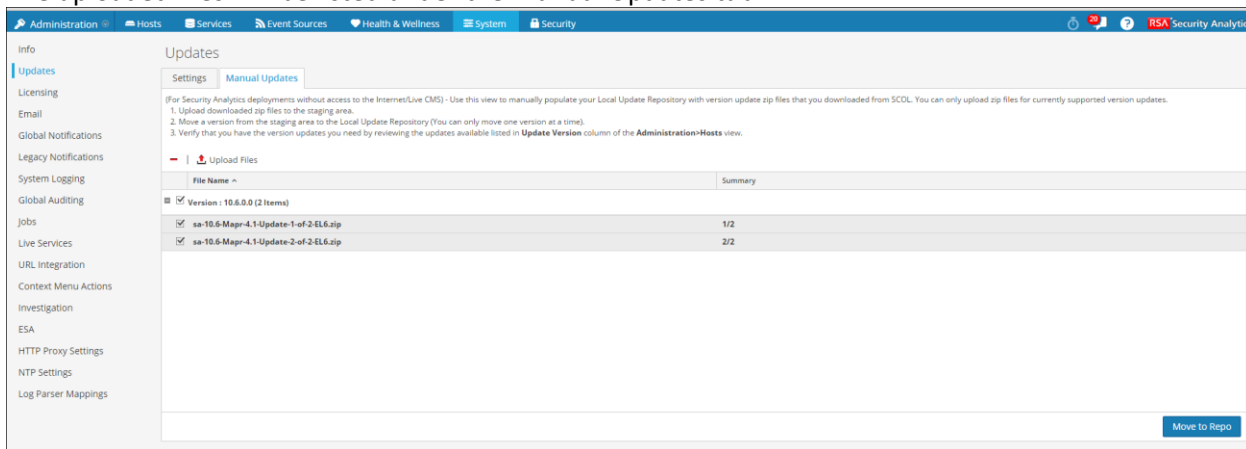


8. Click **Upload**.
The Uploading files view is displayed.


MapR 4.1 Upgrade Instructions



The uploaded files will be listed under the Manual Updates tab.



9. After the .zip files are uploaded, click **Move to Repo**.
The Creating Repository on SA service view is displayed.



After the repository is created on the SA server, the upload process is complete.


Prerequisites

Make sure that:

- You are logged into a MapR component.
- You cancel all the running SAW reports.
- There are no Jobs running on the SAW cluster. To stop all jobs, do the following:
 1. Log on to the Hosts on the SAW cluster.
 2. Enter the following command:


```
hadoop job -list
```

 The command returns the list of jobs that are running on the SAW cluster. Wait for the jobs to complete.

| | |
|---|--|
|  | <p>NOTE You can also forcibly kill the running hadoop jobs using the following command:</p> <pre>hadoop job -kill <job-id></pre> <p>Where <code>job-id</code> is the id of the running jobs that are displayed in step 2.</p> |
|---|--|

3. Run these commands:


```
mkdir /etc/yum.repos.d/repos.oem
mv /etc/yum.repos.d/Cent* /etc/yum.repos.d/repos.oem
```
4. If scp fails in the above step, create the **netwitness.repo** file (using any editor such as vi) with the following content or copy the file from any other device that has been added to SA Admin Devices:


```
[nwupdates]
name=Netwitness-Updates-Repo
baseurl=http://<SA_Server_IP>/rsa/updates/10.6.0
enabled=1
gpgcheck=0
sslverify=1
sslcacert=/etc/pki/CA/certs/RSACorpCAv2.pem
```
5. Copy the certificate from SA to MapR nodes:


```
scp root@<SAServer-IP>:/etc/pki/CA/certs/RSACorpCAv2.pem
/etc/pki/CA/certs/
```
6. Edit the following files with vi, add/edit the lines indicated:


```
vi /etc/yum.repos.d/netwitness.repo
    enabled=1
vi /etc/yum.conf
    http_caching=none
vi /etc/yum/pluginconf.d/fastestmirror.conf
    enabled=0
```
7. Remove all the yum caches from the Host. Enter the following command:


```
yum clean all
```

MapR 4.1 Upgrade Instructions

8. List all the available updates and ensure that the updates are valid. Enter the following command:

```
yum check-update
```

Upgrade the SAW Cluster

To upgrade the SAW Cluster:

1. On all the Hosts on the SAW cluster stop the HiveServer2 by using the command:

```
stop hive2
```
2. On all the Host where the service is running stop MapR-NFS service by using one of the following options:

- a) Use the command line interface by logging on to any one of the Hosts and enter the following command:

```
maprcli node services -nfs stop -filter "[hn==*]"
```

- b) Use the MapR Cluster System user interface, stop the MapR-NFS

Log on to the MapR Cluster System user interface. In a web browser, enter the following URL:

https://<node_ip_address>:8443

| | |
|--------------------|---|
| <p>NOTE</p> | <p>Start the MapR Cluster System user interface on any one of the Hosts in the cluster using the following command:</p> <pre>/opt/mapr/adminuiapp/webserver start</pre> |
|--------------------|---|

In the **Services** panel on the right side of the Dashboard, click **NFS Gateway**.

The screenshot shows the MapR Cluster System user interface for a cluster named 'saw'. The main area displays a 'Cluster Heatmap - 3 Nodes on 1 Racks' with 3 nodes on /data/default-rack (3 visible). An alarm is shown: 'Cluster License Near Expiration Alarm' raised 1m 51.7s ago, with a summary stating 'One or more licenses is about to expire within 21 days'. On the right side, the 'Services' panel is visible, listing various services and their status. The 'NFS Gateway' service is highlighted with a red box, showing 3 active instances and 0 standby, stop, or failed instances.

| Service | Actv | Stby | Stop | Fail | Total |
|--------------------|----------|----------|----------|----------|----------|
| FileServer | 3 | - | - | 0 | 3 |
| HBase RegionServer | 3 | - | - | 0 | 3 |
| NFS Gateway | 3 | - | - | 0 | 3 |
| WebServer | 1 | - | - | 0 | 1 |
| CLDB | 3 | - | - | 0 | 3 |
| TaskTracker | 3 | - | - | 0 | 3 |
| JobTracker | 1 | - | - | 0 | 1 |
| HostStats | 3 | - | - | 0 | 3 |
| HBase Master | 3 | - | - | 0 | 3 |

MapR 4.1 Upgrade Instructions

The Nodes tab is displayed with all the Hosts that are running the MapR-NFS service.

Cluster Name: saw

Dashboard × Nodes ×

Services ▾ Filter Properties Forget Node | Manage Services Change Topology

⊖ Configured Services ▾ is ▾ NFS Gateway ⊕

| Topology | Health | Hostname | Configured Services | Running Services | Physical Topology |
|---|--------------------------|-----------|--|---|------------------------------|
| Cluster data default-rack default-rack | <input type="checkbox"/> | saw-node1 | tasktracker hbmater hbregionserver webserver cdb fileserv nfs hoststats jobtracker | hbmater hbregionserver webserver cdb fileserv nfs hoststats | /data/default-rack/saw-node1 |
| | <input type="checkbox"/> | saw-node2 | tasktracker hbmater hbregionserver cdb fileserv nfs hoststats jobtracker | tasktracker hbmater hbregionserver cdb fileserv nfs hoststats jobtracker | /data/default-rack/saw-node2 |
| | <input type="checkbox"/> | saw-node3 | tasktracker hbmater hbregionserver cdb fileserv nfs hoststats jobtracker | tasktracker hbmater hbregionserver cdb fileserv nfs hoststats | /data/default-rack/saw-node3 |

Select the Hosts and click **Manage services**.

Cluster Name: saw

Dashboard × Nodes ×

Services ▾ Filter Properties Forget Node | Manage services for 3 nodes Change topology for 3 nodes

⊖ Configured Services ▾ is ▾ NFS Gateway ⊕

| Topology | Health | Hostname | Configured Services | Running Services | Physical Topology |
|---|-------------------------------------|-----------|--|---|------------------------------|
| Cluster data default-rack default-rack | <input checked="" type="checkbox"/> | saw-node1 | tasktracker hbmater hbregionserver webserver cdb fileserv nfs hoststats jobtracker | hbmater hbregionserver webserver cdb fileserv nfs hoststats | /data/default-rack/saw-node1 |
| | <input checked="" type="checkbox"/> | saw-node2 | tasktracker hbmater hbregionserver cdb fileserv nfs hoststats jobtracker | tasktracker hbmater hbregionserver cdb fileserv nfs hoststats jobtracker | /data/default-rack/saw-node2 |
| | <input checked="" type="checkbox"/> | saw-node3 | tasktracker hbmater hbregionserver cdb fileserv nfs hoststats jobtracker | tasktracker hbmater hbregionserver cdb fileserv nfs hoststats | /data/default-rack/saw-node3 |

The Manage Services view is displayed.

MapR 4.1 Upgrade Instructions

Manage Services ✕

| | |
|--------------------|-------------|
| CLDB | No change ▼ |
| FileServer | No change ▼ |
| JobTracker | No change ▼ |
| TaskTracker | No change ▼ |
| NFS Gateway | No change ▼ |
| HBase Master | No change ▼ |
| HBase RegionServer | No change ▼ |
| Webserver | No change ▼ |

Node(s) affected:
saw-node1, saw-node2, saw-node3
Note: In order to manage any other services not listed here, please do so on the given node properties page.

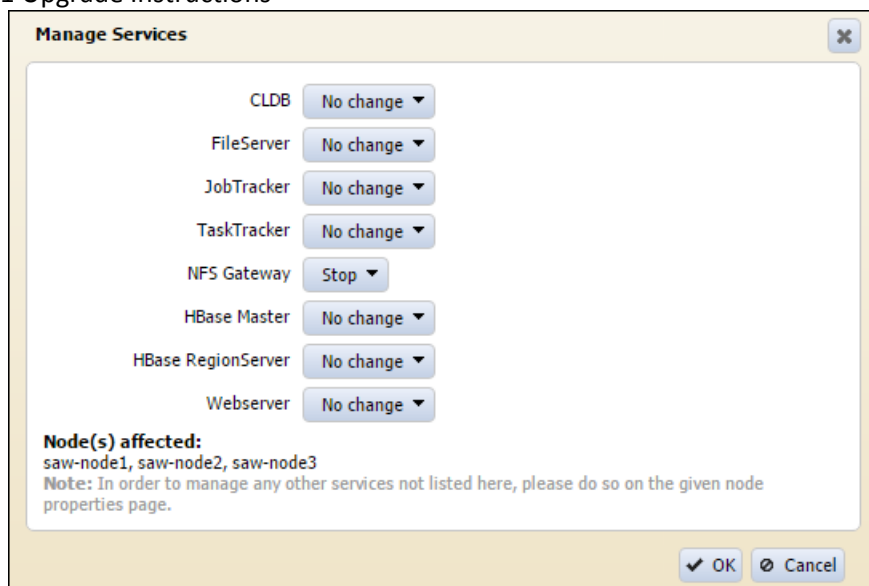
From the **NFS Gateway** drop-down menu, select **Stop**.

Manage Services ✕

| | |
|--------------------|-------------|
| CLDB | No change ▼ |
| FileServer | No change ▼ |
| JobTracker | No change ▼ |
| TaskTracker | No change ▼ |
| NFS Gateway | stop |
| HBase Master | Stop |
| HBase RegionServer | No change ▼ |
| Webserver | No change ▼ |

Node(s) affected:
saw-node1, saw-node2, saw-node3
Note: In order to manage any other services not listed here, please do so on the given node properties page.

MapR 4.1 Upgrade Instructions



3. Click **OK**.
4. Log on to any one of the host that has `nfs-utils` and `nfs-utils-lib` packages installed. By default, these package are installed on the following hosts:
 - Warehouse Connector
 - Log Decoder
 - Decoder
5. On the host, enter the following command for each SAW Appliance that has the NFS service configured.

```
showmount -e <IP_SAW_node>
```

where `IP_SAW_node` is the IP address of the SAW Appliance in the cluster.

This command should give the following message:

```
RPC: Program not registered
```

For example:

If you have a SAW cluster with three nodes configured with NFS service. You need to enter following commands:

```
showmount -e <SAW_node1_IP_Address>
```

```
showmount -e <SAW_node2_IP_Address>
```

```
showmount -e <SAW_node3_IP_Address>
```

These commands should display the following message for each SAW Appliance:

```
RPC: Program not registered
```

MapR 4.1 Upgrade Instructions

6. Stop the CLDB service on the SAW cluster. Perform the following
 - a. Log on to any one of the Hosts on the SAW cluster.
 - b. List all the Hosts that have CLDB service running on it by using the following command:

```
maprcli node listcl dbs
```
 - c. Locate the CLDB Master Appliance from the list of the Hosts running CLDB service by using the following command:

```
maprcli node cl dbmaster
```

NOTE



Please note down the CLDB Master Appliance hostname.

- d. Stop the CLDB service on the SAW cluster.
 You can stop the CLDB service using the **Command Line interface** or **MapR Cluster System user interface**.
 Perform any one of the following:
 - i. To stop the CLDB service on the SAW cluster using **Command Line interface**, perform the following:
 - a) Log on to the CLDB Master Appliance.
 - b) Stop the CLDB service on all the Hosts on the SAW cluster except the CLDB Master Appliance by using the following command:

```
maprcli node services -cl db stop -filter "[hn!=<cl dbmaster-hostname>]"
```

 where, <cl dbmaster-hostname> is the CLDB Master Appliance hostname.

NOTE



Make sure that you provide the correct CLDB Master Appliance hostname.

- c) Stop the CLDB service on CLDB Master Appliance by using the following command:

```
maprcli node services -cl db stop -filter "[hn==<cl dbmaster-hostname>]"
```

 Where, <cl dbmaster-hostname> is the CLDB Master Appliance hostname.

NOTE



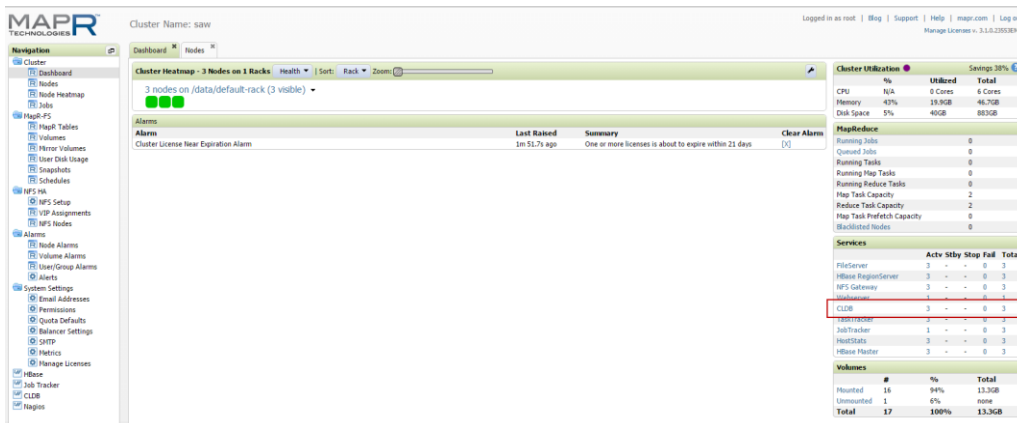
Make sure that you provide the correct CLDB Master Appliance hostname.

- ii. To stop the CLDB service on the SAW cluster using **MapR Cluster System user interface**, perform the following:
 - a) Log on to the MapR Cluster System user interface. In a web browser, enter the following URL:

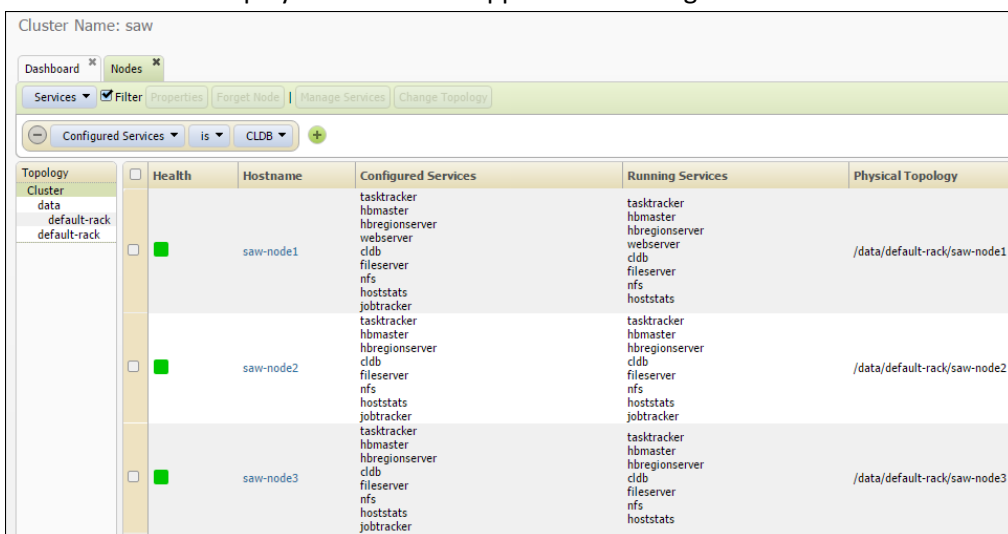
```
https://<node_ip_address>:8443
```

MapR 4.1 Upgrade Instructions

b) In the **Services** panel on the right side of the Dashboard, click **CLDB**.



The Nodes tab is displayed with all the Appliances running CLDB service.



c) Stop CLDB service on all the Hosts other than the CLDB Master Appliance.
Perform the following:

MapR 4.1 Upgrade Instructions

- i. Select the Hosts and click **Manage services**.

The Manage Services view is displayed.

Cluster Name: saw

Dashboard Nodes

Services Filter Properties Forget Node | Manage services for 3 nodes Change topology for 3 nodes

Configured Services is CLDB

| Topology | Health | Hostname | Configured Services | Running Services | Physical Topology |
|---|-------------------------------------|-----------|--|---|------------------------------|
| Cluster data default-rack default-rack | <input checked="" type="checkbox"/> | saw-node1 | tasktracker hmaster hregionserver webserver cldb fileserver nfs hoststats jobtracker | tasktracker hmaster hregionserver webserver cldb fileserver nfs hoststats | /data/default-rack/saw-node1 |
| | <input checked="" type="checkbox"/> | saw-node2 | tasktracker hmaster hregionserver cldb fileserver nfs hoststats jobtracker | tasktracker hmaster hregionserver cldb fileserver nfs hoststats jobtracker | /data/default-rack/saw-node2 |
| | <input checked="" type="checkbox"/> | saw-node3 | tasktracker hmaster hregionserver cldb fileserver nfs hoststats jobtracker | tasktracker hmaster hregionserver cldb fileserver nfs hoststats | /data/default-rack/saw-node3 |

Manage Services

CLDB

FileServer

JobTracker

TaskTracker

NFS Gateway

HBase Master

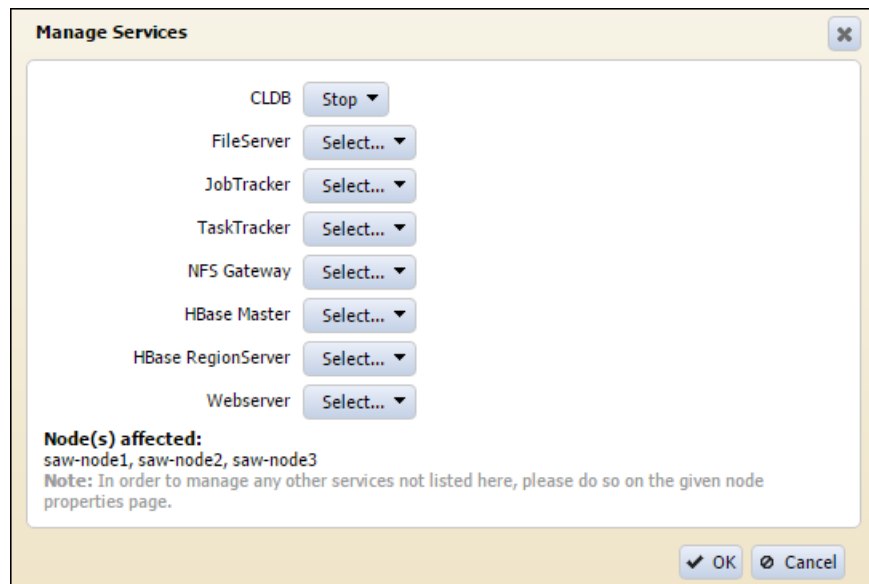
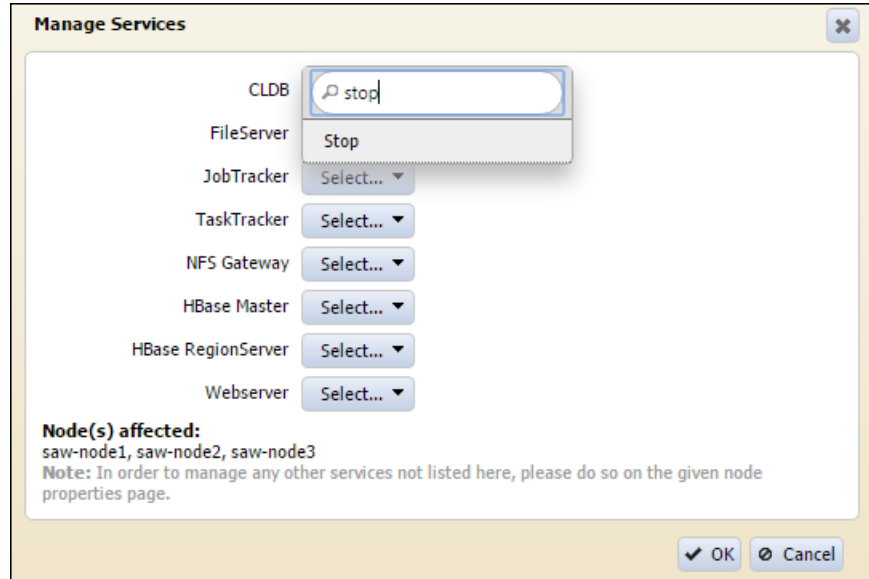
HBase RegionServer

Webserver

Node(s) affected:
saw-node1, saw-node2, saw-node3

Note: In order to manage any other services not listed here, please do so on the given node properties page.

- ii. From the **CLDB** drop-down menu, select **Stop**.



iii. Click **OK**.

d) Stop the CLDB service on the CLDB Master Appliance.

e) Stop the MapR Cluster System user interface on all the Hosts in the cluster by using the following command:

```
/opt/mapr/adminuiapp/webserver stop
```


7. Stop the warden service on all the Host on the SAW cluster in the following order:

- a. CLDB Master Appliance, which was noted in step 6.c
- b. Hosts with the CLDB service
- c. Host without the CLDB service


8. Use the following command to stop the warden service:

```
service mapr-warden stop
```

MapR 4.1 Upgrade Instructions


| | |
|--|---|
| NOTE  | The <code>mapr-warden</code> service may take approximately 10 minutes to stop. |
|--|---|

9. Stop the zookeeper service on all the Hosts on the SAW cluster by using the following command:
`service mapr-zookeeper stop`
10. On each Host on the SAW cluster. Perform the following:


| | |
|--|---|
| NOTE  | Make sure that you do not perform the following instructions simultaneously on multiple Hosts on the SAW cluster. |
|--|---|

Update the MapR packages and the SAW service on the Host using SA repo, using the following command:

```
yum clean all
yum update
```

| | |
|--|---|
| NOTE  | Make sure centos files do not exist in <code>/etc/yum.repos.d/</code> . Refer to step 3 in Prerequisites. |
|--|---|

- a) Set the ownership of all the files in the following folders to `mapr` user:
 - i. `/opt/mapr/pid` (Command: `chown -R mapr: /opt/mapr/pid`)
 - ii. `/opt/mapr/logs` (Command: `chown -R mapr: /opt/mapr/logs`)
 - iii. `/tmp` (Command: `chown -R mapr: /tmp`)
 - b) Set the permission for the `/opt/mapr/conf/daemon.conf` by using the following command:
`chmod 644 /opt/mapr/conf/daemon.conf`
 - c) Execute the following script:
`/opt/mapr/server/configure.sh -R`
 - d) Execute the following script:
`/opt/mapr/server/upgrade2maprexcute`
11. Start the zookeeper service on all the Hosts on the SAW cluster by using the following command:
`service mapr-zookeeper start`

| | |
|--|--|
| NOTE  | Verify the status of the zookeeper service in all the Host by using the following command: <code>service mapr-zookeeper qstatus</code> Make sure that the status is follower or leader . |
|--|--|

12. Start the warden service on all the Host on the SAW cluster in the following order:
 - a) CLDB Master Appliance (Use the hostname that you noted in step 6c)
 - b) Hosts with the CLDB service
 - c) Host without the CLDB service

Use the following command to start the warden service:
`service mapr-warden start`
13. On any one of the Hosts on the SAW cluster, verify the status of the SAW cluster using the following command:
`maprcli node cldbmaster`
14. Once the SAW cluster is up and running, perform the following on any one of the Hosts in the cluster:

MapR 4.1 Upgrade Instructions

- a) Get the MapR version by using the following command:

```
cat /opt/mapr/MapRBuildVersion
```

 The MapR version is displayed. Copy the MapR version number.
- b) Enter the following command:

```
maprcli config save -values {"mapr.targetversion":"<XXX>"}
```

 Replace <XXX> with the MapR version number that you copied in step b.

15. If HBASE is already installed:

- a) Execute `/opt/mapr/hbase/hbase-0.98.12/bin/hbase upgrade -check`
- b) Execute `/opt/mapr/hbase/hbase-0.98.12/bin/hbase upgrade -execute` on all the nodes.

16. To upgrade the schema on all nodes:

- a) Download the `db-derby-10.12.1.1-bin.zip` to `/root/` from SCOL.
- b) Unzip `db-derby-10.12.1.1-bin.zip`.
- c) Run the command to move the Metastore DB from root to MapR.

```
rm -rf /root/metastore_db
```
- d) Run the command to take a backup and change the MapR ownership.

```
cp -r /home/mapr/metastore_db/ /home/mapr/metastore_db_bk
chown mapr:mapr -R /home/mapr/metastore_db_bk
```
- e) Run the command to copy the original metastore from MapR user to root destination.

```
cp -r /home/mapr/metastore_db/ /root/metastore_db
```
- f) Run the command to set the `DERBY_HOME` path.
 - i. `export DERBY_HOME=/root/db-derby-10.12.1.1-bin`
 - ii. `export PATH="$DERBY_HOME/bin:$PATH"`
- g) Run the command to upgrade the metastore DB.

```
cd /opt/mapr/hive/hive-1.0/scripts/metastore/upgrade/derby
```
- h) Run the command to take a backup of the upgrade file.

```
cp upgrade-0.13.0-to-0.14.0.derby.sql upgrade-0.13.0-to-0.14.0.derby.sql_bkup
```
- i) Make sure to remove `'run 019-HIVE-7784.derby.sql'` from `upgrade-0.13.0-to-0.14.0.derby.sql`.
- j) Execute the command to log in to the `ij` prompt.

```
java -jar $DERBY_HOME/lib/derbyrun.jar ij
```
- k) Execute the command to connect to `jdbc`.

```
CONNECT 'jdbc:derby:/root/metastore_db';
```
- l) Execute the script.
 - i. `RUN 'upgrade-0.12.0-to-0.13.0.derby.sql';`
 - ii. `RUN 'upgrade-0.13.0-to-0.14.0.derby.sql';`
 - iii. `quit;`
- m) Run the command to remove the Metastore DB of MapR user.

```
rm -rf /home/mapr/metastore_db
```
- n) Run the command to copy the `metastore_db` to MapR user destination.

```
cp -r /root/metastore_db /home/mapr/metastore_db
```


MapR 4.1 Upgrade Instructions

- o) Run the command to change the ownership of metastore_db to MapR.
`chown mapr:mapr -R /home/mapr/metastore_db`

Configure MapReduce V2 Framework (yarn)

Security Analytics Warehouse upgrade process also includes configuring MapReduce V2 Framework (yarn).

To configure MapReduce V2 Framework:

1. On all the nodes of the MapReduce V2 Framework having job tracker service, enter the following command:


```
yum remove mapr-jobtracker
```
2. To remove the task tracker on all nodes by using the following command:


```
yum remove mapr-tasktracker
```
3. Install the mapr nodemanager on all the nodes:


```
yum install mapr-nodemanager
```
4. Install the mapr resource manager on **ONE NODE**:


```
yum install mapr-resourcemanager
```
5. If HBASE is already removed before upgrading to MapR 4.1. Then upgrade HBASE 0.98.12 version on all the nodes.
 - a) To upgrade HBASE: Copy the following updated HBASE rpm packages to the directory you created (for example, /root/hbase).


```
mapr-hbase-0.98.12.201507081709-1.noarch.rpm
mapr-hbase-internal-0.98.12.201507081709-1.noarch.rpm
mapr-hbase-master-0.98.12.201507081709-1.noarch.rpm
mapr-hbase-regionserver-0.98.12.201507081709-1.noarch.rpm
```
 - b) Move to the directory containing the updated packages with the following command


```
cd /root/hbase/
```
 - c) Upgrade the hbase rpm packages by running the following command as root with the following command


```
rpm -ivh *.rpm
```
 - d) Verify that the upgrade was successful with the following command string.


```
rpm -qa | grep mapr-hbase
```
 - e) The following information is displayed during the verification:


```
mapr-hbase-internal-0.98.12.201507081709-1.noarch
mapr-hbase-0.98.12.201507081709-1.noarch
mapr-hbase-master-0.98.12.201507081709-1.noarch
mapr-hbase-regionserver-0.98.12.201507081709-1.noarch
```
 - f) Execute /opt/mapr/hbase/hbase-0.98.12/bin/hbase upgrade -execute on all the nodes.
6. Execute the following script on all nodes:


```
/opt/mapr/server/configure.sh -C <saw-node1 IP>,<saw-node2 IP>,<saw-node3 IP> -Z <saw-node1 IP>,<saw-node2 IP>,<saw-node3 IP> -RM <saw-node IP> -N <cluster-name>
```

Where,

 - C is the cldb_list
 - Z is the zookeeper_list
 - RM <saw-node IP>, is the mapr-resource manager that was installed in step 4.
7. Execute this command, only if you have installed HBase 0.98 separately as explained in step 5.
 - a)

```
hadoop fs -rmr /hbase/hbase.id
```
 - b)

```
maprcli node services -hbregionserver start -filter "[hn==*]"
```
8. On any one of the Hosts on the MapR V2 Framework, verify the status of the SAW cluster:


```
maprcli node cldbmaster
```

MapR 4.1 Upgrade Instructions

- Restart the warden first on the node running clubmaster followed by the other CLDB nodes and non-CLDB nodes.

```
service mapr-warden restart
```

- Edit the following and set the default mode to yarn.

```
Edit "/opt/mapr/conf/hadoop_version"
```

```
Set "default_mode" to "yarn"
```

- Execute the following command.

```
export MAPR_MAPREDUCE_MODE=yarn
```

- In `/opt/mapr/hadoop/hadoop-2.5.1/etc/hadoop/core-site.xml` file, add the following configuration in all the nodes.

```
<property>
```

```
<name>hadoop.proxyuser.mapr.groups</name>
```

```
<value>*</value>
```

```
<description>Allow the superuser mapr to impersonate any member of any group</description>
```

```
</property>
```

```
<property>
```

```
<name>hadoop.proxyuser.mapr.hosts</name>
```

```
<value>*</value>
```

```
<description>The superuser can connect from any host to impersonate a user</description>
```

```
</property>
```

- Edit the host iptables (`/etc/sysconfig/iptables`) on all the nodes of the cluster and add the port numbers (19890, 8030, 8032, 8033, 8040, 8041, 8088, 8042, 8044, 8090, 19888, and 8031).

The rule syntax is:

```
-A INPUT -m state --state NEW -m tcp -p tcp --dport <port no> -j ACCEPT
```

For example:

```
Generated by iptables-save v1.4.7 on Fri Apr 25 06:22:43 2014
*filter
:INPUT ACCEPT [0:0]
:FORWARD DROP [0:0]
:OUTPUT ACCEPT [0:0]
-A INPUT -i lo -j ACCEPT
-A INPUT -p icmp -j ACCEPT
-A INPUT -m state --state RELATED,ESTABLISHED -j ACCEPT
-A INPUT -p tcp -m state --state NEW -m tcp --dport 22 -j ACCEPT
-A INPUT -p tcp -m state --state NEW -m tcp --dport 50020 -j ACCEPT
-A INPUT -p tcp -m state --state NEW -m tcp --dport 60010 -j ACCEPT
-A INPUT -p tcp -m state --state NEW -m tcp --dport 60020 -j ACCEPT
-A INPUT -p tcp -m state --state NEW -m tcp --dport 20000 -j ACCEPT
-A INPUT -p tcp -m state --state NEW -m tcp --dport 50010 -j ACCEPT
-A INPUT -p tcp -m state --state NEW -m tcp --dport 8443 -j ACCEPT
-A INPUT -j REJECT --reject-with icmp-host-prohibited
COMMIT
```

NOTE



Make sure you add port range 32768:61000 for private or public networks.

If your MapR cluster is configured through private network, do one of the following:
For example (Private network):

- A INPUT -i eth1 -j ACCEPT

Where eth1 is an interface used for PRIVATE network.

MapR 4.1 Upgrade Instructions

- -A INPUT -m state --state NEW -p tcp -m iprange --src-range 192.167.10.1-192.167.10.3 --dport 32768:61000 -j ACCEPT
Where --src range is the range of the IP address of the SAW setup which is in a PRIVATE network.

If your MapR cluster is configured through public network, do one of the following:
For example (Public network):

- -A INPUT -m state --state NEW -p tcp -m iprange --src-range 10.31.246.73-10.31.246.75 --dport 32768:61000 -j ACCEPT
Where --src range is the range of the IP address of the SAW setup which is in a PUBLIC network.
- -A INPUT -m state --state NEW -p tcp -s 10.31.246.73,10.31.246.74,10.31.246.75 --dport 32768:61000 -j ACCEPT
Where -s is the range of the IP address (comma separated) of the SAW setup which is in a PUBLIC network.

14. Start hiveserver2.

```
start hive2
```

15. Set the permissions on the /opt folder:

```
chmod 755 /opt
```

16. After the upgrade process is complete, reboot all the nodes.

17. To upgrade to Hive 1.0 in Reporting Engine:

- Remove and take a backup of “hive-jdbc-0.12.0-with-full-dependencies.jar” available at /opt/rsa/soc/reporting-engine/plugins/hive-jdbc-0.12.0-with-full-dependencies.jar.
- Copy the /opt/mapr/hive/hive-1.0/lib/hive-jdbc-1.0.0-mapr-1508-standalone.jar from MapR node to Reporting Engine’s host /opt/rsa/soc/reporting-engine/plugins.

OR

You can also download it from:

<http://repository.mapr.com/nexus/content/groups/mapr-public/org/apache/hive/hive-jdbc/1.0.0-mapr-1508/hive-jdbc-1.0.0-mapr-1508-standalone.jar>

- Restart Reporting Engine.